Production readiness report for the RICH PDMDB timing and control module (TCM) and data transmission module (DTM)

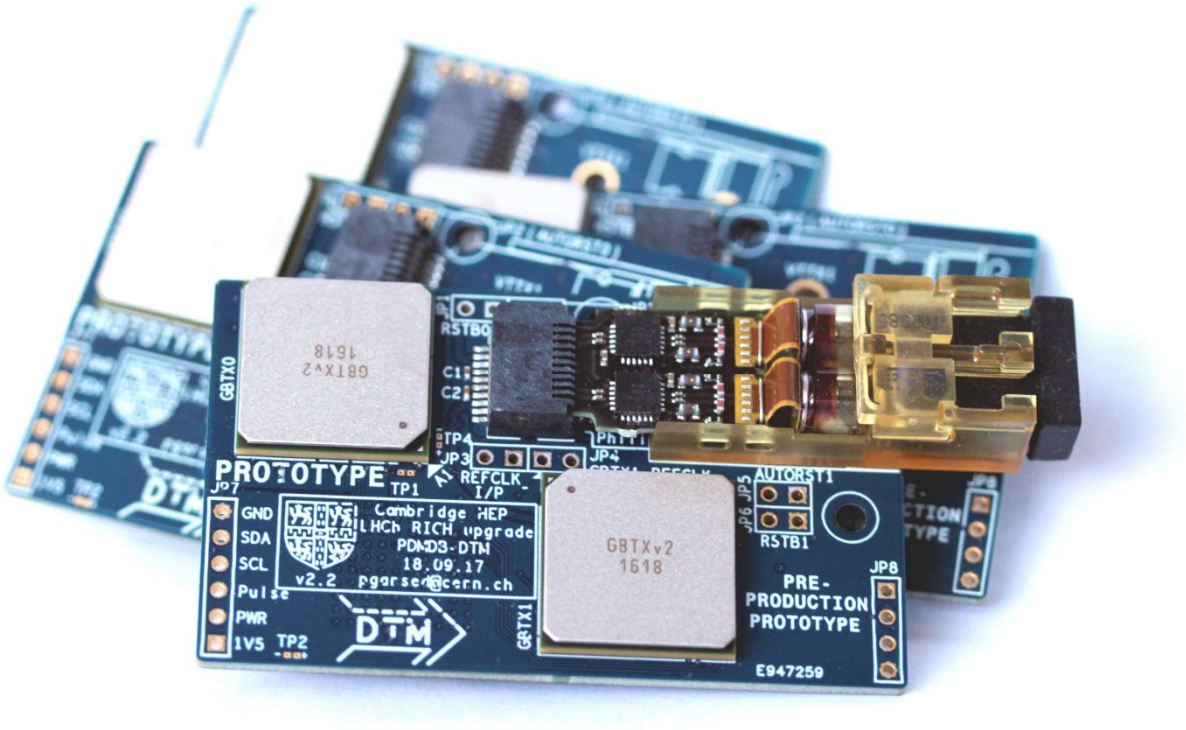

1. Introduction

The connection between the counting room and the RICH detector front-end digital board (PDMDB) is implemented using fibre-optic links. Each PDMDB will require up to 6 × 4.8 Gbps fibre-optic links to transmit experimental data to the counting room, in addition to one RX one TX 4.8 Gbps fibre-optic link to receive and transmit TFC and ECS data to and from the SOL40 module.

These links are implemented using the radiation-hard versatile link components developed at CERN. Thus, experimental data will be transmitted using three identical VTTX dual optical transmitter modules, each driven by a single GBTX chip for each transmitter. TFC and ECS data will be received and transmitted using a single VTRX optical transceiver module, in tandem with a GBTX chip operating in transceiver mode and a GBT-SCA to provide a slow controls interface to the digital board.

Due to concern about the radiation hardness of the PDMDB FPGAs, as well as the inherent duplication of circuitry in the design, a modular approach has been adopted and two plug-in modules have been developed for the PDMDB motherboard. The ‘data transmission module’ (DTM) implements the GBT chipset and versatile link components for a dual optical data transmitter. The ‘timing and control module’ (TCM) implements the GBT chipset and versatile link components for an optical transceiver.

Communication between these plug-in modules and the main digital board takes place using wired E-links and a small number of single-ended control signals. Additionally, the DTM receives a clock

reference from the TCM module[1] and an I2C bus from the digital board, whilst the PDMDB-TCM module provides reference clocks, GPIOs, ADC inputs, DAC outputs, I2C busses, a JTAG interface and an SPI interface to the digital board, PDMDB-DTM modules and front-end elementary cells (ECs)

As well as the reliability, replaceability and duplication advantages mentioned above, this modular approach also improves testability, by making it practical to characterise individual optical link components independently of the digital board. It also follows the approach encouraged in [1], which notes that: “it is efficient to use one duplex GBT link to provide ECS/TFC information to many front-end components”.

This section describes current state of these modules, their readiness for production and plans for testing the completed modules.

1.1. Requirements and specification

The requirements and specifications for the modules are unchanged since the EDR and may be found in that document, available at https://indico.cern.ch/event/515232/

2. Mechanical design

2.1. Layout and mounting

Both modules have been designed with dimensions of 62 x 38 mm. This is the smallest area that is able to comfortably accommodate the main components of each board and allow sufficient area for routing. The components are arranged to first minimise the length of the 4.8 Gbps traces between the GBTXs and the optical modules and then to minimise the length of connections between the GBTX/SCA and the connector the digital board.

Top and bottom views of the pre-production PCBs for the TCM and DTM are shown in Figure 1 and Figure 2 respectively.

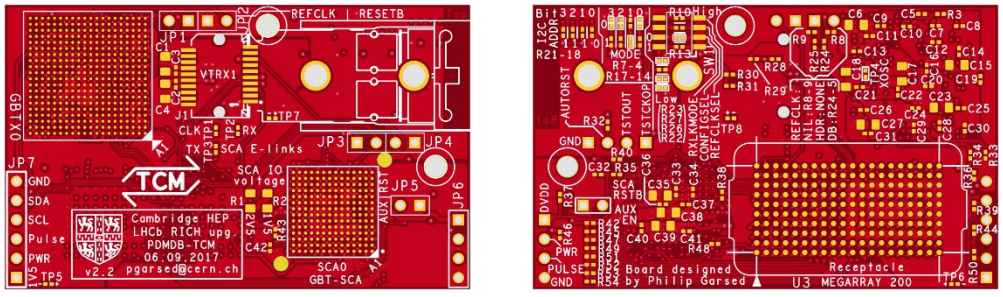

Figure 1: Top and bottom views of the pre-production design for the TCM module, shown actual size.

[1] This reference clock is used to synchronise only GBTX 0 on the DTM. GBTX 1 is synchronised using a reference clock from GBTX 0. This approach was taken to minimise the number of clock references required at the module connectors and on the PCB, as well as to avoid potential problems from split termination of the clock lines.

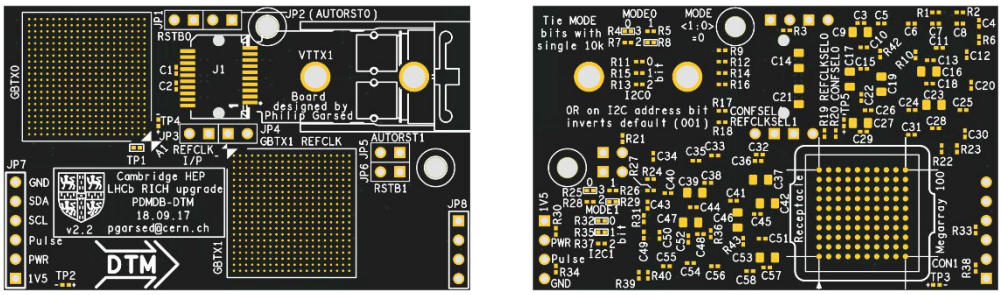

Figure 2: Top and bottom views of the pre-production design for the DTM module, shown actual size.

Each module is connected to the digital board through a 4 mm high Megarray mezzanine connector. The TCM connector is the 200 pin variant, whilst the DTM connector is the 100 pin version. In addition, each module has an M2.5 (2.7 mm diameter) hole on its right hand edge and will have a 3 mm diameter (currently 2.7 mm diameter) hole close to the top edge, adjacent to SFP+ connector for the VTxX[2], partially underneath the VTxX PCB, when viewed from the top.

The right-hand hole is used for securing the TCM/DTM to the PDMDB with a 4 mm spacer, as well as securing the levelling plate to the module. This prevents the Megarray connector taking excessive twisting forces and ensures the space between the top surface of the PDMDB and the bottom surface of the module is a consistent 4 mm.

The hole close to the top edge takes a 4 mm tall plastic clip-in spacer (Ettinger 07.91.340) that rests against, but does not screw in to the PDMDB. This is intended to take some of the load exerted by the levelling plate (section 2.2) on the GBTX, VTxX and SCA, to minimise bending of the PCB and to ensure that the Megarray connector remains flat against its mating half.

VTRX and VTTX modules plug into the 20-pin SFP+ connector on the TCM and DTM respectively. They are attached by small screws from the underside of the modules. The VTxX modules protrude approximately 11 mm from the RICH-DB-VTXX PCBs. This makes it easier for optical fibres to be connected and disconnected with the board in situ.

The position of the TCM and DTMs on the PDMDB motherboard are shown in

[2] The term VTxX is used to refer to both the VTRX optical transceiver module and the VTTX dual optical transmitter module

Figure 3.

Figure 3: Position of TCM (red) and DTMs (black) on PDMDB motherboard (50% actual size)

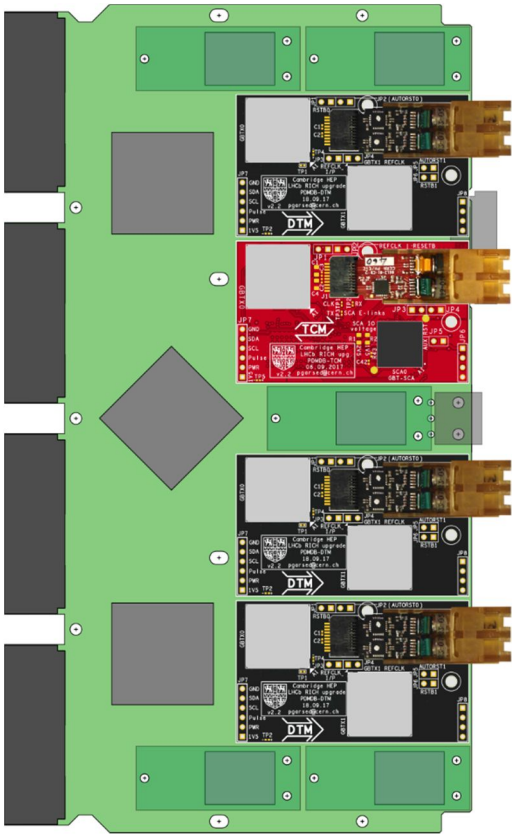

2.2. Thermal design

All components that generate significant heat are placed on the top surface of the TCM and DTM, allowing heatsinking to the cold bar that this side of the modules face. Different components are connected to the water-cooled cold bar as follows:

- The GBTX and SCA chips have 1 mm and 2 mm adhesive thermal pads respectively on their top surface. An aluminium levelling plate presses on this thermal pad and the chip and is screwed into the cold bar.
- The GBLD ICs on the VTTX and VTRX optical modules have a 1 – 1.5 mm adhesive thermal pad on their top surface. This attaches directly to the cold bar.

This arrangement is shown in Figure 4 and Figure 5.

Figure 4: Arrangement of thermal pads on GBTX, SCA and GBLD ICs (Credit: Christoph Frei)

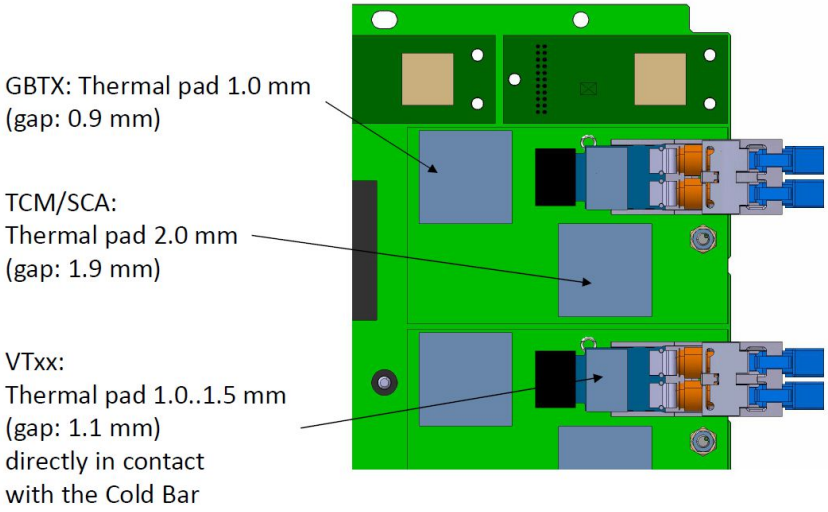

Figure 5: Arrangement of levelling plate on top of modules (credit: Christoph Frei)

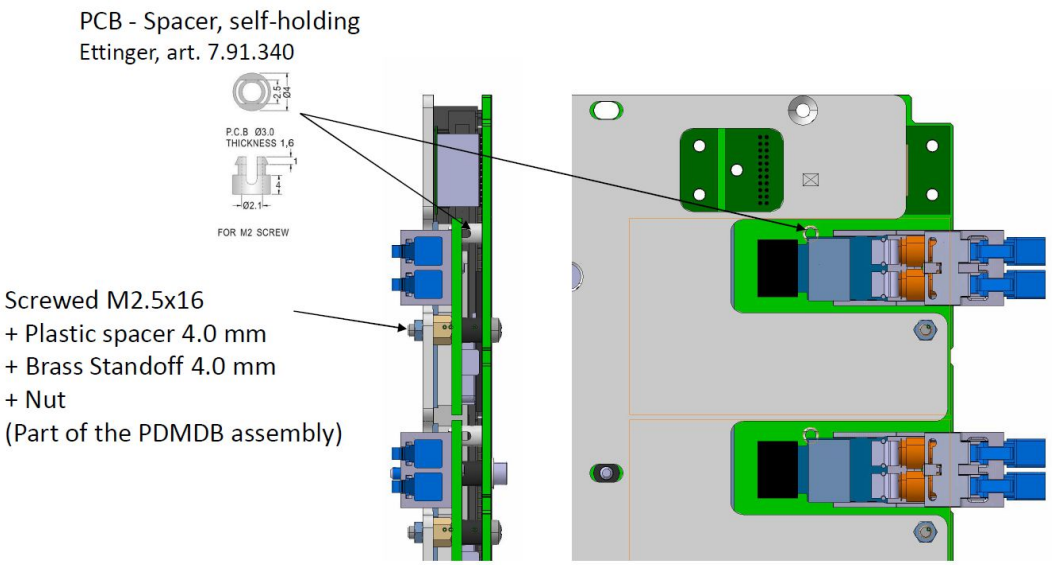

The effectiveness of this cooling arrangement has been tested in the lab using a prototype cooling bar. The operating temperatures for the various cooled components were measured at and ambient temperature of 24.3 °C, both when cooled only by natural convection and when cooled by heatsinking to a cooling bar at 20.5 °C. These measurements are shown in These measurements indicate that the cooling of the modules is effective at reducing operating temperatures by an average of more than 30 °C. This is particularly important given the thermal sensitivity of the MaPMTs at the front-end of the detector. The relatively low operating temperature of the components, combined with the inherently thermally isolated design of these plug-in modules should all act to minimise heat flows from the TCM and DTM to the front end.

Table 1.

These measurements indicate that the cooling of the modules is effective at reducing operating temperatures by an average of more than 30 °C. This is particularly important given the thermal sensitivity of the MaPMTs at the front-end of the detector. The relatively low operating temperature of the components, combined with the inherently thermally isolated design of these plug-in modules should all act to minimise heat flows from the TCM and DTM to the front end.

Table 1: Operating temperature of TCM and DTM components with natural convection and with heatsinking to a cooling bar (measurements made by Simon Juncker)

Component	Typical (worst-case) operating power (mW)	Heatsinking to	Thickness of heatsink material (mm)	Average operating temperature in free air (measured)	Average operating temperature with 20.5 °C cooling bar (measured) (°C)	ΔT to cooling bar best case measured (worst case estimated)
GBTX (wide-bus transmitter)	957 (1602)	Levelling plate	0.9	65.9	23.5	3.0 (5.0)
GBTX (FEC transceiver)	1452 (2097)	Levelling plate	0.9	66.0	22.7	1.6 (1.7)
GBT-SCA	73 (116)	Levelling plate	1.9	55.0	24.7	4.2 (6.7)
VTTX (overall power, split between two ICs)	500	Cold bar (direct)	1.1	62.8	28.1	7.6
VTRX	350	Cold bar (direct)	1.1	54.0	34.0	13.5

3. Radiation hardness

The radiation environment expected in the upgraded RICH detector are shown in Table 2. (Note that the figures quoted above are the expected radiation doses the modules will achieve in their working lifetime, not during one year of running as erroneously suggested in the EDR).

	RICH 1		**RICH 2**	
	Best case	Worst case	Best case	Worst case
Neutrons: 1 MeV n_{eq} [cm^{-2}] ($x10^{11}$)	18.5	30.5	8.5	15.5
Hadrons: >30 MeV [cm^{-2}] ($x10^{11}$)	5.0	11.5	2.9	5.0
Ionising radiation [Gy (krad)]	305 (30.5)	1980 (198)	225 (22.5)	795 (79.5)

Table 2: Estimates of radiation environment of RICH upgrade over the lifetime of the detector [5]. These figures assume 50 fb^{-1} for the upgrade phase after LS2 and include a safety factor of 2.

Both modules are radiation hard by design. The active components used on these boards are all CERN-specific radiation hard parts which have been extensively tested. Passive components including film resistors and ceramic capacitors are not known to be sensitive to radiation and are also used in other radiation components such as the FEASTMP modules. We are also not aware of any issues with the materials used in the proposed PCB substrate or Megarray connector.

An increase in power consumption by the radiation-hard components is expected over the life of the modules due to the effects of irradiation. This has been accounted for in the design of the PDMDB

power supplies and the cooling system, but does not otherwise greatly affect the design of the modules.

The current (broadly similar v2.0/2.1/2.2) versions of the TCM and DTM have been used extensively in beam tests in the North Area at Prevessin, as well as being used on FPGA test boards that experienced very high dose rates in the CHARM facility as well as in the LHCb cavern. To date, we are not aware of any failures or problems due to the radiation doses received by the modules.

4. Power consumption

In the PDMDB EDR, power consumption for the TCM and DTM was estimated from figures provided by the designers of the Versatile Link and GBT chipsets. The power consumption of the pre-production modules has now been measured in a realistic set-up with fully working optical links and a test pattern being sent to the DTM E-links from an FPGA.

All measurements were taken in a lab at 20 °C, with the modules mounted on a prototype PDMDB PCB and force-cooled by a fan. Typical operating temperatures for the modules in these tests were between 20 and 30 °C. The accuracy of the current measurements was estimated to be ± 5mA.

The current consumption of various parts of the modules, as well as the measured and originally estimated power consumption is shown in Table 3. In general, these measurements are in line with the values expected. Helpfully, the overall power consumption of both modules is lower than originally estimated, by a factor of 18% for the TCM and 10% for the DTM. This is illustrated in Figure 6.

It should be noted, however, that the power consumption of the GBT-SCA is nearly five times greater than originally expected. This is because the figures originally used to estimate its power consumption considered only the static power consumption of the chip and not the considerably higher power consumption experienced when its I/Os are active. This also explains why the GBT-SCA was warmer than expected when thermal measurements were taken during operation.

This error is not a cause for concern with respect to the capacity of the power supplies or the cooling of the module as a whole. However, it does significantly affect the likely operating temperature of the GBT-SCA. Fortunately, as described in the Thermal design section on page 4, the existing thermal design appears to be sufficient to keep the operating temperature of the GBT-SCA at around 25 °C.

Table 3: Operating current and power consumption of DTM and TCM

	TCM				**DTM**		
	GBTX	**GBT-SCA**	**VTRX**	**Total**	**GBTXs**	**VTTX**	**Total**
Current on 1V5 supply (mA)	597	0	0	**597**	1140	0	**1140**
Current on 2V5 supply (mA)	0	138	140	**278**	0	197	**197**
Actual power consumption (mW)	896	345	350	**1591**	1710	493	**2203**
Original power estimate (mW)	1452	73	350	**1875**	1914	500	**2414**

Figure 6

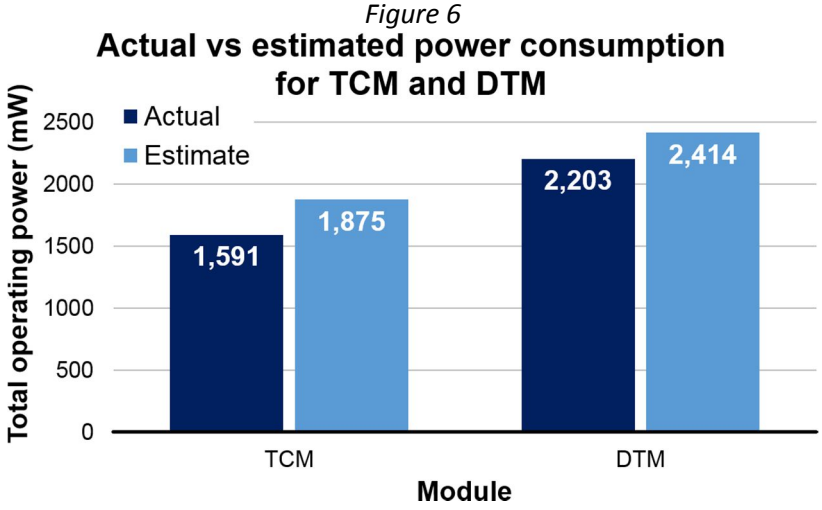

The current consumption of the modules is typically constant during operation, but can change considerably during the steps required to configure them. A typical configuration consists of six steps:

1. Power on
2. Configure TCM (master link) GBTX
3. Configure GBT-SCA on TCM *(provides clock reference to DTMs and TFC to FPGAs)*
4. Configure GBTX0 on each DTM *(provides clock reference to DTM GBTX1)*
5. Configure GBTX1 on each DRM
6. Configure FPGA

The current and power consumption for each module has been measured in each state. The current consumption for the TCM and DTM are shown in Figure 7 and Figure 8 respectively, whilst the power consumption of the TCM and DTM are shown in

Figure 9 and Figure 10 respectively. The last state in each of these figures ‘FPGA configured’ represents the typical operating state of the module.

Note that the power consumption of the GBT-SCA is around 54% higher than its operating value between start-up and when it is configured. It is therefore recommended that the GBT-SCA be configured as soon as possible after it is powered on to minimise its power dissipation.

Figure 7

(a)

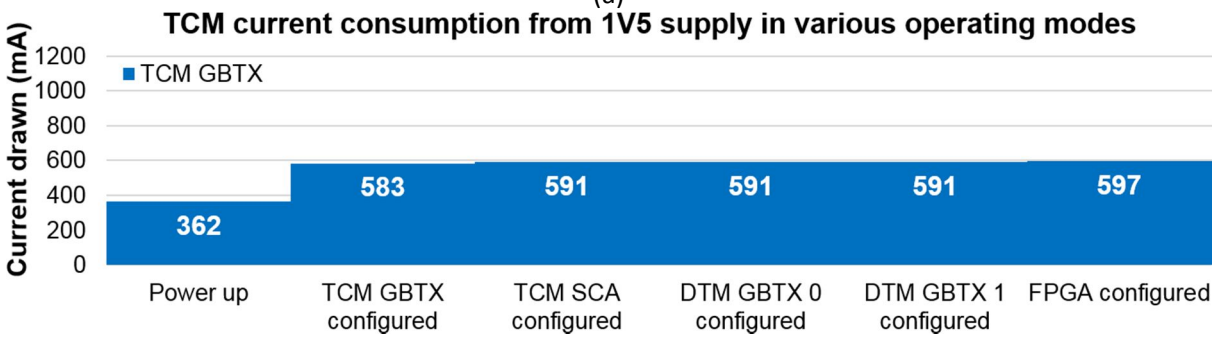

(b)

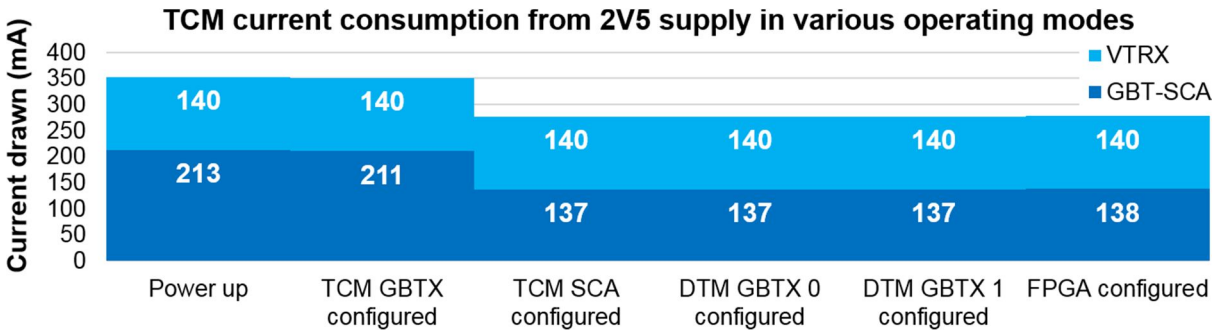

Figure 8

(a)

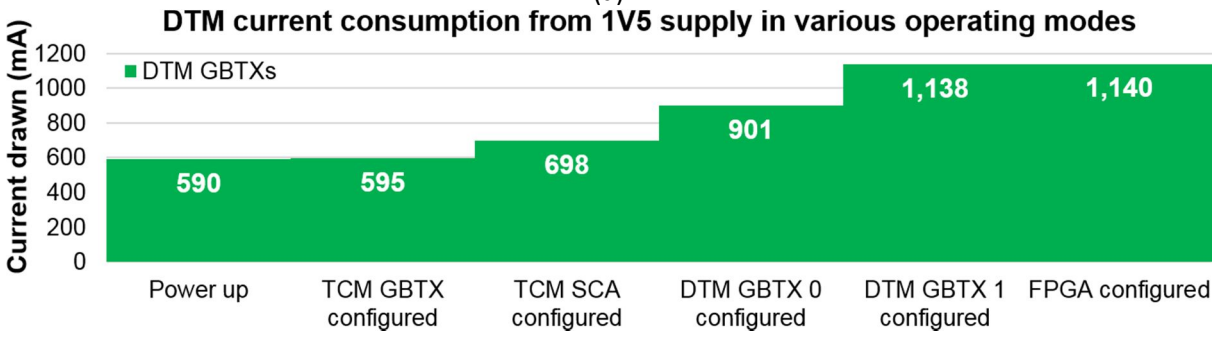

(b)

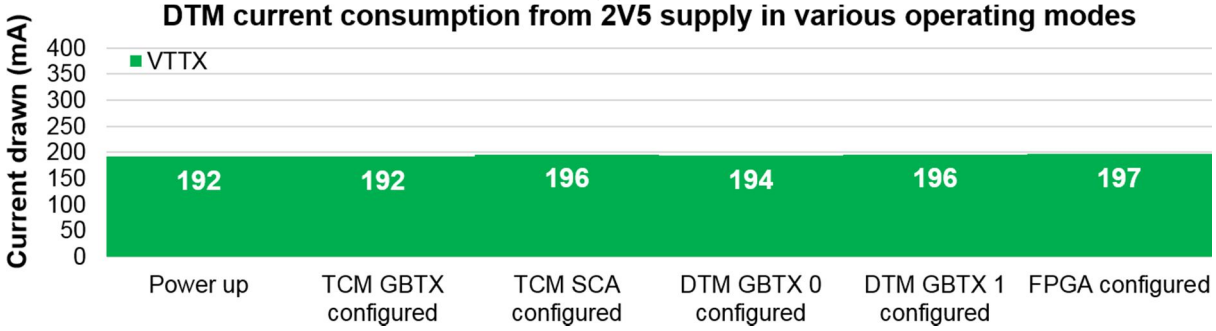

Figure 9

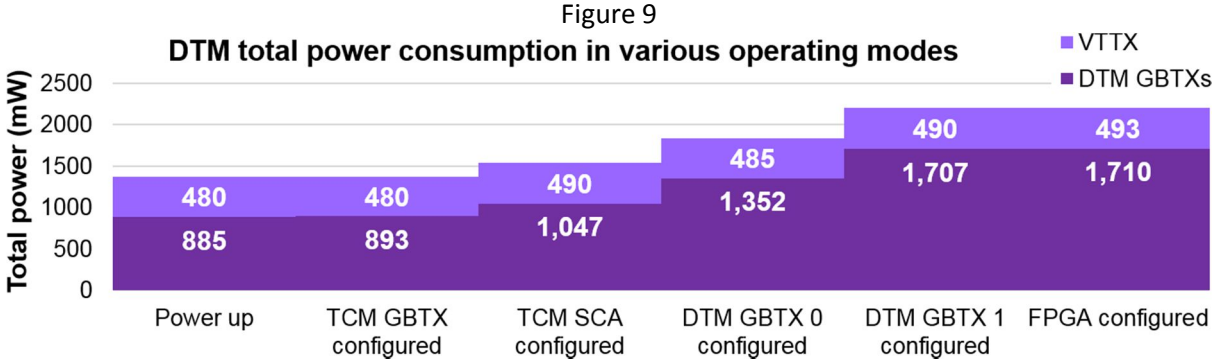

Figure 10

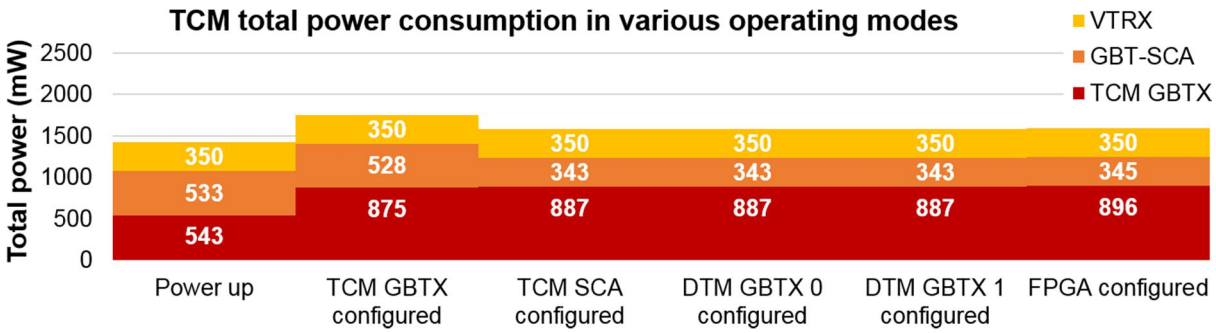

5. Grounding and shielding

Both the DTM and TCM are designed with four signal layers, three ground and shield layers and one power distribution layer. The stackup can be seen in section 8.1.

5.1. High speed signals

High speed signals (40 MHz+) are generally routed as 100R differential striplines on the SIG1 and SIG2 layers, between the GNDT and GNDM, or GNDM and GNDB layers respectively. The exception to this is the 4.8 Gbps TX links on the DTMs, which are routed as differential microstrip lines on the top surface of the PCB, with the GNDT layer acting as its ground plane.

The GNDT and GNDB layers carry power supply return currents, in addition to acting as a ground plane for high-speed signals. To avoid this causing excessive noise on the ground planes, all parts of components that have significant current requirements are extensively locally decoupled (see section 5.3 for further details).

To minimise the size of current loops resulting from decoupling, the GNDT and GNDB layers are cut in several places (see following section for reasoning and implementation). In general, these are arranged such that high speed traces that use these layers as ground planes do not cross these discontinuities. In the few cases where traces do cross these discontinuities, they do so close to one end of the transmission line and for a maximum distance of around 3.3 mm. Based on a rough estimate of the propagation velocity in the transmission line of around $0.5c$ and the rule of thumb that transmission line effects are only significant for lengths greater than $\lambda/20$, we estimate that the effects of these discontinuities should only be apparent for frequencies higher than around 2.2 GHz.

The traces that cross ground plane discontinuities under these circumstances are 40 MHz clock reference and 320 Mbps E-link differential pairs (the 4.8 Gbps links never cross ground plane

discontinuities). Thus, any such clock signals should not see any effects from these discontinuities until at least $55f_0$ – a point at which the magnitude of these and higher harmonics is well below 2% of the fundamental. The quality of the edges of the E-links is not as critical, since they are sampled in the centre of a clock period, but these should not experience any effect from the discontinuities until at least $15f_0$ – a point at which the magnitude of this and higher harmonics is below 6% of the fundamental. We are therefore confident that there are no significant negative effect on signal integrity from cuts made in the ground plane.

The GNDM layer serves as a signal ground plane for both the SIG1 and SIG2 layers. It is therefore kept as continuous as possible and therefore carries no return current. It is only connected to the other ground planes by the vias that link the Megarray mezzanine connector to the other layers - the ground reference for the entire PCB. The continuous nature of this middle ground plane is expected to reduce the severity of any discontinuities experienced by differential pairs on the SIG1 and SIG2 layers caused by cuts in the GNT and GNDB planes.

The PWR layer carries the current supplied by the 1V5 and 2V5 rails of the motherboard to where it is required. It does not act as a ground plane, although it is used to some extent as a shield for signals on the bottom layer.

5.2. Lower speed signals

5.2.1. ADCs and DACs (TCM only)

Although these analogue signals are high impedance, they are also low bandwidth (e.g. <7 ksps for the ADC). These are therefore routed on shielded inner layers from the SCA to the connector, with conductors in close proximity to save space.

5.2.2. GPIOs (TCM only)

The SCA GPIOs are used for functions such as enabling FEBs and SEU detection. They are therefore very low bandwidth (pseudo-DC), so are routed on shielded inner layers from the SCA to the connector, with conductors in close proximity to save space.

5.2.3. Control and monitoring signals

The module control and monitoring signals (such as DATAVALID, TX_READY, TX_DISABLE etc.) are generally low-speed signals that have no specific routing requirements. The exception to this are the two TX_DATAVALID signals on the DTM, which are implemented as 85 Ω microstrip transmission lines, with pads for termination resistors under each GBTX. This is to ensure that the FPGA is able to set the status of each data frame (maximum toggle rate of 40 MHz), without the risk of reflections affecting the integrity of the signal. However, to date, we have not experienced any problems using this signal without the termination resistor.

5.2.4. Communications busses

SPI (TCM only)

The SPI bus is capable of operating at comparatively high frequency (CLK, MOSI and MISO can operate at up to 20 MHz) and its traces pass through a very congested part of the board. Consequently, SPI traces run in very close proximity to other signals. To mitigate this, trace lengths are kept to a minimum. As a result, no two traces run parallel to each other for more than 7 mm and the SPI_CLK does not run near other SPI traces, hence cross-talk is not expected to be a problem.

The SPI clock is a particularly important signal, since it is responsible for synchronising the SPI transactions required to configure the front-end. In the prototype system, problems with configuring the SPI have been identified as being partly due to poor termination in the front-end elementary cells

and partly due to a switching glitch in pre-production SCAs. To mitigate these problems, the SPI_CLK line on the TCM has been fitted with pads for a bypass capacitor, to allow any switching glitches in production SCAs to be shunted to ground before they are able to cause reflections on the SPI_CLK line. The trace has also been implemented as an 85 Ω microstrip transmission line, rather than a simple track, to aid with proper termination.

JTAG (TCM only)

The JTAG traces run on the bottom layer of the TCM, between the SCA and the connector. The length of the JTAG traces on the TCM are short (max. 25 mm). Nevertheless, they run in close proximity to each other for a short distance (around 7.5 mm) between the SCA and connector. Based on estimates from PCB CAD software, the maximum coupling between lines is around 0.9 pF, between TCK and TMS.

Despite this, we have not detected any problems with configuration over JTAG and have been unable to measure any evidence of coupling between JTAG lines on an oscilloscope. Space does now exist to increase the distance between traces in this area, so we will consider increasing the space between these traces in the production boards. However, the small advantages of doing so may be outweighed by the risks of making small changes to an otherwise functioning board.

I2C

By necessity, several of the I2C busses have long parallel traces of up to 50 mm on the modules for the inter-module busses and up to 100 mm for the intra-module GBTX-SCA bus. To minimise cross-talk between data and clock lines where traces are parallel for a considerable distance, signal lines are interleaved with a shield ground line.

5.3. Decoupling

Each power domain of each component on each module is individually decoupled. For example, the GBTX has six power domains, corresponding to different functionality, whilst the SCA has three. Each domain has at least one bulk 4.7 µF decoupling capacitor, as well as a 100 nF capacitor for dealing with higher frequencies. Analogue and high-frequency power domains also have a 10 nF capacitor to further improve high frequency behaviour. In addition, the GBTX VDDIO connections are decoupled by a 10 nF capacitor at each pin.

To prevent interference between different power domains, both the power plane (PWR) and the power ground planes (GNDT and GNDB) have cuts in them to prevent return current from one domain passing across another. For example, it is important that any switching noise from the GBTX IOs is not coupled into the clock manager power domain, where it could cause jitter in the reference clocks.

These cuts are shaped such that supply current from the Megarray connector flows into the centre of each IC, before spreading out from this point to supply each power domain. The ground return currents follow the opposite path. This gives a local supply and ground reference under each IC to further reduce the effects of any power supply noise. The local decoupling of each power domain ensures there is a low impedance path to the local ground through a decoupling capacitor near each power pin, thus minimising any circulating currents between power domains.

The effectiveness of the decoupling for each power plane on the DTM and TCM has been measured using a vector network analyser (VNA). Measurements were made by injecting a signal at the Megarray connector and measuring the attenuation of this signal on the relevant BGA/connector pads at the point where the power domain is connected to the chip or plugin.

The measurements for the TCM (GBTX, SCA and the VTRX) are shown in Figure 11, Figure 12 and Figure 13, whilst the measurements for the DTM (GBTX0, GBTX1 and the VTTX) are shown in ter up to a frequency of 40 MHz.

Figure 14, Figure 15 and Figure 16.

All measurements were taken on boards populated only by decoupling capacitors. Measurements using the FEAST-CLP DC/DC converters (which would show low impedance at low frequency, due to the effect of their feedback loop) were not made, due to the frequency limitations of the DC blocks required to perform the measurement. However, based on their output impedance of 11.3 mΩ[3] and control bandwidth of 150 kHz, we have estimated the low frequency performance of the modules when powered by the DC-DC converters, both with no additional decoupling capacitance and with additional bulk capacitance (220 µF for the TCM and 1000 µF for the DTM) with an ESR of 2 mΩ and an ESL of 10 nH. These are shown as the black and grey lines respectively, between 1 kHz and 1 MHz.

Figure 11

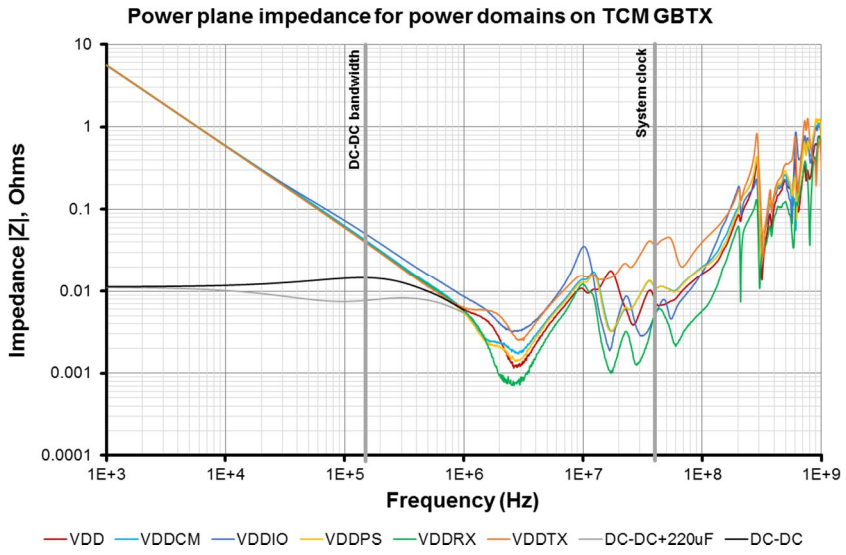

Figure 12

[3] Calculated from output regulation – see datasheet

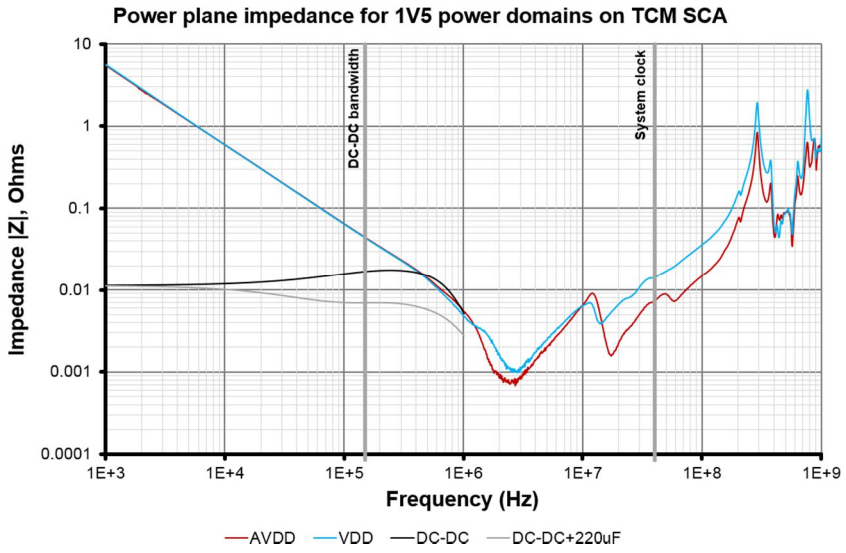

Figure 13

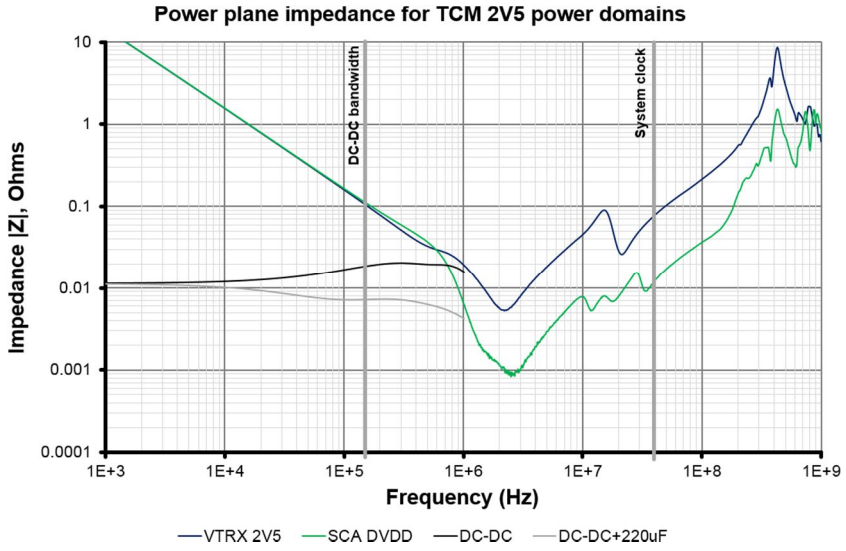

The measurements for the TCM show that the impedance of the power distribution networks (PDNs) for the TCM are generally at or below that of the DC/DC converter between around 500 kHz and 40 MHz. Above 40 MHz, impedance increases roughly by 20 dB per decade, but generally remains below 1 Ω up to a frequency of 1 GHz. This seems adequate, given that higher speed links are terminated in a 100 Ω load.

The exception to this is decoupling for the VTRX, which is less effective at higher frequencies. However, these measurements were taken without the VTRX in place and the optical module has its own decoupling on-board. Additionally, the SFP+ connector will seriously impair any high frequency decoupling on the TCM.

Based on the estimates of the PDN's impedance when including the DC/DC converter, it appears that an additional 220 µF bulk capacitance on the PDMDB would be valuable for keeping the PDN's impedance at or below that of the DC-DC converter up to a frequency of 40 MHz.

Figure 14

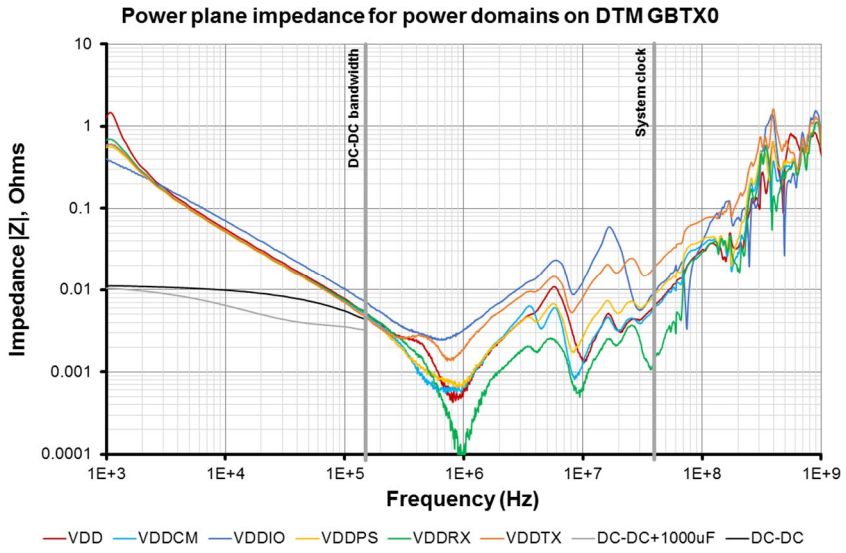

Figure 15

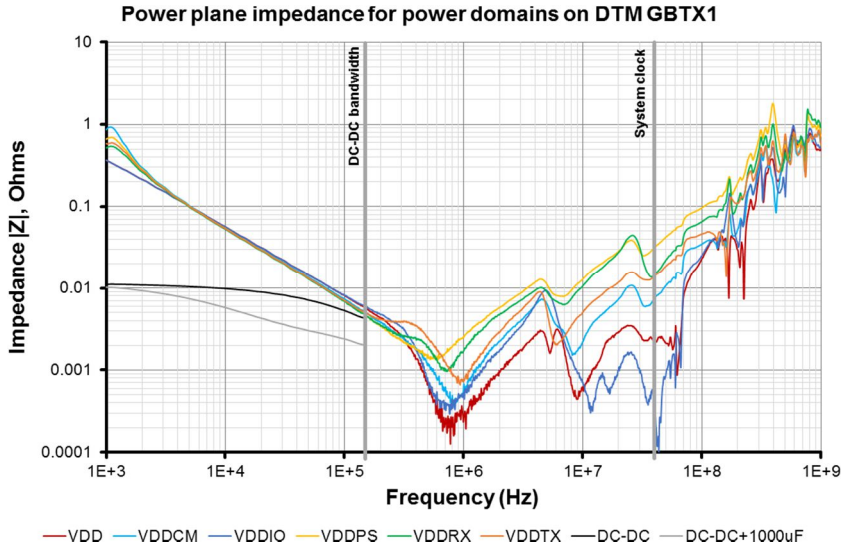

Figure 16

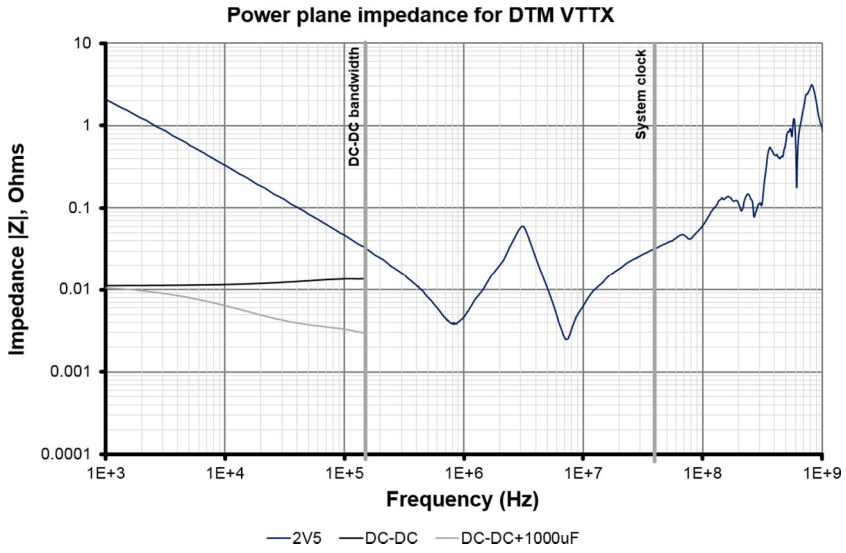

The measurements for the DTM show that the impedance of the power distribution networks (PDNs) for the TCM are generally at or below that of the DC/DC converter between around 50 kHz and 5 MHz. Between 5 MHz and 80 MHz, the impedance is generally below 50 mΩ, whilst above 80 MHz, it increases roughly by 20 dB per decade, but generally remains below 1 Ω up to a frequency of 1 GHz. As with the TCM, this seems adequate, given that higher speed links are terminated in a 100 Ω load. The greater number of 1V5 bulk decoupling capacitors on the DTM is evident from the shift of the minimum impedance to around 800 kHz, compared to around 3 MHz with the TCM.

Based on the estimates of the PDN's impedance when including the DC/DC converter, it appears that an additional 470 - 1000 µF bulk capacitance on the PDMDB would be valuable for keeping the PDN's impedance at or below that of the DC-DC converter up to a frequency of 5 MHz.

6. Signal integrity

6.1. Clock jitter

The reference clocks for the PDMDB are derived from the clock recovered from the incoming data on the RX link of the GBTX on the TCM. This master reference clock is then used to synthesise six reference clocks using the phase/frequency programmable clock outputs of the master GBTX on the TCM. Of these, a single reference clock is routed to each of the three DTMs on the PDMDB. The other three clocks are routed to the connector but are unused.

It is essential that these clocks are a reliable reference for the rest of the system, since they are used by each DTM to synthesise the clocks used by each FPGA to synchronise and latch the incoming signals from the front end. In addition, they are multiplied up by the PLL in each GBTX on each DTM to generate the clock required to drive the two 4.8 Gbps TX links on each DTM. Thus, although the reference clock frequency is comparatively low, its jitter performance is especially important.

The cycle-cycle jitter measured on the 40 MHz reference clocks to the DTM is shown in Figure 17. The reference clock for GBTX 0 in this case is taken directly from the phase/frequency programmable clock on the TCM's master GBTX, whilst the reference clock for GBTX 1 is generated by the programmable clock on GBTX 0 on the same DTM. The reference clock for GBTX0 was measured in two places: at the TCM connector, with the TCM unplugged from the PDMDB and with the differential pair terminated at the connector; and at the RefClk test pads on the DTM near GBTX0. The reference clock for GBTX1 was measured at the RefClk test pads on the DTM near GBTX1.

The jitter measured at both the TCM output and for the GBTX1 RefClk have very close to Gaussian distributions, indicating that the jitter is almost entirely random. However, its magnitude (approximately 70 – 80 ps_{pk-pk}) appears to be almost twice the 41 – 45 ps quoted in the GBTX manual. It is currently unclear whether this is a real effect, or a function of the measurement method used. The values quoted in the GBTX documentation appear to measure the jitter at the first rising edge, nominally 25 ns after the trigger[4]. However, we find that measuring jitter in this way gives a histogram approximately half the width of that produced by the oscilloscope's internal cycle-cycle jitter measurement function. We are currently seeking expert opinion to clarify whether our measurements match the expected jitter of the GBTX.

When the reference clock from the TCM is connected to the DTM via a trace on the PCB, the cycle-cycle jitter distribution becomes flatter and more deterministic (two distinct peaks are visible). This does not significantly increase the peak-peak jitter, but suggests that the signal integrity is reduced by

[4] Based on figures presented in:
https://indico.cern.ch/event/267408/attachments/477815/661090/eseSeminar26november.pdf

the act of connecting the clock output from the TCM through the PDMDB to the DTM. It is possible that this may be due to reflections due to discontinuities in the transmission lines on or between the TCM and DTM and we will continue to investigate this possibility. Interestingly, these data do not support the suggestion that the clock jitter in the GBTX1 reference clock are the cause for the slightly greater error rate observed for GBTX1, (section 6.3.2) so further investigation is required to understand this effect fully.

Whilst cross-talk is discussed in section 6.3.3), our jitter measurements appear to rule out any significant cross-talk between clock signals; the same characteristic is seen even if neighbouring clock signals are disabled, offset by 90 or 180 degrees and their frequency changed.

Despite the as yet unexplained differences between GBTX jitter specifications and measurements, as well as the deterministic component of the GBTX0 input clock jitter, the overall peak-peak jitter is below 40% of the 208 ps unit interval for the 4.8 Gbps DTM TX links. This suggests that the reference clocks as measured should provide an adequate reference for the GBTX PLLs to generate the high speed link reference clocks.

Figure 17

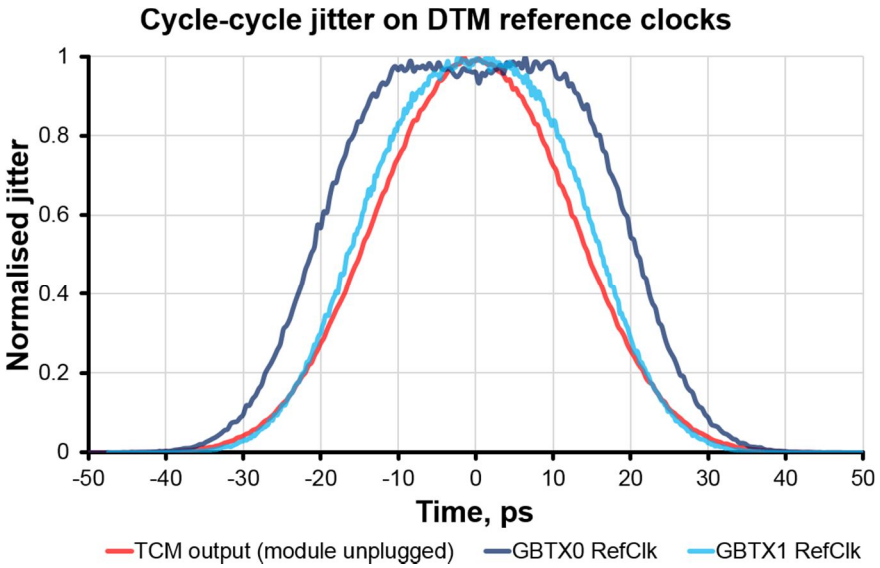

(26k samples were taken for the GBTX 0 measurement and 83k samples were taken for GBTX 1)

6.2. Signal spectra

Knowing the frequency content of clock signals is important for establishing the likely severity of effects such as crosstalk and reflections in transmission lines. This is particularly important for the 40 MHz reference clocks of the DTM, which are fed into an internal PLL to generate the clocks to drive the 4.8 Gbps links. Thus a small percentage jitter on these clocks has the potential to corrupt data on these links. Data lines tend not to be as severely affected, since they tend to be sampled in the middle of a clock period. In addition, their power spectrum is not as clearly defined as for a clock signal.

Time domain and frequency domain plots of the SLVS clock signals are shown in figures Figure 18 and Figure 19 respectively. For reference the LVDS data signals from the FPGA to the DTM GBTXs are also plotted in figure Figure 18.

The clock signals measured in figure Figure 18 are generally clean and whilst there is some evidence of reflections after each transition, these are small (<25%) relative to the peak to peak value of the waveform and do not appear to be a cause for concern.

The frequency content of the clock signals shown in figure Figure 19 appears to drop a little faster than predicted by the Fourier terms for an ideal square wave, with the result that there is little significant high-frequency content to the clock waveforms above around $13f_0$ (520 MHz for the GBTX reference clocks and 2.08 GHz for the FPGA reference clocks)

Figure 18

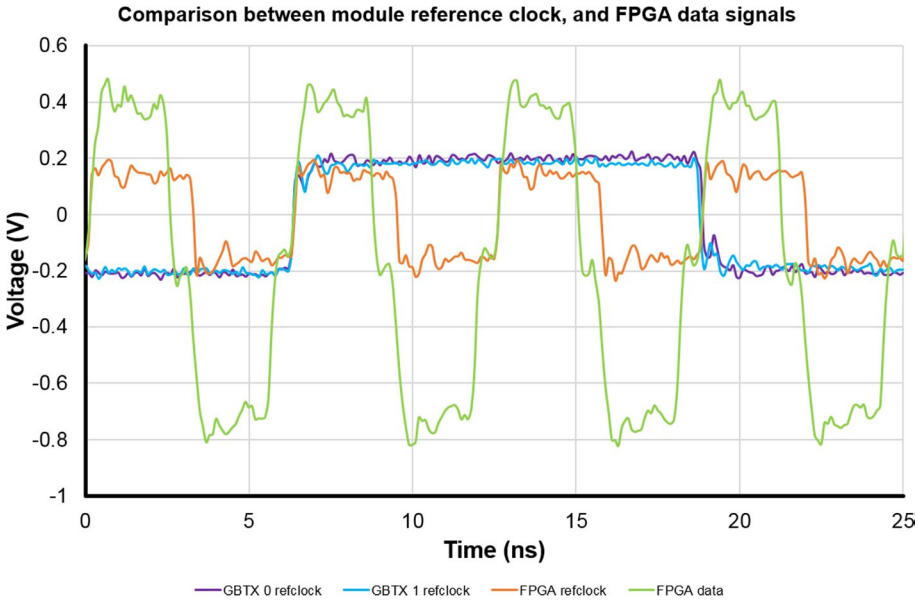

Figure 19

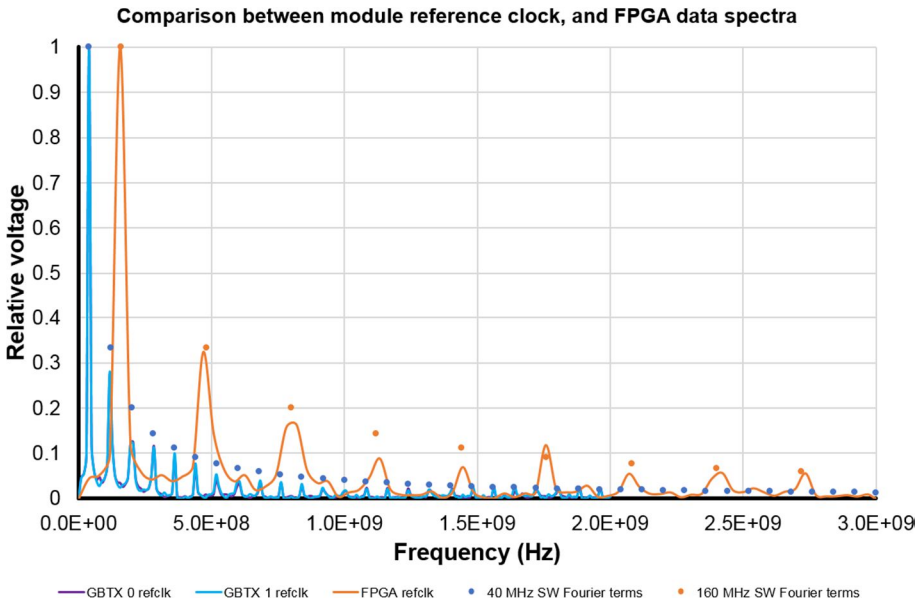

6.3. Bit error rates

The performance of the high speed links on the modules was characterised by using a serial data analyser to measure the differential electrical signals on the copper transmission lines, as a well as a

loopback test to produce a statistical eye diagram of the complete optical and electrical links. It is particularly critical to understand the BER high speed links of the DTM, since unlike the high speed links on the TCM, there is no forward error correction.

6.3.1. Performance of copper high-speed links

The performance of the 4.8 Gbps copper links on the modules was measured using a LeCroy SDA6000 serial data analyser.

For the DTM, the bit error rate was measured by modifying a pre-production VTTX module by soldering in a connector for a differential probe at the termination of the transmission line. This was found to have the smallest effect on the signal quality compared to other techniques such as using a contact probe with adjustable pins.

For the DTM, modifying a VTRX in a similar way was not practical, since only one of the links is terminated at the VTRX and a working VTRX is required to provide the clock reference for the system. Instead, point contact probes were used on the solder pins at the base of the 20-pin SFP+ connector on the module. This slightly resulted in a slightly degraded signal, which should be taken into account when interpreting the results.

Eye diagrams, bathtub curves and jitter histograms for the two high speed links on both the TCM and DTM are shown in figures Figure 20 to Figure 23 . Estimates of jitter based on these measurements are given Table 4.

Figure 20: Eye diagram for TCM GBTX RX copper link after 1.1 M eye samples

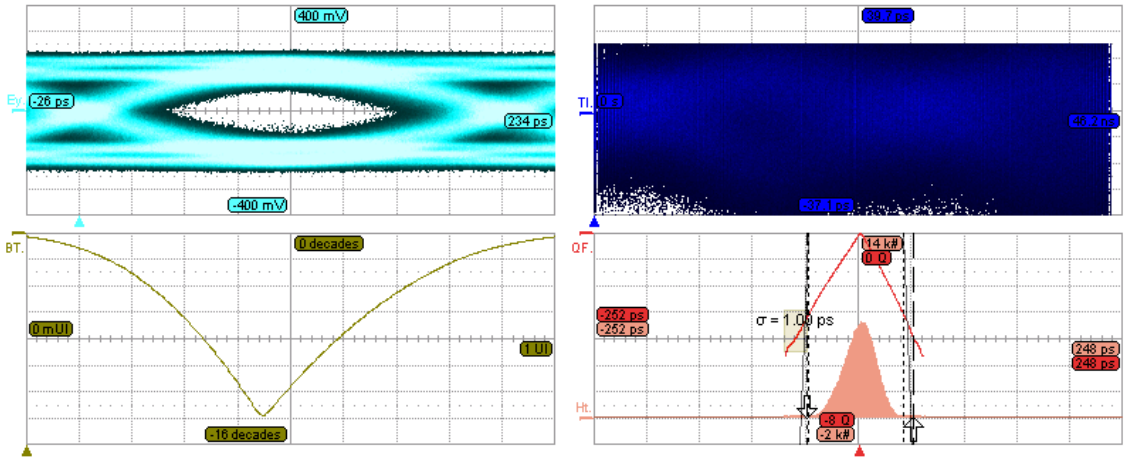

Figure 21: Eye diagram for TCM GBTX TX copper link after 1.1 M eye samples

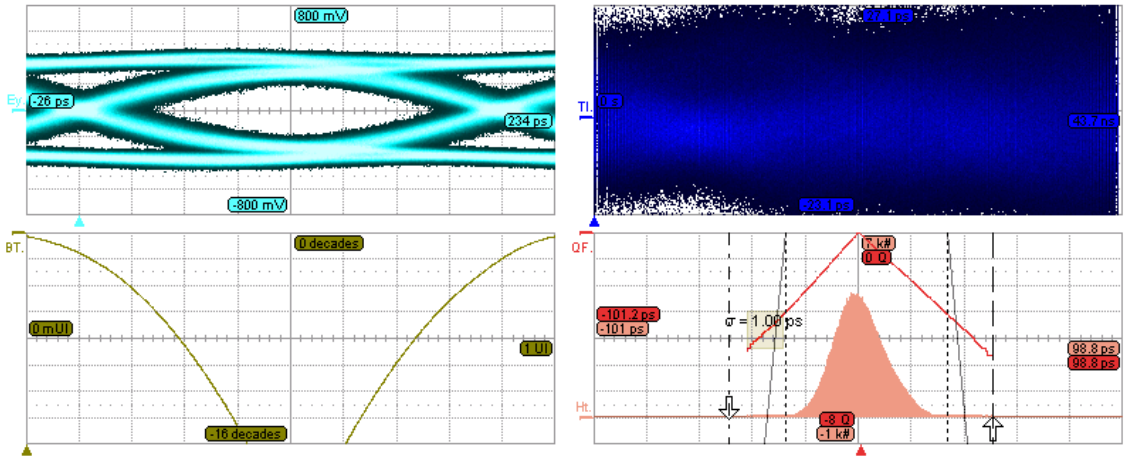
800 mV
-26 ps
234 ps
-800 mV
0 decades
0 mUI
1 UI
-16 decades
-101.2 ps
-101 ps
σ = 1.00 ps
98.8 ps
98.8 ps
-8 Ω
-1 kΩ

Figure 22: Eye diagram for DTM GBTX0 TX copper link after 1.7 M eye samples

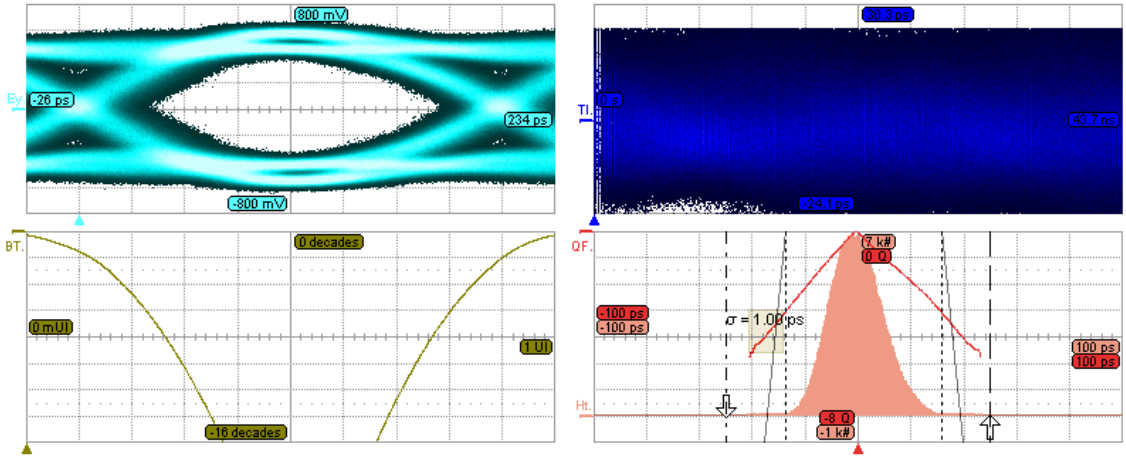

Figure 23: Eye diagram for DTM GBTX1 TX copper link after 1.2 M eye samples

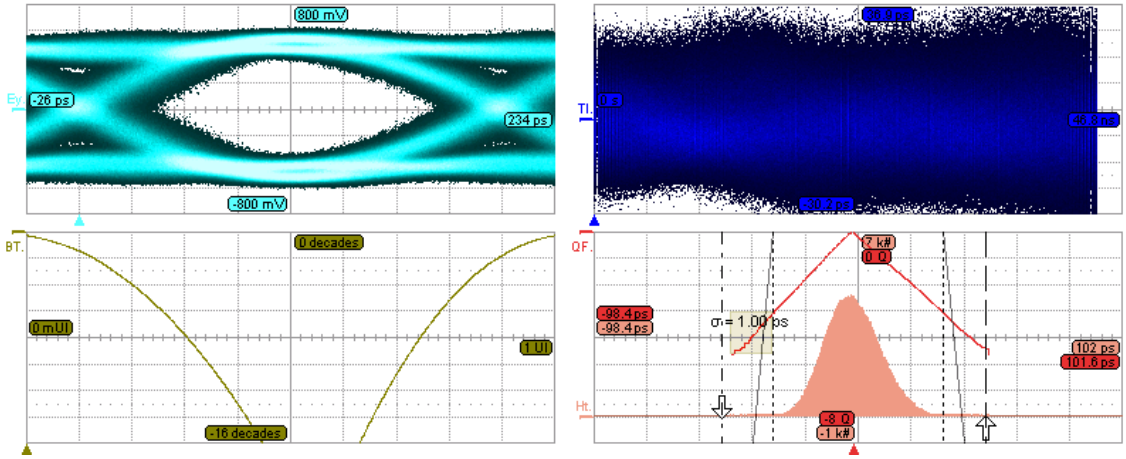

Table 4: Jitter and bit error rates for high speed links. (Note: 1 unit interval (UI) is 208 ps). DTM measurements were made using a modified VTTX, allowing 'scope probe to be attached directly to terminations. TCM measurements are made using needle probes on the SFP+ connector.

Module	Link	Mean random jitter, Rj (ps)	Mean deterministic jitter, Dj (ps)	Mean total jitter @ BER = 10^{-12}, Tj (ps)	Estimated BER from bathtub curve (for central 20% of UI)*
TCM	GBTX RX	13.6	0	192	$<10^{-8}$ †
	GBTX TX	10.3	0	145	$<10^{-10}$ †
DTM	GBTX0 TX	9.1	0.3	128	$<10^{-16}$
	GBTX1 TX	10.4	0.3	146	$<10^{-16}$

*A BER of 10^{-16} corresponds to 15 errors per year

†Measurement technique (needle probes on SFP+ connector) is likely to lead to higher BER measurement, due to length of unterminated line attached to transmission line.

The jitter measurements of the 4.8 Gbps exhibit almost purely random jitter, which suggests that the links are relatively free from reflections, crosstalk and other distortions. For the DTM, the eyes are fairly wide open, indicating a good signal to noise ratio and sufficiently low jitter.

The TCM measurements have eyes that are less open, suggesting a lower signal to noise ratio. Whilst the jitter is slightly higher, it is not severely worse than for the DTM links. Part of this difference can be explained by the probing technique used on the TCM module, which can be expected to increase the noise on the measurement. It is also expected that the signal to noise ratio could be improved by increasing the drive strength of the line drivers on these links.

Despite the difference in the measurements between the DTM and TCM high-speed links, we do not believe there is a cause for concern about the TCM links for the following reasons:

1. It appears likely that signal-noise ratio can be increased by changing the device drive strengths, which can be done through programmable registers in the GBTX and VTRX.
2. The measurement method used is expected to exaggerate noise in the link.
3. The TCM data links use forward error correction (FEC). Under most circumstances, a bit error should not lead to a data error.

6.3.2. Statistical eye diagrams

The Kintex 7 FPGA on the 'µDAQ' test board is capable of generating statistical eye diagram based on observing a test pattern generated in the PDMDB FPGA.

The eye diagrams generated by this tool are coloured images, with a width representing 1 UI. Areas within the UI with 100% error are coloured white, whilst areas with 0% error are coloured black. Areas with intermediate numbers of errors between 0 and 100% are coloured in a spectrum of blue, green, yellow, orange, red and pink. The larger the black area in the centre of the eye diagram, the better the signal-to-noise and jitter of the link.

The data in these eye diagrams clearly includes contributions from many components outside the scope of this PRR. However, useful information can be gained from looking at the differences in between them. In particular, anecdotal evidence from testing in the system suggested a slightly higher error rate in the data from GBTX 1 on the DTM, compared to GBTX 0.

When comparing statistical eye diagrams for the data loopback between GBTX 0 and GBTX 1 on several DTMs, it was noted that the eye was considerably more 'closed' for GBTX 1. Since measurements on the high-speed links showed very good performance and low bit-error rates, it was hypothesised that excessive jitter on the clock reference for GBTX 1 (which is supplied from GBTX 0) could be a possible cause.

It was noted that a pair of 0.1" pitch test pads on the input reference clocks to both GBTX 0 and GBTX 1 on the DTM, were a possible cause of impedance mismatch and hence reflections. Statistical eye diagrams were therefore plotted for a loopback using DTM GBTX 1, with these pads in place as designed and with the pads drilled away (Figure 24). The board used for this test was suspected to have worse than average performance on GBTX 1.

From these plots, it is clear that removing the test pads resulted in a measurable improvement to the bit error rate of the modules when operating in a loopback system. This suggests that the GBTXs may be more sensitive to clock quality than previously realised, so we are currently looking at measures to mitigate this, such as enabling the internal PLLs to filter the signal. Despite this improvement, it remains noticeable that the bit error rate is still lower when GBTX 0 is used for loopback, rather than GBTX 1. As yet, we are unsure of the cause of this and will continue to investigate this issue.

As a result of these measurements, we propose to remove the 0.1” test pads from the production boards and instead replace them with 0402 size test pads, connected to the signal traces with small vias (0.15 mm finished hole size).

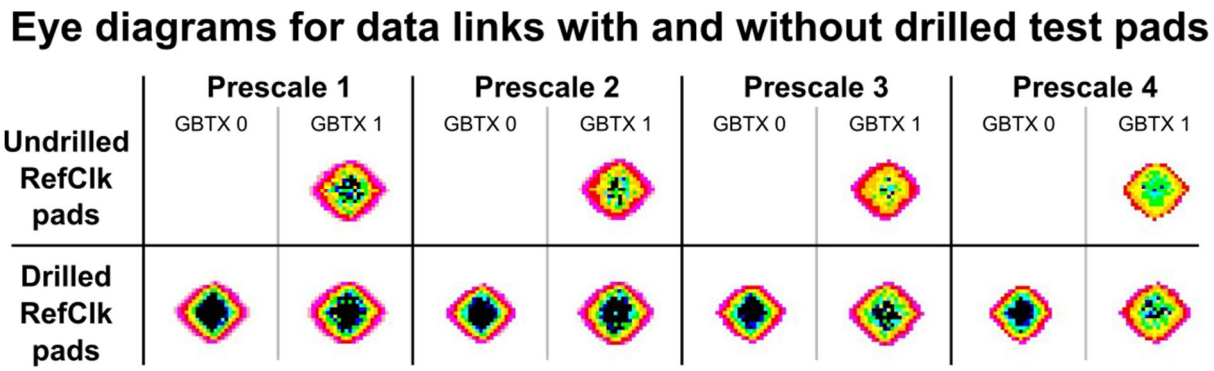

Figure 24: Statistical eye diagrams for loopback through DTM GBTX0 and GBTX1, with an without test pads on RefClk lines drilled out.
***Note:** Prescale is a parameter used in optimising link performance. The use of different prescale factors is intended to show that these results are repeatable over a wide operating range.*

6.3.3. Crosstalk measurements

During the EDR, concern was raised that the ground connections in the Megarray connector between the TCM and the motherboard were arranged to optimise power return and that the reference clock signals were therefore not shielded by ground connections from adjacent analogue, digital and clock lines.

To investigate whether this arrangement was likely to cause problems with crosstalk, the crosstalk between pairs of differential connection in various arrangement on a mated 100-pin Megarray connector in various arrangements was measured. To do so, each differential pair was terminated on one side of the connector, and the ‘aggressor’ differential pair connected to port 1 of a vector network analyser, whilst the ‘victim’ differential pair was connected to port 2. By measuring S_{21}, the amount of crosstalk could therefore be directly measured.

The pins on the Megarray connector are rather flat, and therefore different amounts of coupling may be expected between adjacent pins that are ‘long-side coupled’ (wide edges facing, but fairly widely spaced) compared to those that are ‘short-side couple’ (short edges facing, but close together)

The different arrangements of aggressor and victim differential pairs investigated is shown in

Figure 25. The percentage of the voltage coupled into the victim differential pair, relative to the amplitude of the voltage on the aggressor is shown over a frequency range of 10 MHz to 1 GHz in Figure 26. The spectrum of the GBTX reference clock is also plotted on this figure.

These measurements made it possible to order each arrangement of differential pairs from that which causes the most cross-talk to that which causes the least. This is shown in in

Figure 25. Unfortunately, the arrangement chosen for the clock signals on the TCM connector (adjacent short-side coupled pins – yellow trace on Figure 26) is the worst arrangement for crosstalk.

Nevertheless, Figure 26 indicates that crosstalk at f_0 and $3f_0$ is comfortably below 1% and remains below 5% until at least $13f_0$. Therefore, we believe that although the arrangement of these clock signals is not ideal, we are not aware of any particular problems that can be attributed to this problem and there is little to be gained from reconfiguring the connector at this stage.

Figure 25

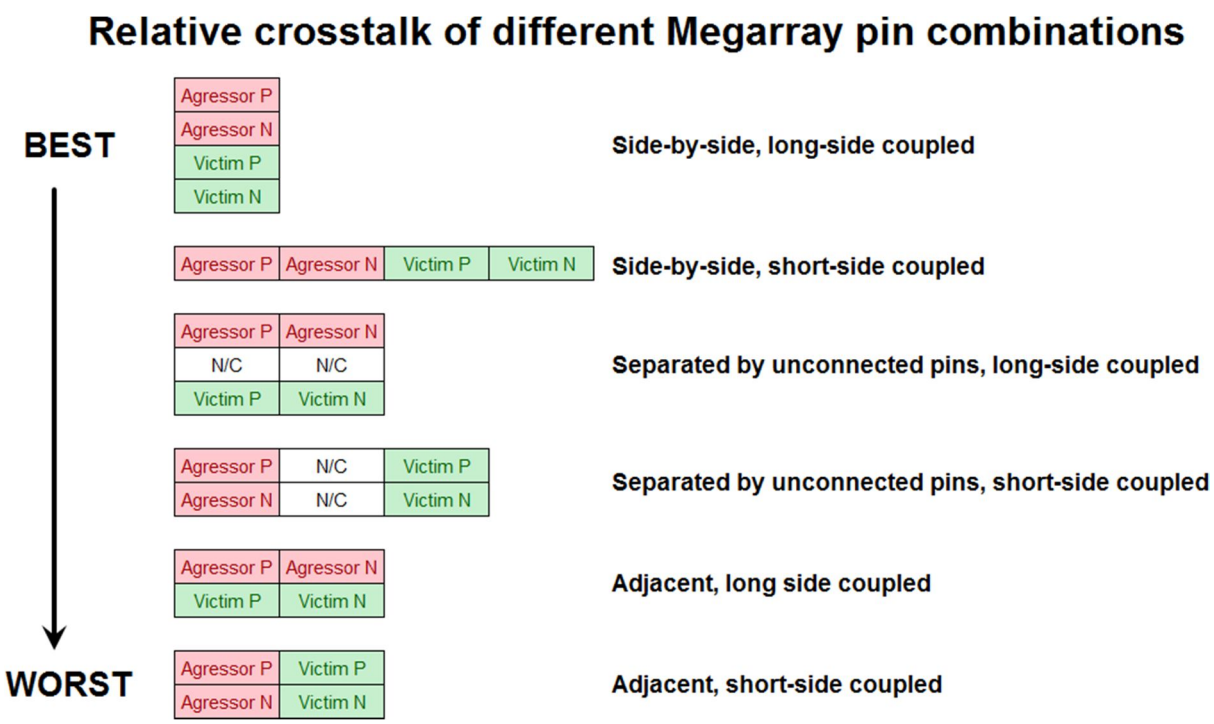

Figure 26

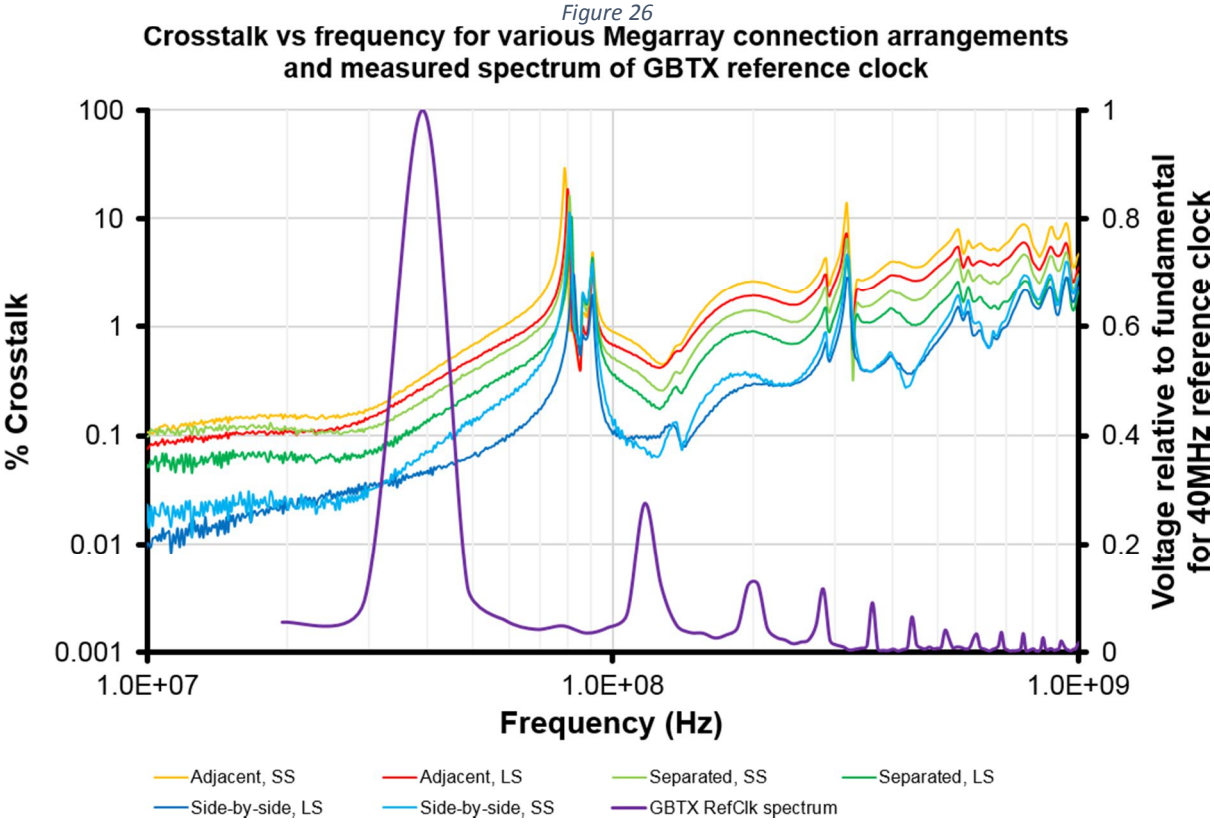

7. Pinouts

The pinouts for the TCM and DTM connectors are generally unchanged from those presented in the EDR. The exception is pin C4 on the TCM connector, which will be changed from NC to CONFIGSELECT in the production boards. This has not yet been implemented on any physical PCBs. See section 10 for more details. The pinouts for the TCM and DTM are shown in Figure 27 and Figure 28 respectively.

Figure 27

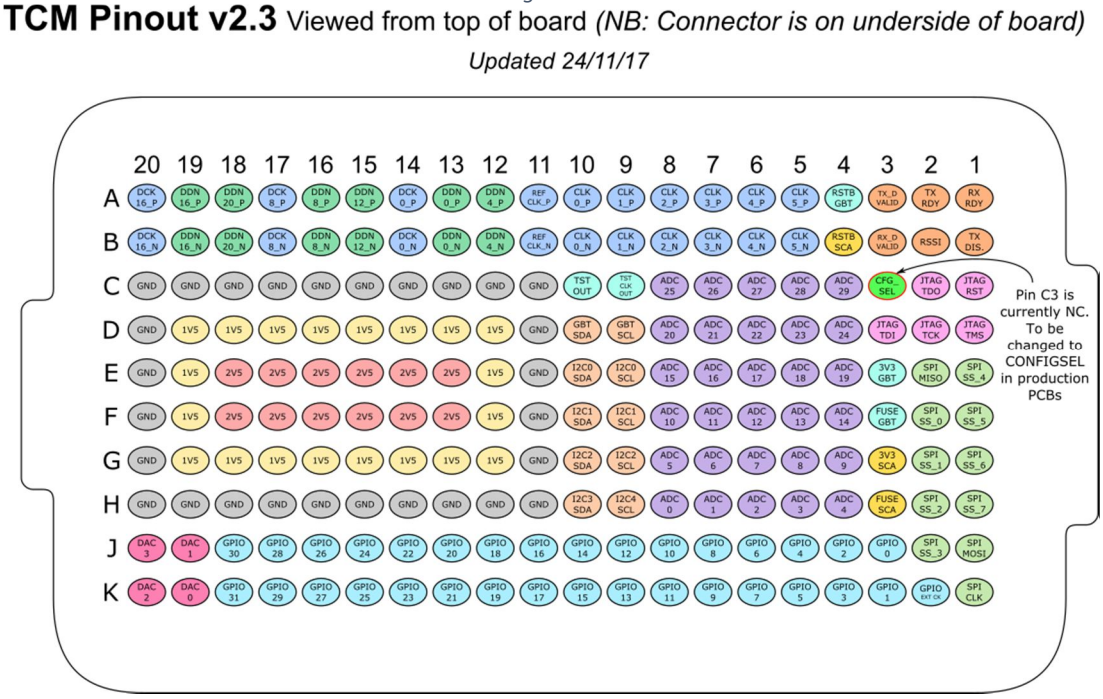

Figure 28

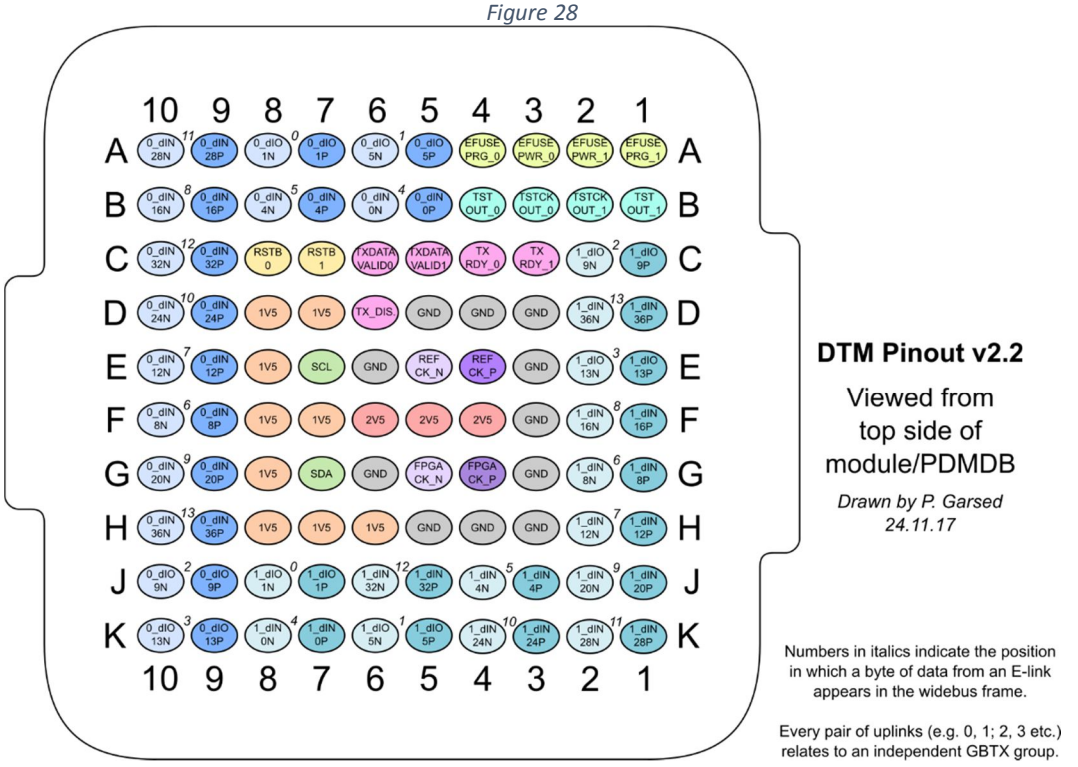

8. Manufacturing and costs

8.1. Stack-up

The module PCBs consist of 8 layers, of which 4 are ground/power layers. The outer layers are to be finished with electroless nickel gold (ENIG), which has good oxidation resistance and is very well suited to reliable soldering of high-density BGAs.

FR-4 has been used for development prototypes for simplicity and cost. However, halogen-free materials will be used in production boards, due to CERN safety requirements. Typically, halogen free PCB substrates have slightly higher relative permittivity than FR-4, so tweaks to the width and spacing of high speed lines will be required to ensure the intended characteristic impedance is maintained.

The stack-up of the current module PCBs is shown in Figure 29. The spacing between layers will be maintained as far as possible in the production boards, but will be dependent on the capability of the manufacturer. Any changes will be accounted for by modifying the width and spacing of high speed lines as necessary. Stack-ups with very asymmetric layer spacing will not be used.

If possible we will have the production boards manufactured with a red soldermask for the TCM and a black soldermask for the DTM. This is intended to make the modules immediately distinguishable and will also match the soldermask colour of the VTRX and VTTX modules, to minimise the possibility of the wrong optical modules being used on the wrong PCB.

Figure 29: Stack-up for DTM and TCM PCBs

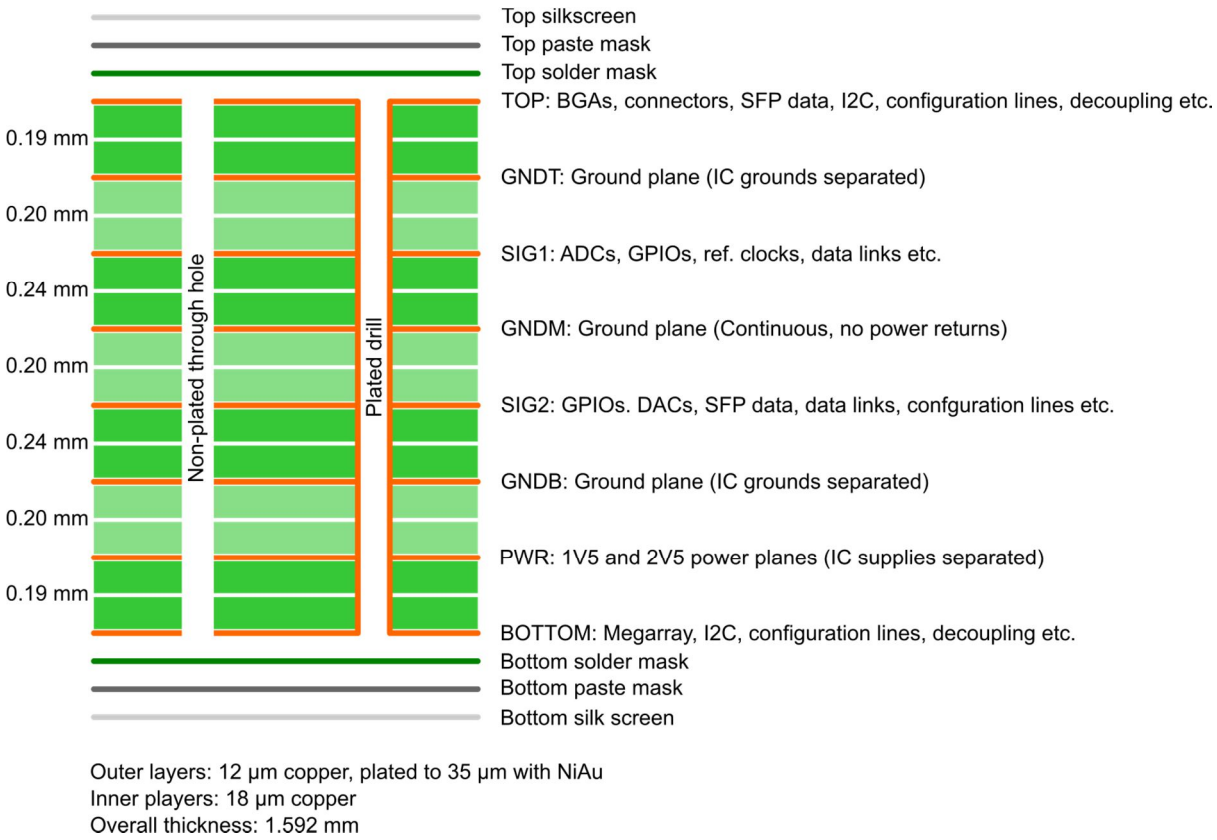

8.2. Manufacturing

Several manufacturers have been approached for budgetary quotation for both manufacturing the module/PDMDB PCBs and for populating them with components. We requested quotations for manufacturing both FR-4 and halogen-free PCBs from PCB manufacturers. On average, we were

quoted 11% more for zero-halogen PCBs, which equates to an overall increase of less than 5% in the cost per board once assembly is also taken into account. We therefore propose to make all the PCBs from zero halogen material, since the increased cost (typically £1.25k) is negligible compared to the overall cost of manufacture.

It should be remembered that all quotations were explicitly for budgetary purposes, therefore the final price may be expected to change and small cost variations between manufacturers should not be given great emphasis.

8.2.1. PCB manufacture

Budgetary PCB manufacturing costs for five manufacturers we approached are shown in Figure 30. Of these, we have the most confidence in Wrekin, Graphic and Amphenol Invotec for delivering PCBs of sufficient quality. Of these, Wrekin manufactures off-shore (although a more expensive on-shore service is also available), whilst Graphic and Amphenol Invotec manufacture in the UK.

We feel that Clarydon is possibly too small for the quantities we require, whilst Argus appears to rely heavily on subcontracting, such that we do not have confidence in their ability to oversee the entire production process.

We propose to manufacture a production quantity of TCM PCBs from Wrekin in the first instance. This is a comparatively small production run and will allow us to assess the quality of Wrekin's production process. If quality proves adequate, orders for the DTM, then the PDMDB will be placed with the same company. If quality is insufficient, we will seek to approach our second choice of manufacturer, most likely Amphenol Invotec.

Whichever manufacturer is chosen, we will request flying probe testing of all boards, to ensure connectivity of the bare boards.

Figure 30

Relative total costs of halogen free PCB manufacture (DTM + TCM + PDMDB) for various manufacturers

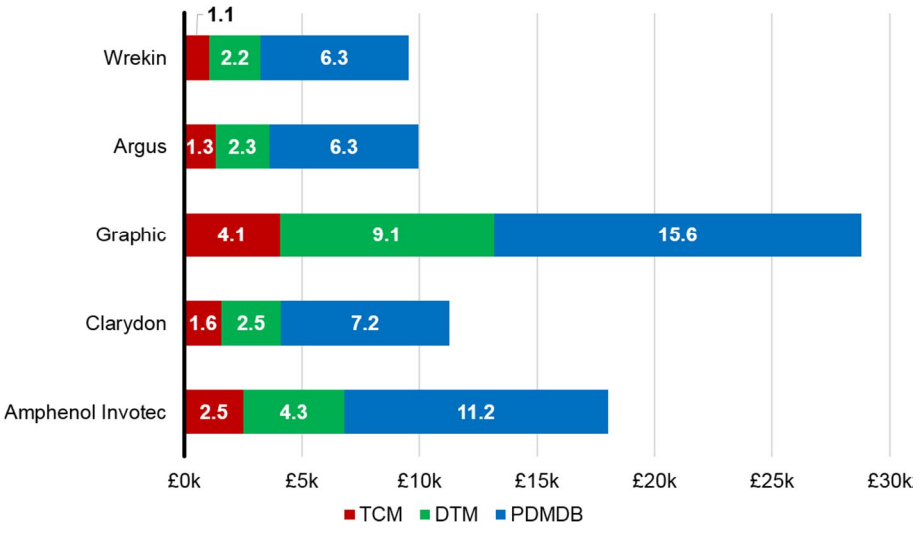

8.2.2. PCB assembly

Budgetary PCB assembly costs for three manufacturers we approached are shown in

Figure 31.

It is essential that the chosen assembly house is able to correctly assemble all the PCBs without mistakes such as rotated components that could cause the loss of irreplaceable one-off radiation hard components. We therefore propose that we agree with the chosen manufacturer a production schedule where boards are assembled in batches, allowing us time to check each batch for correct functionality before triggering assembly of the next batch. We will also require the assembly house to use automated optical inspection on all assembled boards to minimise the risks of mistakes.

Figure 31

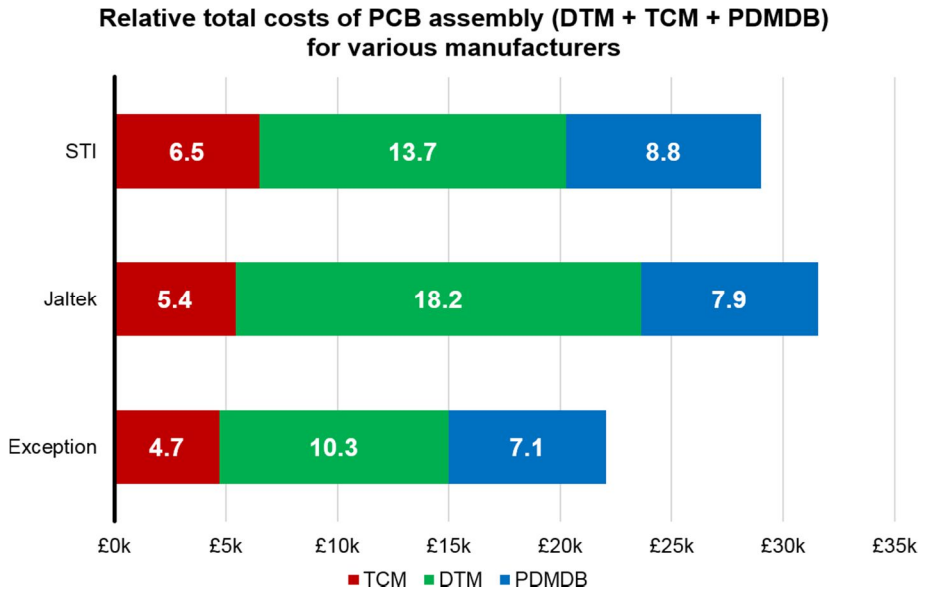

9. Testing strategy

Since the TCM and DTM will be manufactured before the PDMDB, they will be tested separately. Once the PDMDBs are complete, the modules will be attached semi-permanently to these boards and the PDMDB assemblies tested as a system. Here, we describe only the testing of the modules as individual components.

Two types of testing will be performed:

1. Manufacturing quality tests will be performed on a sample of boards, to check that the electrical performance of the boards as manufactured meets expectations. These will often involve specialist test equipment, making it impractical to test all boards in this way. For example, the signals on the clock and data lines of a communications bus may be measured to check their signal integrity.
2. Production tests, carried out on all boards. These will check that all functions of the boards work as expected, but without measuring the details of every signal. For example, a communication bus will be tested by checking that data can be transmitted and received as expected, without measuring the signals on the data and clock lines. A specific test PCB will be manufactured for the production testing of the modules. This is currently being designed.

9.1 Manufacturing tests

Manufacturing quality tests that will be carried out on a sample of boards include:

- Assembly X-ray inspection
- E-link impedance and eye diagrams
- 4.8 Gbps copper link impedances and eye diagrams
- Optical eye diagram tests
- I2C, SPI, JTAG signal integrity tests (edges, crosstalk, timing etc.)
- GPIO limitations test (TCM only)
- Power plane impedance measurement
- Reference clock amplitude frequency and jitter measurements
- RESETB rise/fall time test
- TCM E-fuse burning procedure test
- Long term loopback signal integrity test: BER and eye diagram with high statistics

9.2 Production tests

Production tests that will be carried out on all boards include:

- Automated Bare board optical and electrical tests*
- Assembly optical inspection*
- Read unique board production barcodes
- Check power consumption on 1V5 and 2V5 lines in various operating modes
- SCA E-link functionality test (TCM only)
- RESETB functionality test
- Module I2C tests
- Read unique electronic board IDs and link with barcodes
- DTM reference clock frequency (and phase?) test
- VTxX I2C tests
- JTAG functionality test (TCM only)
- Test pattern functionality test (TCM only)
- Downlink eye diagram/BER test
- SPI functionality test (TCM only)
- Analogue functionality test (ADC and DAC, TCM only)
- GPIO test (TCM only)
- Test GBTX/SCA control signals
- Check RSSI for input optical signal (TCM only)
- Burn TCM E-fuses (TCM only)
- E-fuse burning success test (TCM only)

All tests, except those marked with * will be performed on the module test PCB

9.1 Tracking

Boards will be tracked in a database containing all production testing measurements. Each board will be identified by a QR code sticker or silkscreen printing. This will be cross-referenced to a numerical ID derived from the unique GBTX ID that can be accessed over the I2C link.

10. Final modifications before manufacturing

Although the pre-production modules were intended to be the final design, a few issues have come to light during final testing, which we intend to correct in the production boards. These are:

- **DTM and TCM: Enlargement of stand-off mounting hole:** The mounting hole under the VTxX modules was originally designed for M2.5 screws, but will now be used with a clip-in 4 mm high plastic stand-off (Ettinger 07.91.340) requiring a hole diameter of 3 mm. The hole will be increased in diameter from 2.7 to 3.0 mm, with no additional changes required.
- **DTM and TCM: Transmission line impedance:** Once the final stack-up and board materials are known for certain, the track width and spacing of transmission lines will be tweaked to ensure the correct characteristic impedance is obtained.
- **DTM and TCM: Modification of VTxX connector paste mask:** When assembling PCBs in Cambridge, we have observed that the 20 pin SFP+ connector for the VTxX has a tendency to displace solder paste during mounting, resulting in occasional short-circuits between pins. We will double-check the manufacturer's recommendations for the size of the paste mask openings for this connector and consider using two smaller squares of solder paste at either end of each pin, rather than a single long, thin rectangle of solder paste, in an attempt to reduce this problem.
- **DTM: Removal of 0.1" headers on RefClk lines:** We believe these may be degrading the reference clock signals, so these will be removed and replaced with surface mount test pads.
- **TCM: Removal of C36:** This 0402 10 nF decoupling capacitor is one of four (1x4µ7, 1x100n, 2x10n) on the SCA analogue power supply AVDD/AGND. Due to space constraints, it is positioned between two much larger 0805 4.7 µF capacitors. We consider this a risk to the soldering process, due to the possibility of this capacitor being pulled against one of the larger capacitors by the surface tension of the solder. Additionally, its location makes spotting problems and performing rework harder.

  We do not believe removing this capacitor will have a significant overall effect on the decoupling of this power domain. The combined effect of the two 10 nF capacitors can be seen in
  Figure 12, as a ~30% reduction in power domain impedance between around 50 and 200 MHz. This is well above the operating frequency of the analogue parts of the SCA. To avoid the risks associated with redesigning this complex section of PCB, C36 will simply not be fitted to production boards, but the pads will remain.
- **TCM: Spacing of JTAG lines:** There is scope to space out the JTAG signal lines (TCK, TDI, TDO, TMS) where they run close together, without otherwise affecting the rest of the board. Whilst the current arrangement is not thought to cause any problems, this modification is easy to do and does not have any obvious drawbacks.
- **TCM: Addition of CONFIGSELECT signal to Megarray connector:** To enable the TCM's default configuration to be stored on the modules, the E-fuses on its GBTX must be burned. To do so, the correct configuration must be loaded and the E-fuses burned using I2C, since the optical link cannot be used until this is done. To avoid the user needing to manually change the state of DIP switches on the PCB, we will add a new track between R26 and the unused pin C3 on the Megarray connector to allow the configuration source to be selected electronically. This track can be routed such that no vias, re-routing of other tracks or splitting of power planes are required (see Figure 32), making this a low-risk modification.

Figure 32: Approximate location of new CONFIGSELECT track on bottom layer of TCM

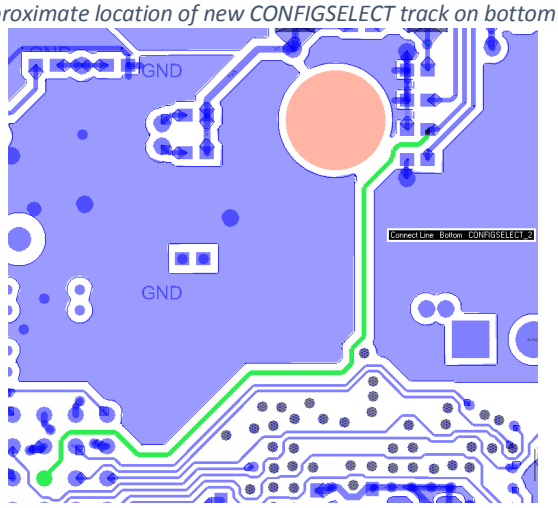

11. Risks and mitigations

The overall risk of the production of the TCMs and DTMs is thought to be reasonably low. Several batches of each module have already been made in small quantities and the design of the modules is mature, having been used in several test beams without serious problems. The main risks to the production of the modules are currently thought to be:

11.1 Errors on DTM GBTX 1 signals

As described earlier in this document, data on the optical link from GBTX 1 on the DTM has both a measurably worse eye diagram and also a higher rate of header errors than the similar link from GBTX 0. Based on the rate of frequency errors, we estimate that the overall error rate is likely to be of the order of 1 – 100 Hz, although this measurement is simplistic and may underestimate the rate.

The cause of these errors is currently unknown. The eye diagram for the 4.8 Gbps link looks acceptable, as does jitter on the reference clock. To date, the best theories we have for these difference are related to clock jitter, clock reflections and impedance mismatches, but we have found no clear evidence for any of these being the actual cause.

It is important to note that these measurements have been performed using the µDAQ system and it is possible error rates may be lower with the MiniDAQ 2. We also need to perform tests with much longer fibre-optic cables, to check that there is sufficient margin in the system that these do not cause problems.

At the moment, this effect is an annoyance, rather than a critical problem. We have successfully taken data with the current DTM boards, without obvious problems from these errors, so it is possible that the noise due to these errors is masked by other noise sources such as dark counts. It also seems likely that the higher-end hardware in the MiniDAQ 2 will, if anything, help to reduce these errors. However, since the problem is not fully understood, we have no way of fully quantifying its severity or overall effect on the data.

11.2 Halogen-free PCBs and manufacturing by new company

Our proposed approach to manufacturing means that we will simultaneously be switching to a new manufacturer for the PCBs as well as a new halogen-free substrate. The new substrate is likely to have a slightly higher relative permittivity, whilst the new manufacturer is likely to have different capabilities, equipment and may not be able to offer precisely the same stack-ups we have used to date.

The main risk from this is changes to the impedance of transmission lines causing signal integrity problems. We will therefore re-calculate and modify the track width and spacing of transmission lines to match the PCB material and stack-up used. There is also a small risk of unexpected manufacturing errors due to changes in the manufacturing process. However, we would hope that these sort of problems would come to light before production when the data is analysed by the manufacturer. Whilst the design of the PCB is complex, it is capable of being built by prototype manufacturers using standard technology, removing a lot of the risk associated with more exotic manufacturing.

We will further mitigate this risk by having the TCMs (500 boards) manufactured first, followed by the DTMs (1,400 boards), once the TCMs have been tested. If serious problems come to light, the TCM production could be re-run, at a cost that is not prohibitive (a few £k). This allows us to be sure of the production process before placing the much larger DTM order. We will use a similar approach for the assembly of the boards: an initial order of a few tens of boards will allow us to verify their performance, before an order is placed for the remaining boards in several, larger batches.