# PDMDB Motherboard

P.J. Garsed & S. Wotton

The RICH PDMDB is the on-detector electronics module that is the interface between the CLARO front-end ASICs and the LHCb common readout environment. The module is implemented using FPGAs and the Versatile Link chipset and uses FEASTMP DCDC converters to locally regulate the module power.

There are in fact two variants of the PDMDB, one (PDMDB-R) for the RICH1 and RICH2 detector regions instrumented with 1 inch MAPMTs and another (PDMDB-H) for the parts of RICH2 instrumented with 2 inch MAPMTs. The two variants are functionally almost identical except that the PDMDB-H has 256 instead of 512 input channels and these are distributed across 2 FPGAs and 2 PDMDB-DTMs instead of 3 FPGAs and 3 PDMDB-DTMs in the case of PDMDB-R. In this document these variants will both be referred to as PDMDB except where necessary to draw a distinction.
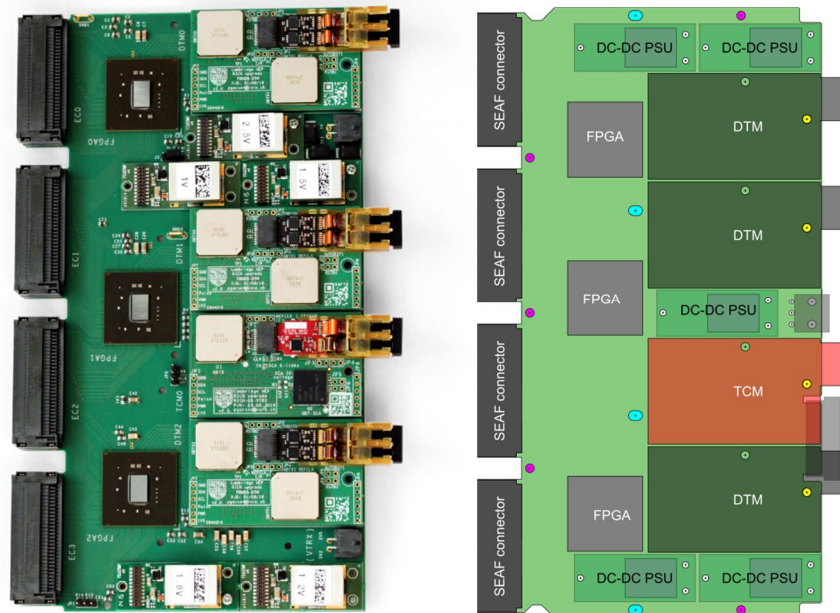
Figure 1 A fully equipped R-type PDMDB module. The photo (left) is in the prototype configuration. The sketch (right) shows the final symmetrised layout.

1   FPGA choice

The main active components on the PDMDB are the FPGAs. They contribute significantly to the cost of the system. Therefore, the number required per PDMDB and the specific part are optimised to minimise costs.

The xc7k70t676 part is the smallest FPGA in the Xilinx Kintex 7 series with sufficient IOs and IO banks to provide the necessary connectivity to the single-ended 2.5V CMOS CLARO signals and the LVDS/SLVS and single-ended 1.5V CMOS signals of the GBTX ASIC. There is a penalty associated with

selecting larger FPGAs, as the increased cost reflects the additional high performance features embedded in the FPGA which are not used in this application. The increased physical size would also require more routing layers on the PCB, which would also increase costs. With this choice of FPGA series, it appears that 3 smaller FPGAs are more cost-efficient than two larger FPGAs.

2   Master link (TCM)

The bi-directional master link is implemented by the PDMDB-TCM plug-in using Versatile Link components (GBTX, GBTX-SCA and VTRX). It provides a configuration and monitoring interface for the FE electronics and a fast command path to control and synchronise the data acquisition. All the connections from the TCM GBTX to the motherboard are through a 20×10pin Megarray connector (FCI-84516-002LF).

The GBTX-SCA ASIC provides the interfaces required for the control and monitoring towards the FE and is connected to the LHCb controls network via the master GBTX ASIC using its dedicated EC port.

The master GBTX ASIC is programmed in *transceiver mode* (i.e. bi-directional) with *forward error correction.* The downlinks operate at 80Mbit/s. Two such e-links are routed to each FPGA which allows up to 4 synchronous bits to be sent to each FPGA per LHC bunch-crossing. If required, the design allows for these links to be operated at up to 320Mbit/s without hardware changes. The uplinks (other than the SC link) are not used in the LHCb upgrade architecture.

The master GBTX uses its integrated crystal oscillator for power-up sequencing and monitoring functions before the master link is established. The master GBTX e-fuses will be programmed during production testing to allow automatic establishment of the master link during power-up.

No external programmable reset is provided. Reset of the master links will require power cycling of the entire column.

The motherboard is designed to allow initial configuration of the master GBTX through its I2C interface for prototyping, production testing and maintenance. Once installed, access to the master GBTX registers is possible only through the GBTX IC channel (i.e. through the master link).

3   Clocks

The PDMDB-TCM is the master clock source for the PDMDB. The PDMDB-TCM GBTX recovers the LHC clock and GBTX phase shifter clock outputs are routed to the PDMDB-DTMs and FPGAs using 100 Ohm-nominal differential traces. An additional clock from the PDMDB-DTM is routed to its corresponding FPGA. This clock topology is shown in Figure 2 for one of the DTMs. The 40MHz clock to each FPGA allows to resolve the phase ambiguity between the 160MHz data transmission clock and the 40MHz system reference clock.
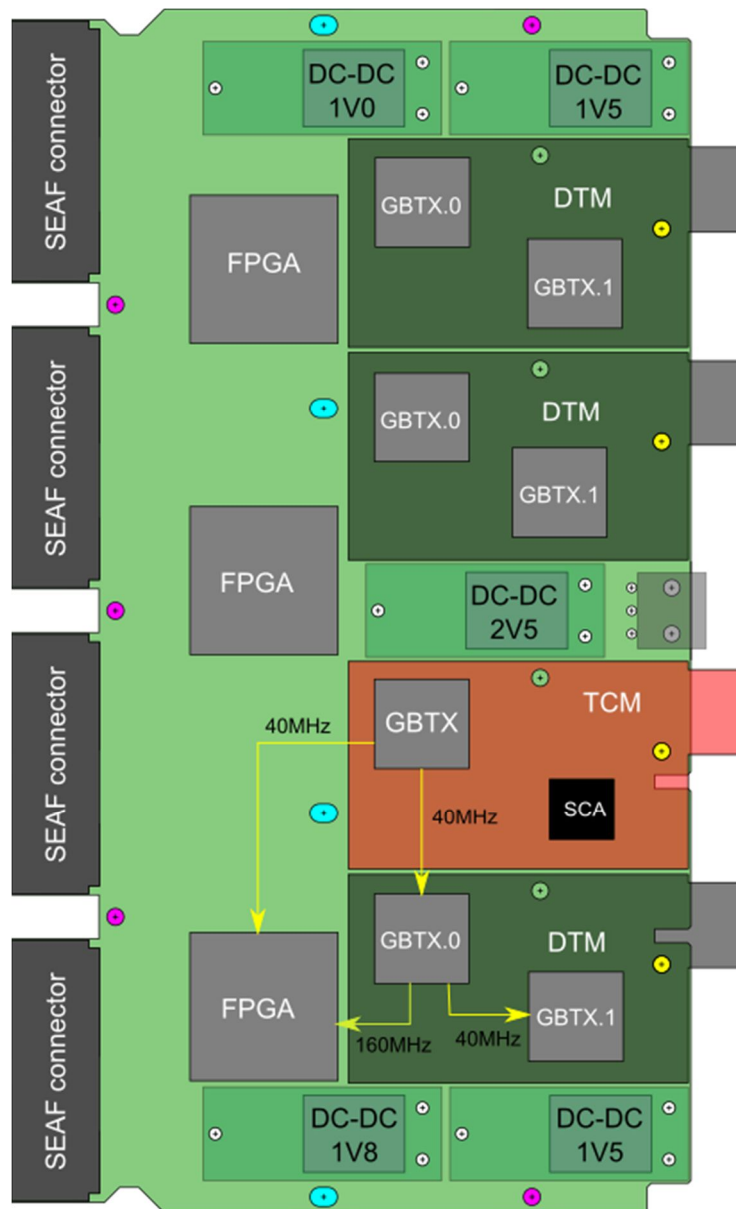
Figure 2 Clock distribution scheme

Waveforms for the principal clocks are shown in Figure 3. The 40 MHz GBTX.0 reference clock is driven from a GBTX phase-shifter output clock on the TCM. This SLVS signal is routed from the TCM GBTX pins through the TCM Megarray connector, across the motherboard, through the DTM Megarray connector and finally to the DTM GBTX.0 reference clock input pins which have an internal 100 Ohm differential termination. The clock trace is a point-to-point connection. The 40 MHz FPGA reference clock is also sourced by a TCM GBTX phase shifter output clock and is similarly routed through the TCM Megarray connector to the PDMDB motherboard but then goes to FPGA IO pins with internal 100 Ohm differential termination. The FPGA reference clock input is configured as LVDS in one of the HP (High Performance) IO banks powered at 1.8 V. In order to reduce the routing complexity and to avoid multi-drop loading, the reference clock of GBTX.1 on each DTM is driven by a phase-shifter output clock of GBTX.0. DTM GBTX.0 also drives a second 160MHz IO clock that is used in the FPGA in the serialisation of the data onto the data transmission e-links. The assignment of the master GBTX clock outputs is as follows:

Clock source	Clock destination
0	DTM2.GBTX0 40MHz reference clock
1	FPGA2 40MHz (aux) clock
2	DTM1.GBTX0 40MHz reference clock
3	FPGA1 40MHz (aux) clock
4	DTM0.GBTX0 40MHz reference clock
5	FPGA0 40MHz (aux) clock

Table 1 Master GBTX clock assignments
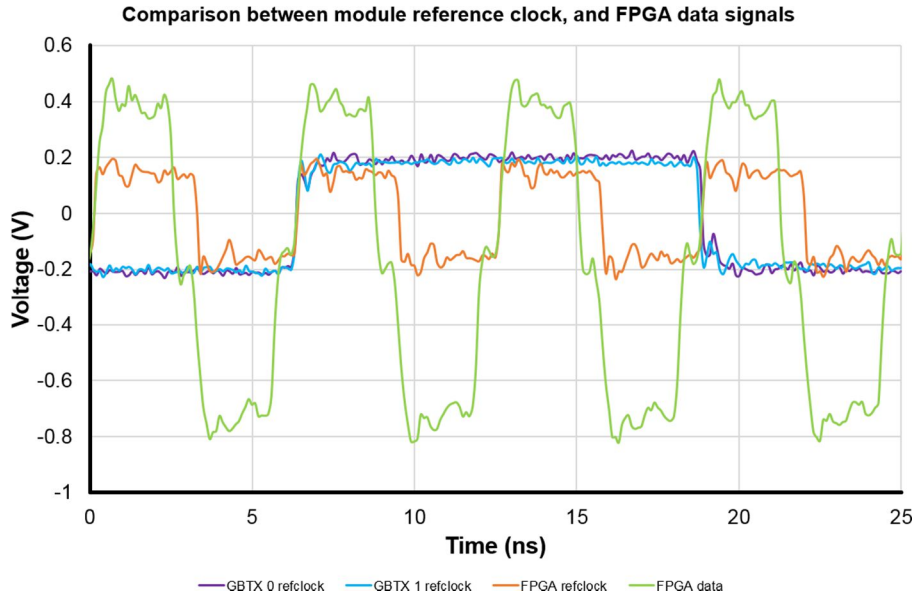
Figure 3 Measured reference clock waveforms

A comparison of the cycle-to-cycle jitter is shown in Figure 4. The GBTX.0 reference clock shows worse jitter performance than GBTX.1 which is likely due to the filtering effect of the GBTX reference clock PLL.
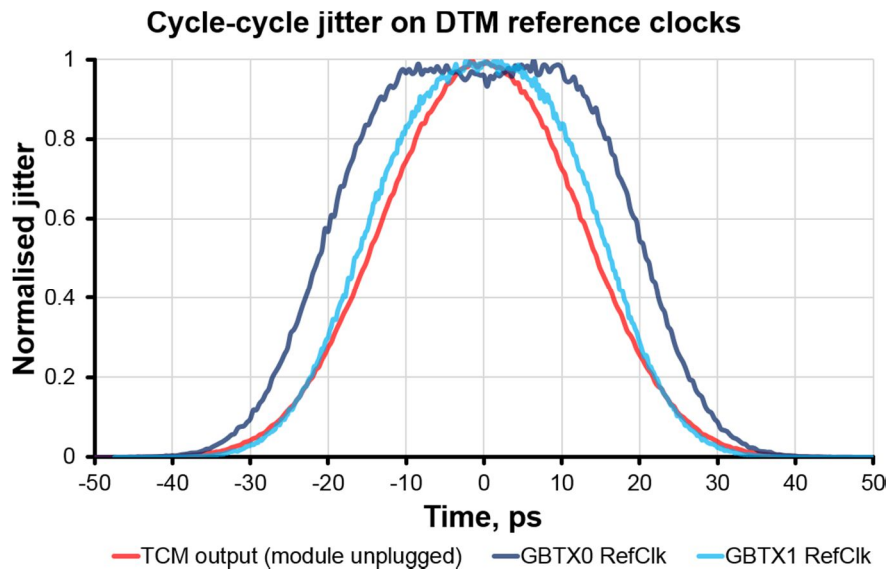
Figure 4 Measured cycle-to-cycle clock jitter

4   EC connections

All connections to and from the EC are routed through the four 30×6pin Samtec Searray connectors (SEAF-30-01-L-06-2-RA).

Control and monitoring signals are connected directly to the GBT-SCA ASIC located on the PDMDB-TCM plug-in. The possibility of configuration and monitoring via the FPGAs was rejected since the FPGAs are expected to be susceptible to radiation-induced upsets and a fail-safe configuration and monitoring interface at the front-end was preferred.

PMT outputs from the CLARO ASICs are point-to-point connections to FPGA inputs.

4.1   EC power

Each SEAF connector has 10 power return pins connected to the PDMDB ground planes and 4 power source pins connected to the 2.5V DCDC that also powers the FPGA IO banks used for the EC signals.

4.2   CLARO PMT signals

The PDMDB receives 512 (4 ECx2 MAPMTx64 channel) digital single-ended 2.5V CMOS signals from the 4 connected elementary cells.

To minimize component costs taking account of the number of FPGA IOs available and the bandwidth and modularity of the output links, the signals are distributed across 3 FPGAs, each driving two data links. Each data link therefore transmits 85 or 86 bits of PMT anode data per GBTX frame leaving 26 bits available in GBTX wide-frame mode for header data.

The CLARO ASIC channels are independent and asynchronous and the digital outputs are not clocked. The digital signals that arrive at the FPGA will therefore have a latency relative to the interaction time that depends on a number of factors: particle time-of-flight, photon path in the detector volume, time response of the MAPMT, time-walk in the front-end ASIC and propagation delays through the electronic chain. The effect of these is that the 512 signals at each PDMDB will arrive with a characteristic average latency and some spread about the mean value. Since each PDMDB receives signals from a small and localised area of the image plane, path length variation amongst these tends to be small. Any variation across the detector plane can then be compensated by clock phase adjustment per PDMDB (or even per FPGA).

The electrical characteristics of the FPGA IOs might be expected to have some effect on the photon detection efficiency since, for very small MAPMT signals, the digital outputs from the CLARO may not satisfy the set-up and/or hold time specification of the FPGA IOs. Although not separately measured, tests with beam demonstrate that this effect combined with the effect of the CLARO threshold does not have an unacceptable impact on the overall photon detection efficiency.

The routing of the PMT signals uses the top layer and two internal layers of the PCB stack-up (see Figure 13). The positioning of the FPGAs and the DTMs is mainly dictated by the aim to keep these traces as short as reasonably possible but is also constrained by the position of the mechanical mounting holes.

4.3 CLARO configuration

The CLARO ASIC has a number of features that must be programmed for correct operation. Each CLARO has a single 128-bit configuration shift register that is accessible through an SPI interface. Additional chip-select signals are used to identify the addressed CLARO.

The role of the PDMDB is therefore to route the configuration signals from the GBT-SCA chip to the EC SEAF connectors. The GBT-SCA has a dedicated SPI port which includes SCLK, MOSI, MISO and a set of 8 slave-select IO's. A few additional GBT-SCA GPIOs are used to further extend the slave-select address space so that every connected CLARO can be individually selected. There are no active components between the GBT-SCA and CLARO SPI interface. A sketch of the signal topology is shown in Figure 5.
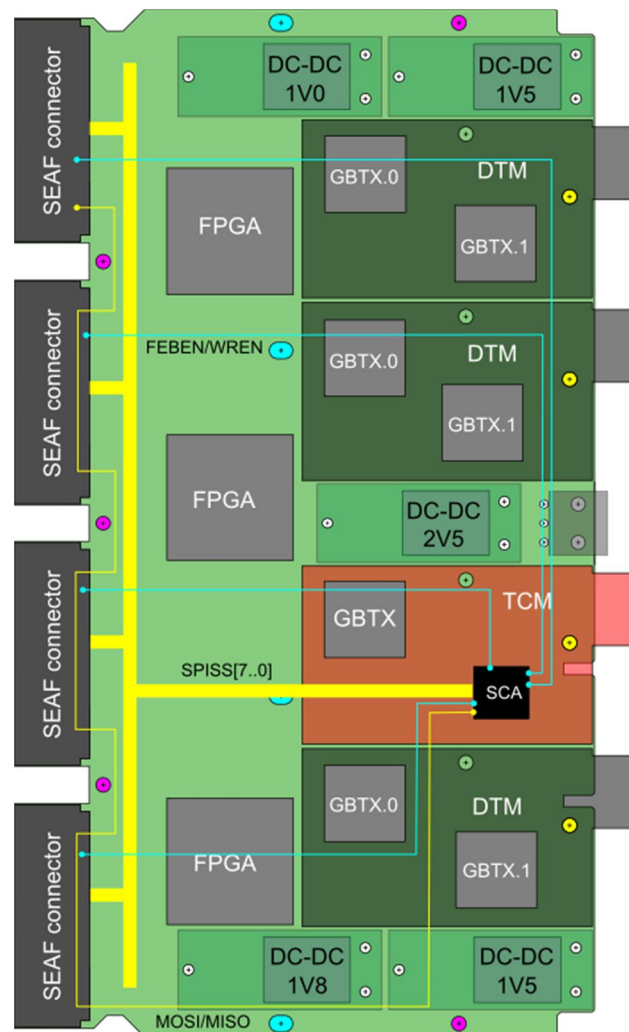

Figure 5 CLARO SPI signals. SCA SPI signals in yellow. Additional control signals in cyan. SPICK network not illustrated.

4.3.1 Clock signal

The path lengths of the signals are relatively long and several of them have multiple loads, being connected to several CLAROs. Lab measurements have been made to verify the signal integrity in a mock-up using a GBT-SCA and multiple CLARO chips. These measurements indicate that the signal integrity is satisfactory at clock speeds up to about 5MHz. However, to further optimise the clock

scheme, the latest PDMDB prototype implements an SPI clock network that allows different configurations to be studied. These include the original scheme of a single trace with short stubs at each load, an alternative 4-fold passive fan-out and an option to buffer the fanned out clocks using FPGA IOs.

For the passive fan-out, on leaving the TCM, the SPI clock signal is immediately fanned out to four separate transmission lines, with a series termination resistor on each, as close to the TCM as possible. The fanned-out traces are routed past optional FPGA buffers to each elementary cell. The use of FPGA buffers introduces the risk that a radiation-induced upset in the FPGA might result in uncontrolled behaviour of the clock signals although this is mitigated by the fact that the SPI interface of the CLARO ASICs is only active for a fraction of the time. The implemented clock scheme is illustrated in Figure 6
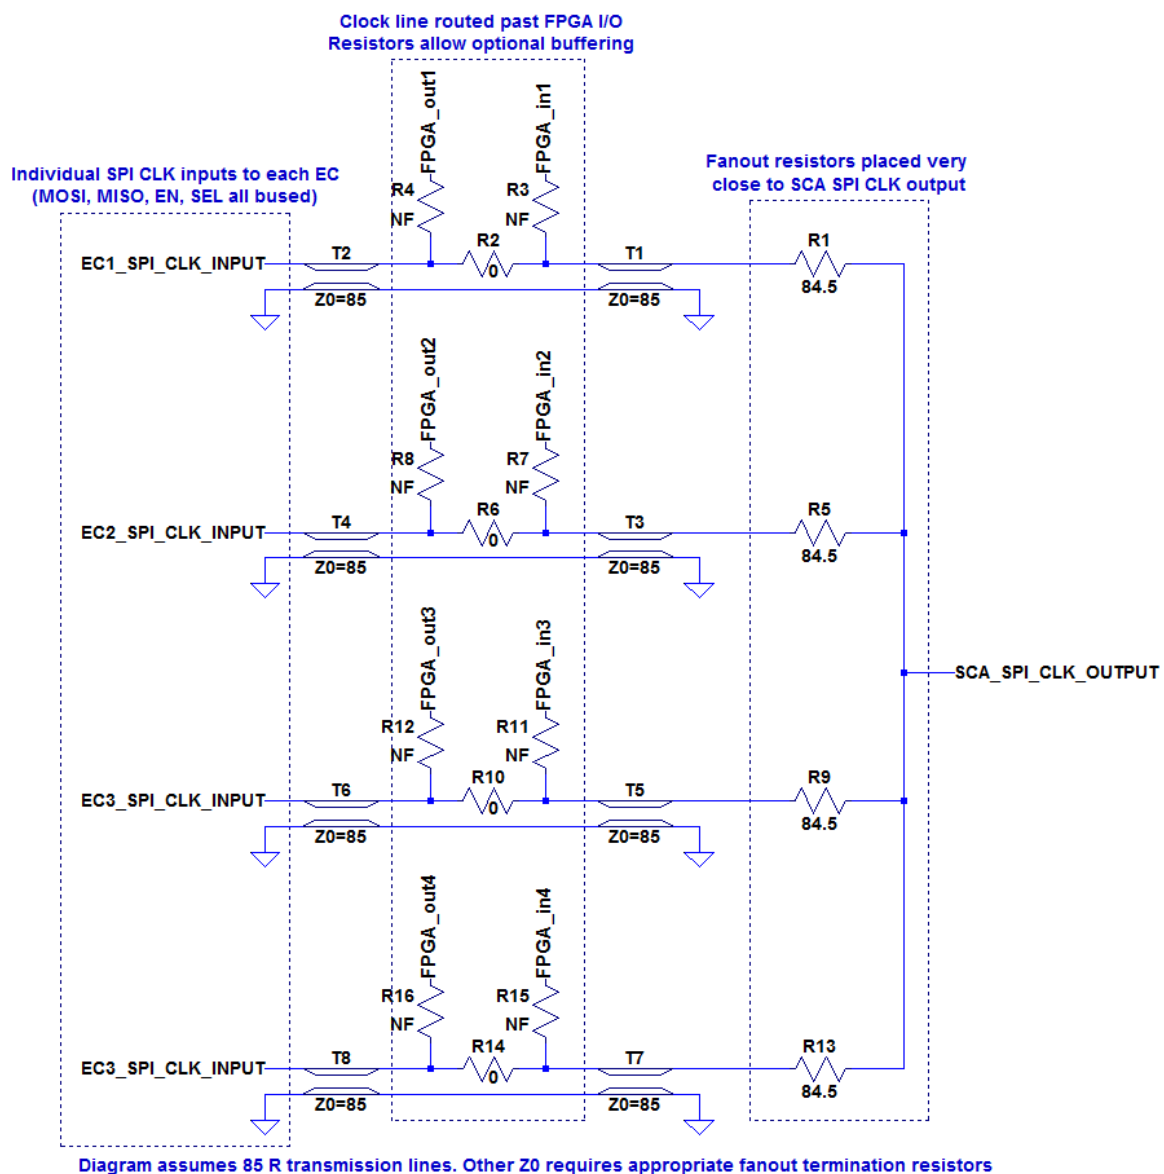
Figure 6 SPI clock distribution network

The resulting measured clock signals can be compared in Figure 7 and Figure 8. The resistance values in the figures refer to the backboard series resistance (R17 or R37).
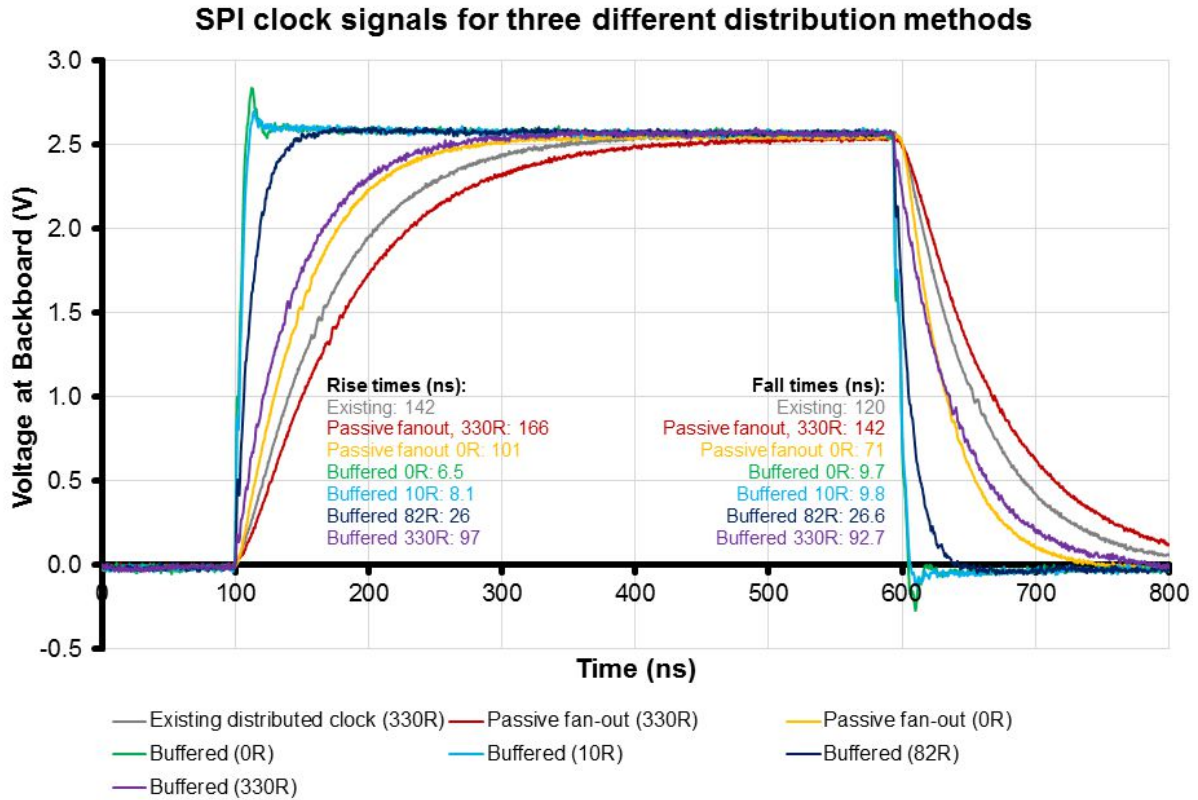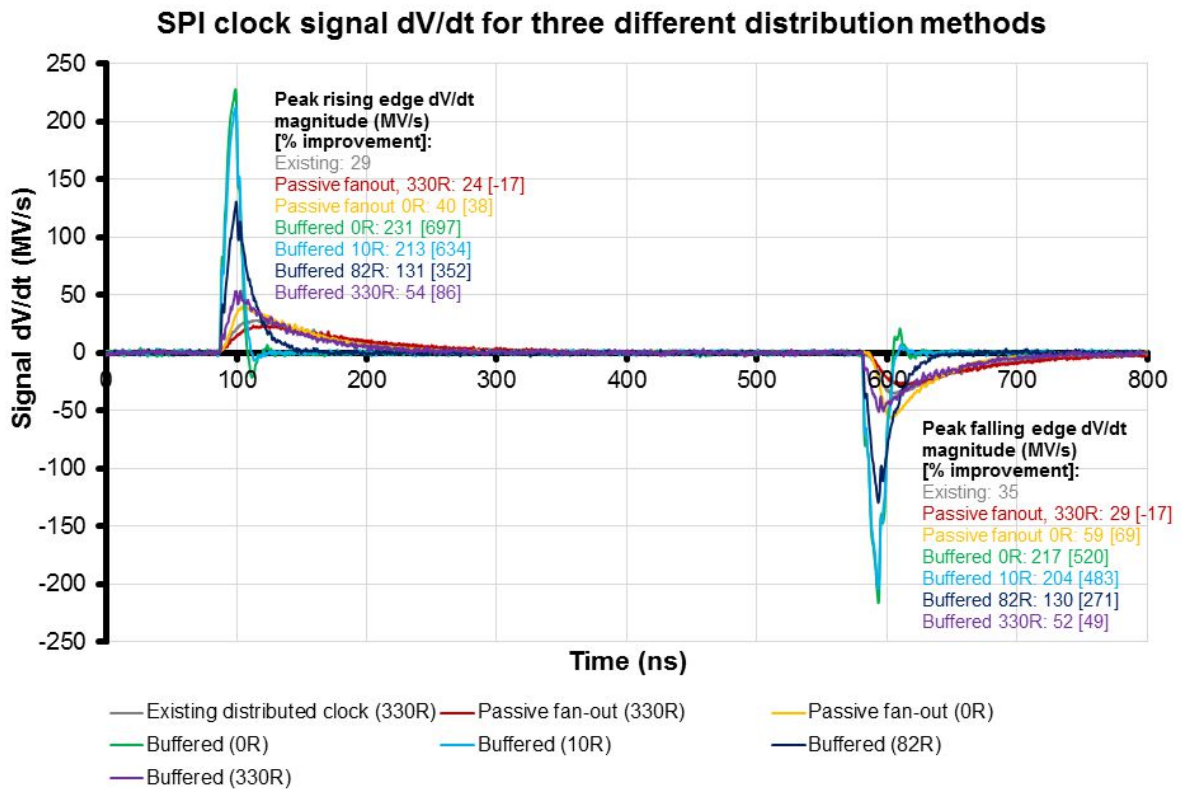
Figure 7

Figure 8

4.3.2   SPI timing

The CLARO reacts only to the rising edge of the SPI clock. The first data bit is driven by an internal pseudo-clock generated by the chip select signal before the first SPICK rising edge. The GBT-SCA

allows to separately choose which clock edge samples MISO or drives MOSI. The clock idle state can also be chosen to be either high or low. Empirically, the clock idle high mode does not appear to be compatible with the CLARO. Figure 9 shows the MISO signal (green) measured near the CLAROs at position EC3.
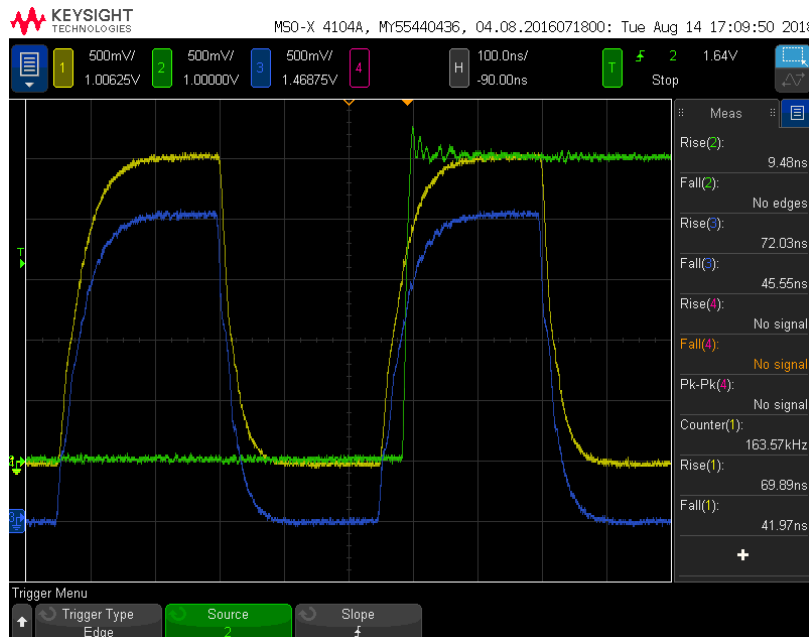

Figure 9 MISO timing

The yellow trace is the SPI clock measured at the backboard series resistor and the blue trace is the same clock near the SCA chip. These two measurements are nearly indistinguishable on this timescale so are slightly offset vertically for clarity. The SPI internal clock rises not later than the beginning of the rise of the external clock signal and MISO transitions roughly 30ns after this. It is therefore optimal for MISO to be sampled in the SCA on the rising edge of the clock since this maximises the settling time.

The corresponding MOSI trace is shown in Figure 10. In this case it is the timing of the yellow clock trace relative to the green trace that is important.
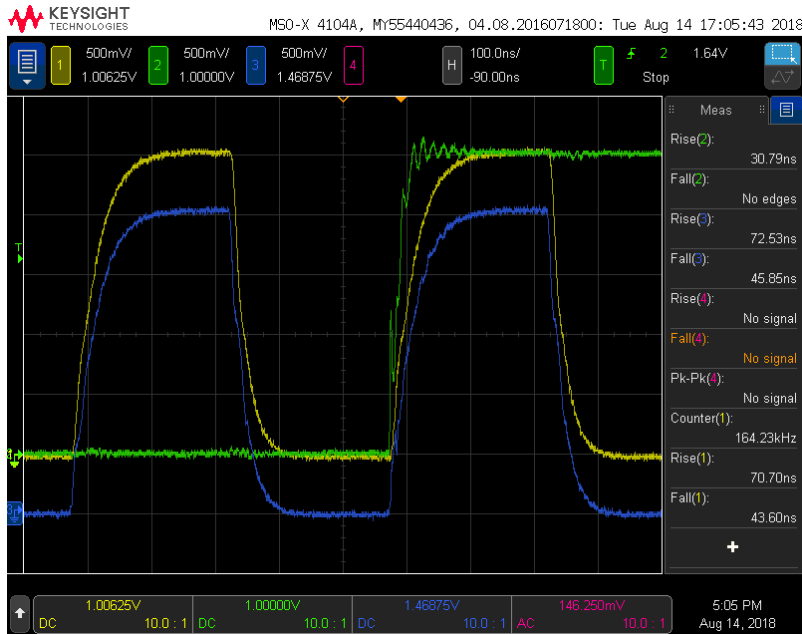
Figure 10 MOSI timing

Because the rising edge of the clock is slower than the MOSI transition, it is also apparently safe to drive MOSI on the rising edge of the clock. However, the MOSI edge is less clean so it may be more conservative to drive MOSI on the falling edge to centre the clock edge in the data at the CLARO. Both possibilities have been tested in digital tests without errors.

4.3.3 Digital tests

Digital tests have been done for some configurations (Table 2). Each test cycle corresponds to programming 16 CLAROs on two FEBs (2048 bits), reading back the registers (2048 bits) and comparing. Each cycle therefore corresponds to transferring 4096 bits. The programmed pattern changes from cycle-to-cycle but is not random (or even pseudo-random). No bit errors have been observed in any digital tests in the "reasonable" configurations. For most of these tests the SPI clock frequency was 1MHz. Larger resistors give slower edges but clock speeds up to about 5MHz are possible in some configurations (e.g. R1=0R, BB=10R).

R1,R5,R9,R13	R2,R6,R10,R14	R3,R4,R7,R8,R11 ,R12,R15,R16	BB	Cycles without errors	Clock
82R	0R	DNF	10R	>20000	1MHz
0R	0R	DNF	330R	>240000	1MHz
0R	0R	DNF	10R	>20000	1MHz
0R	0R	DNF	10R	>800000	5MHz

Table 2 Digital test results for tree topology

The clock "tree" topology is now preferred over the original "snake" topology. It gives faster clean edges and helps to isolate the CLAROs on different ECs. This may be important because there is some evidence that dynamic effects in the CLARO SPI clock IO inject glitches back into the clock trace.

4.4   Test pulse

The ability to inject a test charge at the CLARO input is an important feature that allows the read-out chain to be tested *in situ* in the absence of HV or Cherenkov photons and also provides a way of

monitoring the response of the system to a known stimulus. The test pulse feature of the CLARO is enabled through its configuration register and allows the discharge of an internal capacitor when triggered by an external signal. The GBT-SCA DACs are used to set the voltage level across the capacitor that defines the injected charge. These levels are static or very slowly changing therefore no special precautions are needed in the routing. There are 4 available DAC outputs which will allow the independent setting of the test-pulse level on each of the connected half-ECs. The external signal that triggers the test pulse injection must be synchronous with the LHCb readout. For this reason, this digital signal is driven by an FPGA IO in response to a synchronous command arriving on the master link and decoded in the FPGAs. The signal is driven by separate buffers for each EC and is fanned out within the EC to the CLARO ASICs.

4.5   FE temperature monitoring

The MAPMT photocathode is sensitive to temperature and must be operated within acceptable limits (less than 35°C). Temperature sensors (pt1000) are foreseen at various locations on the EC and these are connected directly to GBT-SCA ADC inputs. A precisely calibrated temperature measurement is not required – a resolution of a few °C is sufficient to monitor the expected temperature variation during operation. The SCA ADC inputs implement a programmable 100uA current source and this results in a voltage of around 0.1V at room temperature at the ADC input to be compared with the full-scale range 0-1.0V. The 12-bit conversion results in a sensitivity of roughly 2°C per LSB which is sufficient for our requirements.

4.6   FE SEL monitoring

SEL has been observed at a very low rate under heavy irradiation of the CLARO ASICs. In order to detect the presence of latch-up, the FEBs implement a circuit that combines the CLARO digital SEL detection signal into an analogue output. The PDMDB routes these signals to SCA ADC inputs where a change in the analogue voltage would indicate a possible SEL on one of the CLAROs.

Some additional sensitivity to FE SEL is provided by the current sensing network of the 2.5V DCDC converter on the PDMDB as described later in Section 7.1.

5   FPGA programming interface (JTAG)

The Kintex-7 FPGAs are programmed remotely through the SCA JTAG TAP. The SCA implements a single TAP and all three FPGAs are connected in a single chain as shown in Figure 12. The Kintex-7 uses the TCK, TMS, TDO and TDI signals. TRST is not used. The TCK signal is terminated at the end with a split termination network. With this configuration, the FPGAs can be programmed at 20MHz, the maximum speed permitted by the GBTX-SCA. Figure 11 shows the TCK and TMS signals measured at the edge connector with the SCA TAP operating at the maximum clock speed. Sampling is at the rising edge of TCK so the data is well-centred on the clock edge and comfortably satisfies the set-up (3ns) and hold (2ns) time requirements of the FPGA. The FPGA JTAG interface is specified to operate at up to 66 MHz. The signal amplitude is 2.5V for TMS but slightly less (about 2V) for TCK due to the split termination network. The FPGA switching threshold (VCCO/2) is well matched to the TCK signal which is symmetric about VCCO/2. The signals shown in the figure are expected to be worst case since they are far from the termination. TDI and TDO have the same characteristics as

TMS but, since they are daisy-chained through the FPGAs, the signal traces are shorter and the reflections seen on the TMS transitions are even less prominent for these.
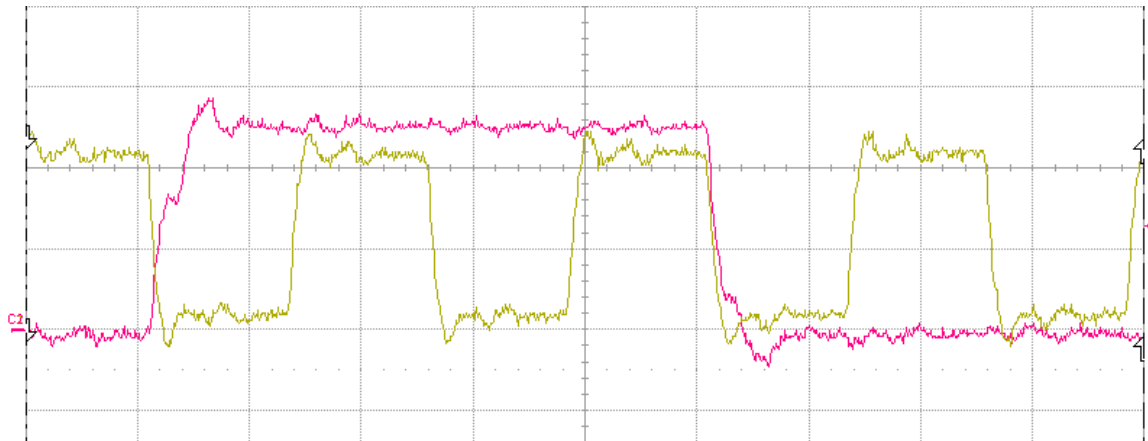
Figure 11 TCK and TMS measured at the edge connector.

The FPGA logic occupancy is very low (few %) therefore the Xilinx "compressed" bit-stream format is effective at reducing the bit-stream size. This helps to minimize the re-programming time and is an important consideration because of the need to perform scrubbing in the event of un-recoverable radiation-induced CRAM upsets.

With the SCA JTAG interface operating at maximum speed, the programming of each FPGA takes a few seconds.
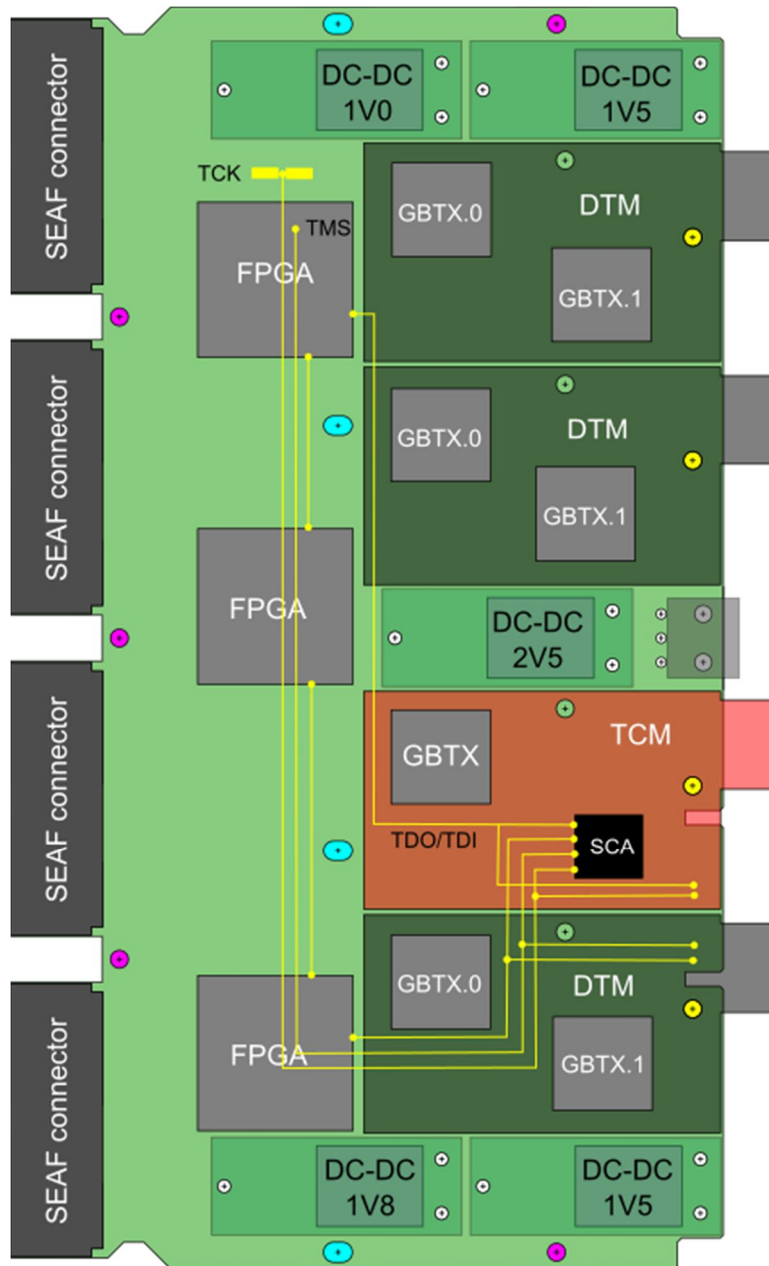
Figure 12 JTAG topology

6 Data transmission

The PDMDB captures digital data from the ECs, formats them and forwards them for onward transmission to the PCIe40 via the PDMDB-DTM plug-in. All the connections from the DTM GBTXes to the motherboard are through a 10×10pin Megarray connector (FCI-84512-002LF, mates with FCI-84513-001LF).

The underlying data transmission is implemented using the Versatile Link hardware whose core element, the GBTX transceiver ASIC, defines the low-level protocol. A higher protocol layer [LHCBDAQ] has been defined by LHCb in common for all subdetectors in which the GBTX ASICs operate as simplex transmitters. Some limited flexibility is allowed within this protocol including the choice between using the GBTX *wide frame mode* or *forward error correction mode*. Since the cost of

the links is high and the *wide frame mode* has less overhead, this mode is used for the RICH detectors. Occasional bit errors in the front-end data part of the data frame are expected to have negligible effect on the data quality.

Due to the likelihood of radiation-induced upsets, the role of the FPGAs is minimal by design. They function essentially as level translators between the CLARO outputs and GBTX inputs, latches for the incoming data to synchronise with the system clock and serialisers to multiplex the data into the uplinks. Because the data path behaves as a fixed latency transparent connection to the remote readout modules (PCIe40), it has been possible to relocate most of the architectural and synchronisation logic into the remote modules. In this way, it has been possible to implement a largely stateless design with a very small logic footprint.

The GBTX serialiser multiplexes the 112 bit data frames arriving on the *e-links* into the serial output link with a choice of multiplexing ratio. In 160 MHz DDR mode, 14 data *e-links* are sufficient to implement one transmission link. There are sufficient FPGA IOs in the *xc7k70tfbg676* to implement two such data links (allowing also for clocks) and this, together with the fact that the VTTX is a dual-transmitter module, motivates the modularity chosen for the PDMDB and its PDMDB-TCM plug-ins. The reference clock input for the PDMDB-DTM is sourced by the GBTX ASIC on the PDMDB-TCM.

Figure 13 gives an indication of the signal paths on the PDMDB between the CLARO outputs and FPGA (single-ended) and the data uplinks between the FPGA and DTM connector (differential). The figure shows the top layer only. A similar number of traces are routed on internal layers. The PDMDB-H module follows the same principle but with two FPGAs instead of three and repositioned to optimise the layout.
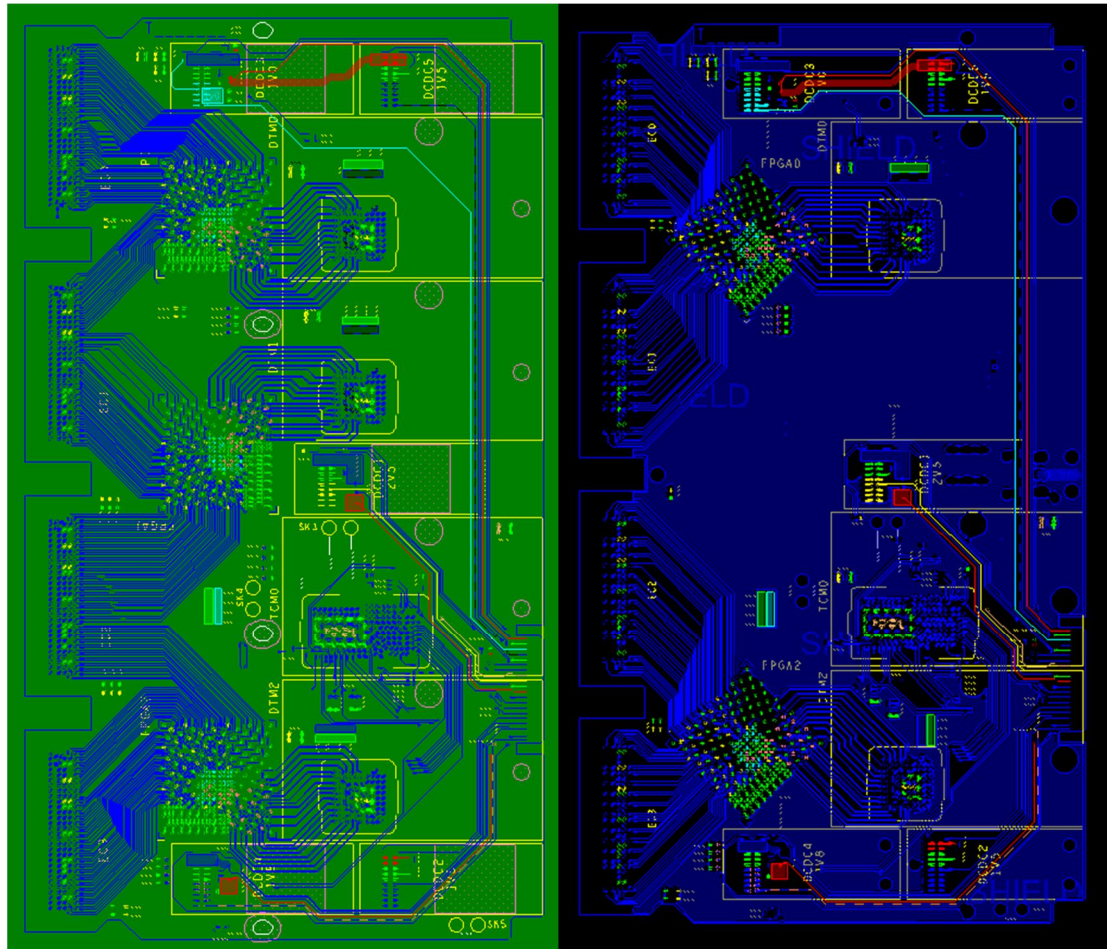
Figure 13 FPGA PMT and uplink signals. Left PDMDBR, right PDMDBH.

7   PDMDB monitoring

The role of the PDMDB as the control and monitoring interface for the ECs has been covered in Section **Error! Reference source not found.**. In addition, there are a few parameters on the motherboard itself that are also monitored.

7.1   Current monitoring

In order to detect SEU/SEL in the FPGAs it is desirable to monitor the current drawn on the critical supplies. The DCDC regulators do not themselves monitor the current, so an indication of the currents is obtained by using a small-value resistor in series at the input of each of the 3 DCDC converters that supply the FPGAs as shown in Figure 14. The voltage drop across these series resistors is read using SCA ADC channels. The resistor networks have been selected taking into account the maximum expected input current (after irradiation) to give an acceptable sensitivity and acceptable voltage drop at the DCDC inputs. It is not practical to use this type of passive network on the output side since good sensitivity cannot be achieved without unacceptable voltage drop.

Current sense resistors are only required at the inputs to the 1V0, 1V8 and 2V5 DC/DC converters. To ensure that the ADC input is kept within its input range (0-1V), the output from each sense resistor

must be passed through a potential divider. A 1/7.8 divider (6k8, 1k0) ensures that the ADC input is within acceptable limits over all possible operating conditions.

To minimise the number of ADC connections required, only one high-side ADC connections will be used. Thus the four required connections to the ADC are:

- High side (5-7 V input)
- Low side of 1V0 input
- Low side of 1V8 input
- Low side of 2V5 input

Heat dissipation will be improved by adding as large an area of copper connected to each end as possible.

Output voltages will not be monitored during operation, but will be tested during acceptance of the digital boards using the test connector.

Suitable part numbers are:

0R33 resistor (1V0, 2V5): Welwyn LR2512-R33FW (2512, 2W)
1R0 resistor (1V8): Welwyn LR2512-1R0FW (2512, 2W)

The required values and the estimated sensitivity and voltage drops are tabulated in Table 3.

PSU	Sense resistor (ohms)	Power dissipation (W)		Overall sensitivity		Voltage drop (mV)	
		Typical	Worst-case	mV/A	ADC bits/A	Typical	Worst-case
1V0	0.33	0.29	0.44	10.0	41.2	312	383
1V8	1.00	0.07	3.24	47.3	193.6	273	1800
2V5	0.33	0.15	1.83	20.4	83.5	226	776

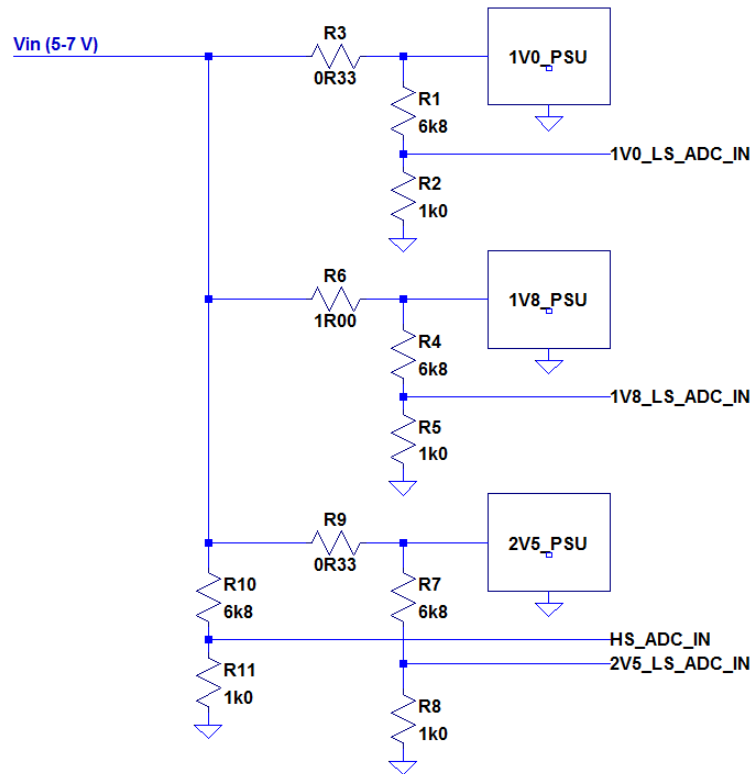**Table 3 Current sense resistor values**

Figure 14 Current monitoring resistor network

7.2   FPGA temperature and core voltage monitoring

The Kintex-7 FPGA has an internal ADC that can sample an internal temperature sensor and the core and auxiliary voltages. The sampled values can be read through the SCA JTAG TAP. These are not considered to be critical parameters for monitoring but can be used to supplement other monitoring.

8   Powering

The PDMDB and EC components require several different voltages. These voltages will be generated locally to the PDMDB using FEASTMP-CLP DCDC converter modules plugged onto the PDMDB and powered remotely with Wiener Maraton modules. The different voltage levels define the minimum number of distinct powering domains. However, for operational reasons, it is desirable to independently power the PDMDB-TCM allowing it to remain on while the ECs and FPGAs are powered down.

Since the 2.5V rail is common between the controls interface and the FE electronics this requires the use of more than one 2.5V DCDC converter. Furthermore, the load on the 1.5V rail is such that two DCDC converters are required to source enough current but it is not permitted to connect the outputs of the two DCDC converters in parallel. This requires that the 1.5V domain is also split into two sub-domains. The design of the PDMDB power domains is therefore based on a compromise between minimising the number of required DCDC converters and satisfying the operational demands. Table 4 summarises the proposed powering scheme with the "mandatory always on" domains highlighted.

Domain	Voltage [V]	Current [A]	Scope
TCM-IO	2.5		GBT-SCA IO, VTRX, (3 VTTX)
GBT-core/1	1.5		4 GBTX
GBT-core/2	1.5		3 GBTX, SCA core
EC2V5	2.5		4 half-EC, 3 FPGA HR bank IO, (3 VTTX)
FPGA-core	1.0		3 FPGA
FPGA-IO	1.8		3 FPGA HP bank IO

**Table 4 PDMDB power domains**

Due to space limitations on the PDMDB and the low power requirement, a single DCDC converter per column that powers the **TCM-IO** domain will be used and it will be located at the end of the column.

8.1 Input filter capacitors

Analysis of the power requirements of the PDMDB was prepared at the time of the PDMDB EDR and PDMDB-TCM and PDMDB-DTM PRR. Following this analysis and recommendations during these reviews some additional local capacitive filtering has been implemented for the raw power supply input of each PDMDB. A suitable layout of these capacitors is illustrated in Figure 15.

Capacitors at the power input to the digital board are important for two reasons: to attenuate noise conducted into the board from the cables, and to improve the immunity of the DC/DC converter to power supply glitches, by giving their control loops time to respond to changing input voltages.

To meet these requirements, the following input capacitors are used (16 V rated) and located on the underside of the boards. The maximum height of these capacitors is 2mm.
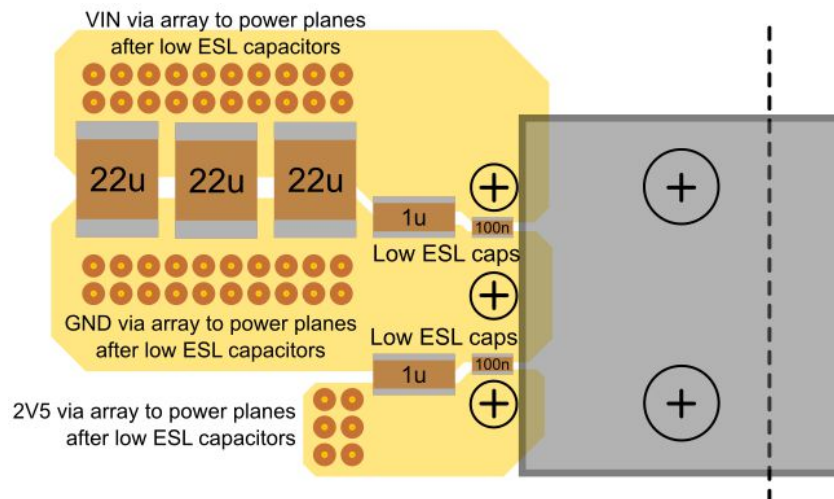
Figure 15 Layout of input filtering capacitors

8.1.1 High frequency filter capacitors

To attenuate high frequency conducted interference, two low ESL capacitors are used: 100nF || 1µF. These have very low impedance at the most sensitive frequencies of the detector. The 100nF capacitor is placed as close as possible to the legs of the input power connector. The 1µF capacitor is immediately adjacent to it, such that the entire supply current passes both capacitors before being

distributed to the rest of the board. These capacitors are placed on both the 5-7 VDC and 2.5 VDC inputs.

The preferred part numbers are:

- 100nF Use LLL185R71A104MA01L or similar 100nF X7R 10V part
- 1µF: 2 x 0603YC105KAT2A or similar 1µF X7R 16V part if space allows, otherwise 1x

8.1.2 Bulk capacitors

Disturbances on the main power supply can be compensated for by the DC/DC converters with negligible effect on the rest of the board, provided the disturbance happens within the control bandwidth of the converters.

Taking into account the existing input capacitance of the DC-DC converters, it is estimated that the following input capacitance should be adequate for the PDMDB:

- 3x C4532X7R1C226M200KC (1812), or equivalent 1812 X7R ≥16 V 22 µF part. If availability of this size is difficult, replace with 4 x 1210 X7R ≥16 V 22 µF.

8.2 Power sequencing

To comply with the recommended power-up sequence for the FPGA, the DCDC *power-good* and *enable* signals are used. The **FPGA-core** domain acts as the master for the **FPGA-IO** and **EC2V5** domains. The *power-good* output of the DCDC supplying **FPGA-core** is connected to the *enable* inputs of the DCDCs supplying **FPGA-IO** and **EC2V5** ensuring that the IO banks are powered after the FPGA core voltage is applied.

When the Maraton raw power to the PDMDB module is applied, **TCM-IO**, **GBT-core/1** and **GBT-core/2** power up immediately. **FPGA-core** and its slaves **FPGA-IO** and **EC2V5** remain off until **FPGA-core** is enabled through the controls interface.

8.3 Plug-in power filtering

The TCM/DTM PRR identified a lack of sufficient bulk capacitance near the modules for the 1.5 V and 2.5 V supplies. This is corrected in the updated PDMDB. Extra 1206 bulk capacitors are fitted underneath the modules to avoid conflict with other components and to allow them to be placed close to the Megarray connector. Based on measurements of power plane impedance, the following bulk capacitors have been fitted near the modules to optimise the impedance over a wide frequency range:

TCM:

- **1.5 V rail:** 2× 68 µF
- **2.5V rail:** 1× 68 µF

DTM

- **1.5 V rail:** 4× 68 μF
- **2.5V rail:** 1× 68 μF

8.3.1 Grounding and shielding

The grounding strategy is unchanged from the EDR, except that the ground plane on the bottom of the PCB is only connected at the input connector and has no ground return currents flowing across it. This change is intended to minimise the possibility of noise being coupled between the HV cables and the PDMDB.

9 Reset and configuration

9.1 Master GBTX

The master GBTX link must come up in an operational state on power-up. This will be done by burning a minimal set of e-fuses during production of the PDMDB-TCM modules. Once the link is up, all registers can then be programmed through the GBTX link. The master GBTX must be configured first as it provides the reference clocks for the data transmission GBTXes. There is no external reset signal. Reset of the master link requires a power cycle of the entire column. An RC network on the TCM provides the necessary reset time constant.

9.2 TX GBTX

The data transmission GBTX chips operate in simplex TX wide-frame mode and therefore cannot be configured through the GBT link. These will be configured through their I2C interfaces with the TCM GBT-SCA acting as the I2C masters. Figure 16 shows the first byte of an I2C transaction for DTM2 with 400 kHz SCL frequency. The I2C protocol specifies that SCL should change only while SDA is high and this can be seen to be satisfied with some margin for bit 6 of the first (slave address) byte. The small runt pulse on SDA is where the slave releases the acknowledge. The speed of the rising edges is determined by the strength of the pull-up resistors (4k7 for this example) and the capacitance/loading of the bus. The transactions have been measured to be reproducibly reliable in this configuration even at the maximum I2C speed supported by the SCA (1MHz). However, to increase the margin we will most likely use stronger 2k2 pull-ups for the final design. In any case, the I2C speed is not a critical parameter since the amount of data transferred per GBTX is small (a few kBit) and it will be possible to configure all of the PDMDBs in parallel with the LHCb upgrade control system.
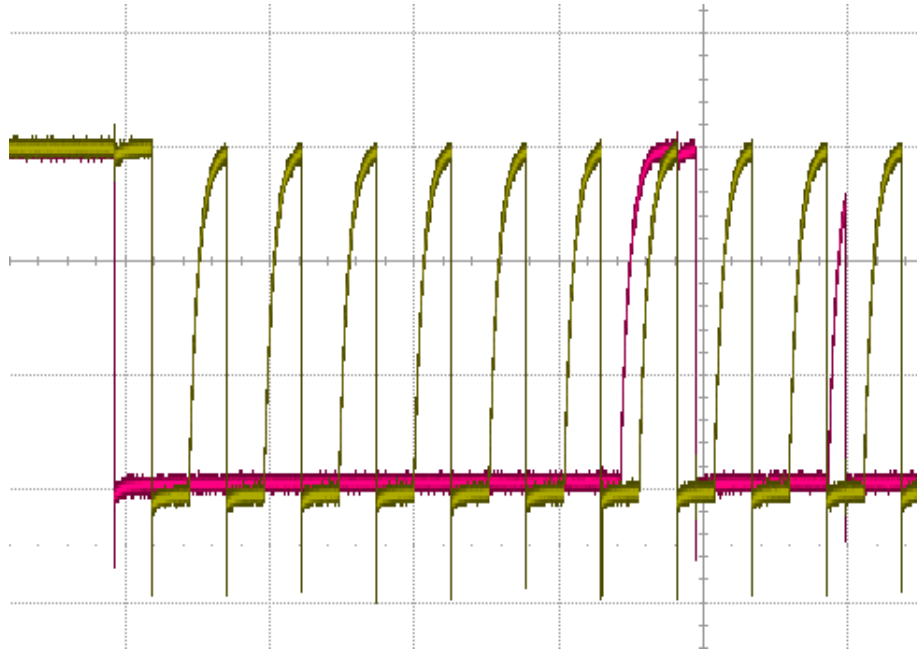
Figure 16 DTM2 I2C SCL & SDA at 400kHz

Two GBTXes belonging to a DTM and its FPGA share the same SCA I2C bus. In our configuration, the DTM GBTXes may power-up in the absence of an external reference clock. Under these conditions it has been observed that a GBTX may block the I2C bus to which it is attached (by e.g. pulling SCA low even if it is not addressed). Therefore, the following reset sequence is used:

- Assert and hold **both** GBTXes reset (GPIO).
- Deassert reset of the first GBTX **only**.
- Configure the first GBTX.
- Deassert reset of the second GBTX.
- Configure the second GBTX.

The DTMs may be reset and configured independently of each other.

9.3 FPGAs

There is no dedicated external reset for the FPGAs. However, each FPGA has one IO connected to a TCM GBT-SCA GPIO. An FPGA is reset by remotely reprogramming the device with or without a preceding power cycle.

10 FPGA firmware

A very minimal firmware is proposed in order to minimise the probability that radiation-induced upsets will interfere with the operation. In this scheme, the FPGAs perform level translation between the input (2v5 CMOS) and output (LVDS) IO standards, pack the data into fixed-format GBT frames, and drive the data transmission e-links through 8:1 serialisers. The FPGAs also decode TFC data fanned out by the TCM in order to generate a synchronous front-end test-pulse signal.

10.1 Clocking

The FPGA clocking scheme must preserve the synchronisation with the experiment-wide clock. This is achieved by using 40MHz TCM GBTX phase-shifter clock outputs as the FPGA reference clocks. However, the 8:1 e-link serialisation requires also a 160MHz IO clock. A 160 MHz FPGA reference clock synchronous with the experiment clock is driven by a phase-shifter output from GBTX0 of its corresponding DTM. The different propagation delays of the 40MHz and 160MHz reference clocks means that the phase relationship between them, although constant, may not be the same at the FPGA and the DTMs. It is therefore necessary to tune the phase relationship in the FPGA to align with the DTMs. Typically, this would be done with a PLL but an alternative scheme, using the clock-dividers built in to the Xilinx regional clock buffers has also proven effective. Such a scheme is illustrated In Figure 17.
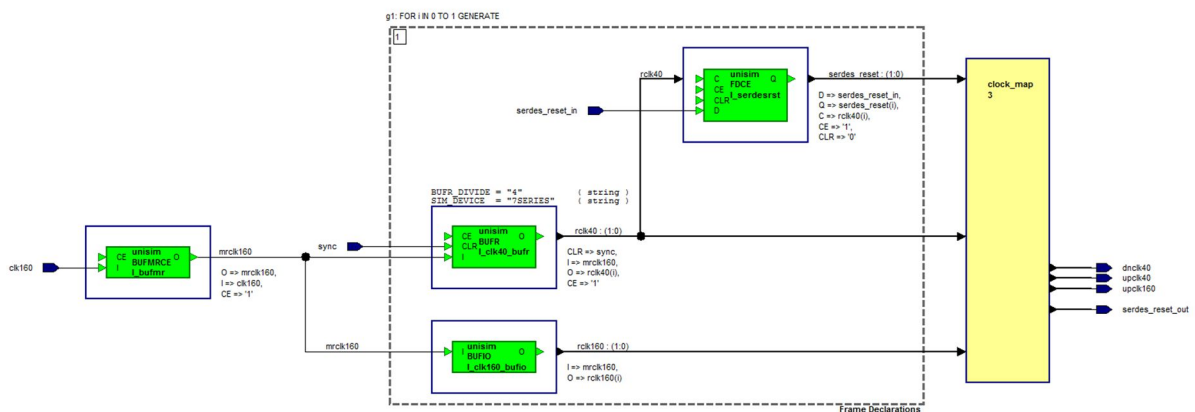
Figure 17 Data transmission clocking scheme

Since the e-links are spread across two clock regions, two regional clock buffers (BUFR) are generated and these are themselves driven by the 160MHz reference clock through a multi-region clock buffer (BUFMRCE). The BUFRs are configured to divide the input clock by 4 to provide a new 40MHz clock. The synchronous CLR pin then allows the phase of these regional 40MHz clocks to be adjusted relative to the 40MHz FPGA reference clock. IO clock buffers (BUFIO) drive the 160MHz clock to the serialisers. A reset for the serialisers, resynchronised to the 40MHz regional clock is also generated. A static mapping is then used to select the appropriate BUFR, BUFIO and serialiser reset output for the e-link serialiser depending on its physical location.

The above logic requires the generation of a synchronisation signal to reset the BUFR clock dividers. This is easily done with a simple chain of shift registers clocked by the global 160MHz clock and triggered by the 40MHz global clock. Using the shift register length to adjust the delay allows the BUFR divider output phase to be adjusted relative to the global reference clock.

Having aligned the data transmission clocks for DTM GBTX0 it is then necessary to also align DTM GBTX1. The simplest way to achieve this is to use the phase adjustment of the GBTX0 phase shifter output that provides the 40MHz reference clock for GBTX1.

The same static settings are found to be effective for all FPGAs and all DTM GBTXs on a single PDMDB but this is still to be verified for all boards.

Finally, the sampling phase at the receiver must be tuned by phase aligning the GBTX e-ports. These are operated in widebus mode at 160MHz DDR so only e-ports 0 and 4 of groups 0 to 6 need to be adjusted. Figure 18 shows the result of scanning the phase for a selection of e-links. The system is set up to transmit a 10101010… pattern on each e-link. The range of the phase adjustment is 2 data

bits. A very sharp transition from 1 to 0 or 0 to 1 is observed in all cases. From the figure it is clear that using a value of 5 for the phase setting centres the sampling edge in the middle of the stable region with a margin of at least ±2. Although only a small selection of e-links is shown, it has been found that the variation across all links amounts to only about ±1 leading to the conclusion that the same setting can be used on all links.
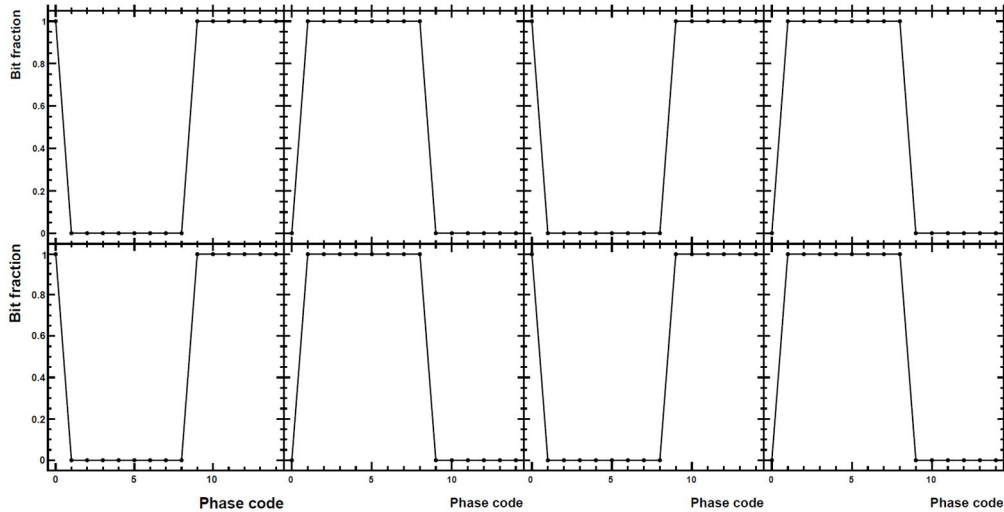

Figure 18 Data transmission e-link phase scan

10.2 Front-end data capture

The CLARO front-end ASICs are asynchronous – the digital outputs are asserted whenever the analogue input exceeds a programmable threshold. These signals must be captured and synchronised to the global 40MHz clock while preserving good efficiency but also minimising spillover into the following clock cycle. Figure 19 illustrates a possible solution by sampling the input with the 160MHz clock and generating a 25ns output pulse. The output pulse width never exceeds 25ns even for longer input pulses. This capturing scheme can be expected to lose efficiency when the input pulse is less than 6.25ns.
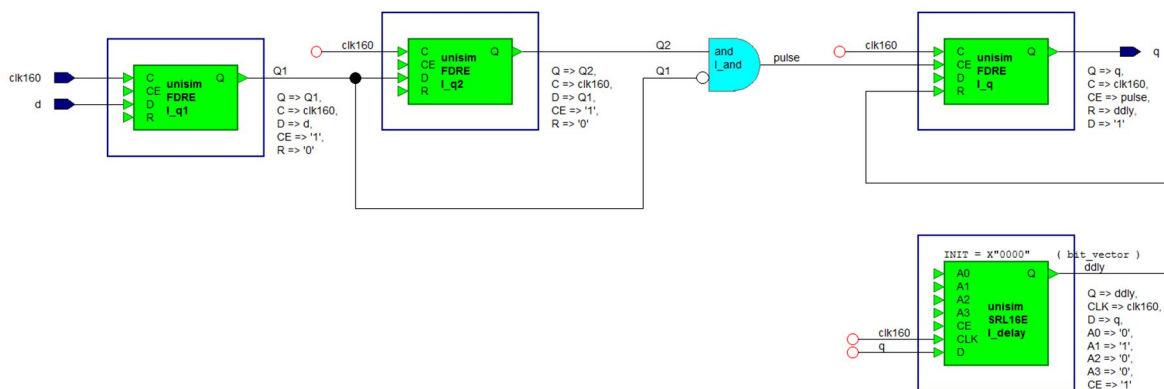

Figure 19 Example front-end signal capture logic.

Studies have shown that signal Cherenkov photons from the same bunch crossing arrive within a few nanoseconds of each other. Therefore variants of this scheme in which the output pulse width is reduced to 12.5ns or even 6.25ns might be beneficial in helping to reduce background.

11 Test connector

An edge connector is implemented. Not used after installation, this connector allows some additional parameters to be monitored during production testing.

The edge connector is a 40-pin edge connector for mating with Molex part number 45984-0183 (Figure 20). Note that this version of the connector does not have any guides. The connector datasheet (which includes the specification of the PCB contacts) may be found at http://www.molex.com/pdm_docs/sd/459840183_sd.pdf
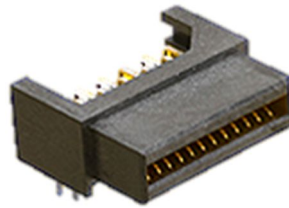
Figure 20 Molex 45984-0183

The pinout for the connector is given in Table 5. Note that the circuit numbering is such that pins 1 – 20 are on the bottom row (underside of PDMDB) and 21 – 40 are on the top row (top side of PDMDB) and run in the same direction. I.e. pin 1 is on the bottom of the PDMDB, with pin 21 on the top surface above it. Then follows pin 2 on the bottom with pin 22 above, 3 and 23 and so on. On the PDMDB pins 1 and 21 are closest to the top edge of the PCB.

Since the connector is not keyed, the pinout is designed to be able to be capable of being connected backwards without damaging the PDMDB or the test board. This is subject to the caveat that all DC/DC enable pins must be forcibly held low by the test PCB whilst the polarity is checked. Polarity checking will be implemented by attempting to set CONFIGSELECT high, by applying 1.5V to the pin via a 100Ohm resistor. If the pin voltage remains at zero, then CONFIGSELECT is shorted to ground due to incorrect polarity, so the test board will keep everything powered down and alert the user.

Pin numbers			
Assignment	Top	Bottom	Assignment
VIN1V0 I+ (Vcc)	21	1	VIN1V0 I-
GND	22	2	PG1V0
1V0	23	3	EN1V0
B1V5	24	4	EN1V5B
GND	25	5	PG1V5B
TDI	26	6	TMS
TDO	27	7	TCK
2V5VTRX	28	8	GND
2V5	29	9	PG2V5
GND	30	10	VIN2V5 I-
VIN2V5 I+ (Vcc)	31	11	EN2V5
GBTFUSE	32	12	GBTFUSE3V3
CONFIGSELECT	33	13	FPGA_0
VTRX_SCL	34	14	FPGA_1
VTRX_SDA	35	15	FPGA_2
PG1V5A	36	16	GND
EN1V5A	37	17	A1V5
EN1V8	38	18	1V8
PG1V8	39	19	GND
VIN1V8 I-	40	20	VIN1V8 I+ (Vcc)

**Table 5 Pin assignments for the test connector**

12   PCB Technology

Prototypes have been fabricated with Eurocircuits using an 8-layer FR4 stack-up with minimum conductor feature size of 0.1 mm and minimum drill size of 0.15 mm. The maximum signal speed on the PCB is 160 MHz DDR over short 100 Ohm nominal differential traces. The nominal impedance and employing good routing practice gives satisfactory results without the need for impedance matching during PCB manufacture.

For production, Wrekin International will be used and the boards will be fabricated from halogen-free materials. The same company and technology is used for the production of the PDMDB-TCM and PDMDB-DTM and has produced satisfactory results.

13   Mechanical

The PDMDB mechanical envelope is strongly constrained by the close-packing of the photon detector array. The board height (220mm) is dictated by the pitch of the elementary cells and allows for a small gap between boards when mounted on the columns. The board width (124mm) is constrained by the space behind the RICH columns for cables and other services. The board outline and principal dimensions are shown in Figure 21. The outline includes cut-outs between the SEAF connectors to allow the connectors to be inserted through holes in the column structure and profiling at the top and bottom edge to allow for access to mounting holes.
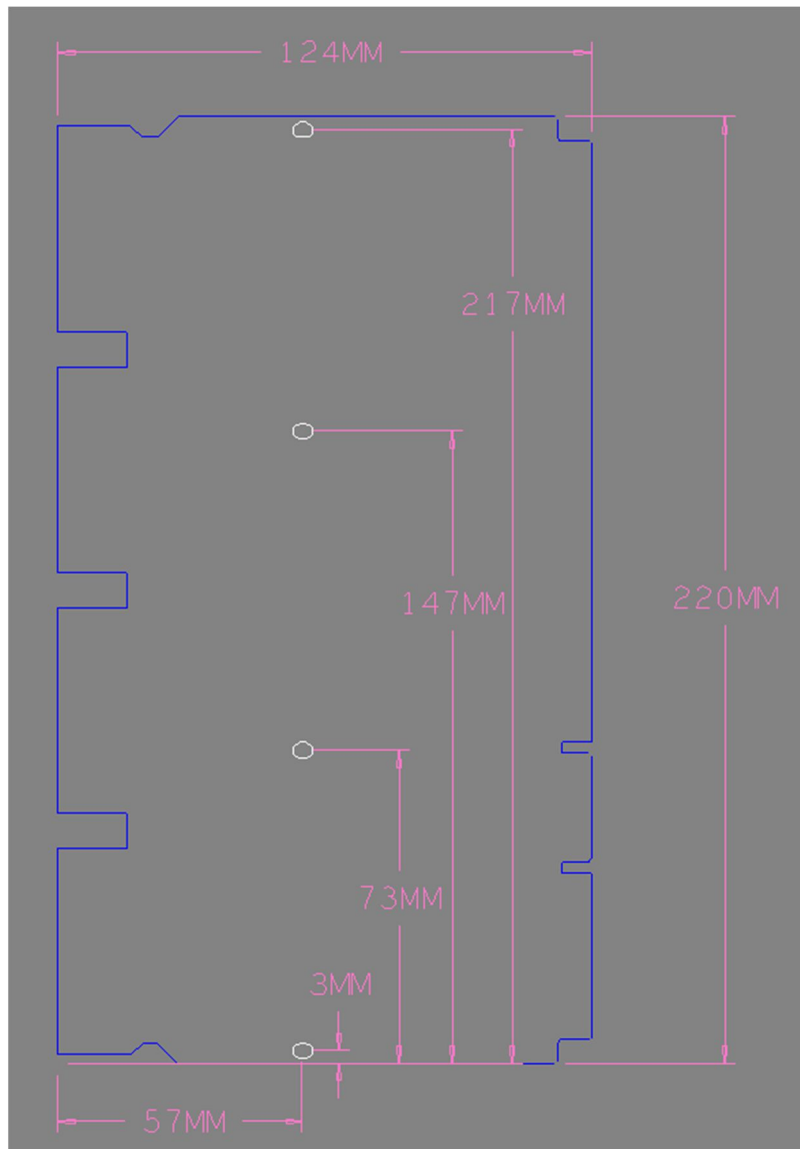
Figure 21 PDMDB dimensions and outline

More details of the mechanical, power and thermal integration of the PDMDB onto the photon detector columns can be found in a supplementary document [INTEGRATION].