A BMT Layer-1 technology demonstrator card Hardware and Firmware

Authors : S. Mallios[1], K. Adamidis[1], G. Bestintzanos[1], C. Fountas[1], G. Karathanasis[2], P. Katsoulis[1], N. Manthos[1], I. Papadopoulos[1], S. Sotiropoulos[3], P. Sphicas[2] on behalf of
University of Ioannina[1], University of Athens[2] and Institute of Accelerating Systems and Applications (IASA), Athens[2]

Barrel Muon Trigger Layer 1 Demonstrator board

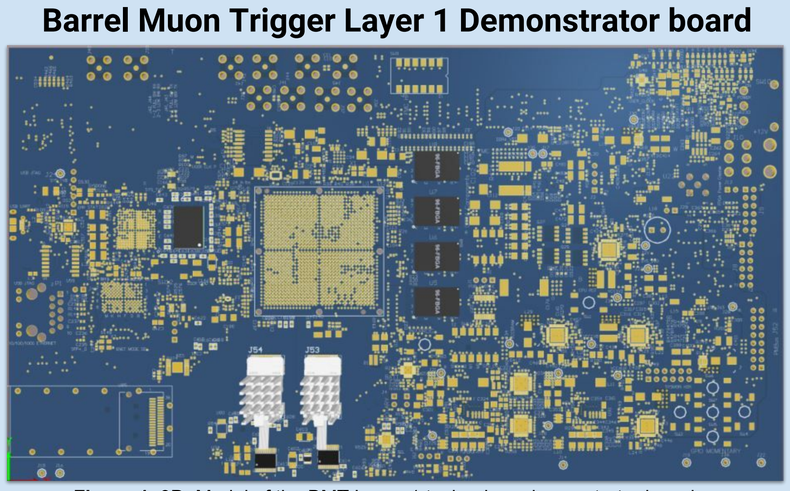

Figure 1. 3D Model of the BMT Layer-1 technology demonstrator board

The board is **powered by a Kintex UltraScale FPGA**, that provides the best price per performance per watt at 20 nm technology in a mid-range device. The Kintex UltraScale FPGA provides 20 next-generation GTH transceivers that reach speeds up to 16.3 Gbps. The board comes with state-of-the-art fibre optics technologies, from Samtec. The high performance interconnect system uses active optical engines, that provide 12 full-duplex channels, at data rates up to 16 Gbps. Furthermore, 4 FPGA transceivers are routed to a QSFP28 connector, allowing data rates of up to 28 Gbit/s per channel over 4 channels. In total the board's 16x16 Gbps links add up to a **total optical bandwidth** of approximately **256 Gbps in each direction**, making it a **high-performance all-optical data-stream processor**.
A Xilinx ZYNQ System-on-Chip (SoC) device will be used as the control interface for the Kintex UltraScale FPGA. The system controller sets up or queries on-board resources, such as the power controllers and programmable clocks.

		BMT-L1 Demo BOARD
FPGA	TYPE	Xilinx XCKU040
	Logic Cells (K)	530
	DSPs	1920
	BRAM (Mb)	21.1
	16.3G Transceivers	20
	I/O Pins	520
OPTICS	QSFPs	Finisar FTLC9551REPM
	Fireflies	Samtec ECUO-Y12-16
System Controller		Zynq XC7Z010
RAM	DDR4 (Mb)	2048
	DDR3 (Mb)	512
OPTICAL Links	Number	16
	Total Bandwidth (Gbps)	256
Other features	Ethernet PHY, JTAG, USB UART, SD Card, HDMI Multiple programmable clock generators	

Optics overview

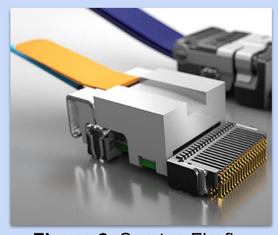

Figure 2. Samtec Firefly ECUO-Y12-16

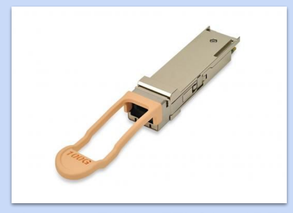

Figure 3. Finisar FTLC9551REPM QSFP28 Optical Transceiver

FireFly optical cable system
The FireFly optical flyover assembly is designed for flexibility and is interchangeable with the FireFly copper assembly using the same connector system. It is available with x12 simplex or duplex optical transceivers to achieve 16 Gb/s per channel.

Features
- Data "flies" over the PCB for easier routing
- Interchangeable with FireFly copper
- 16Gbps maximum data rate

QSFP28 transceiver module (limited to 16Gbps)
- 4x25Gb/s 850nm VCSEL-based transmitter (103.1Gb/s aggregate limited to 64Gb/s)
- Single 3.3V power supply and 3.5W maximum power dissipation
- Maximum link length of 100m on OM4 Multimode Fiber (MMF)
- Hot-pluggable QSFP28 form factor
- Digital diagnostics functions are available via an I2C interface

A GLANCE AT THE PHASE 2 TRIGGER RATES

At HL-LHC we will be dealing with high pileup beam conditions (200 PU).
The required trigger bandwidth will increase more than 7 times with respect to the current system.

- Level 1 trigger rate = 750 kHz
- HLT rate = 7.5 kHz
- Level 1 Trigger Latency = 12.5 μsec

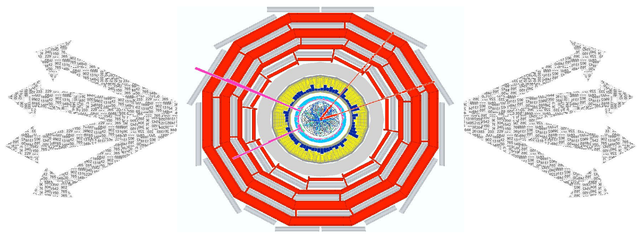

Clocking

The board includes 5 clock sources. The GTH transceivers connected to the high speed Firefly modules are clocked by a dedicated low jitter quad clock generator (Si5338). A low-jitter frequency generator (Si570) is connected to the QSFP transceivers and can be used as a secondary clock source to the Firefly transceivers. A jitter attenuator (Si5328B) is used to reduce the jitter of an RX recovered clock. A fixed frequency clock source can be used as a free running clock for reset and initialization FSMs. Finally an SMA external clock input is also included. All programmable clocks are accessed through a dedicated I2C bus.

The new CMS Drift Tubes Front End and Back End electronics

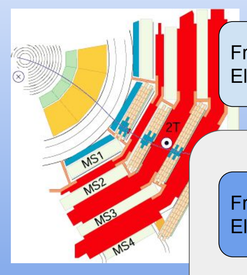

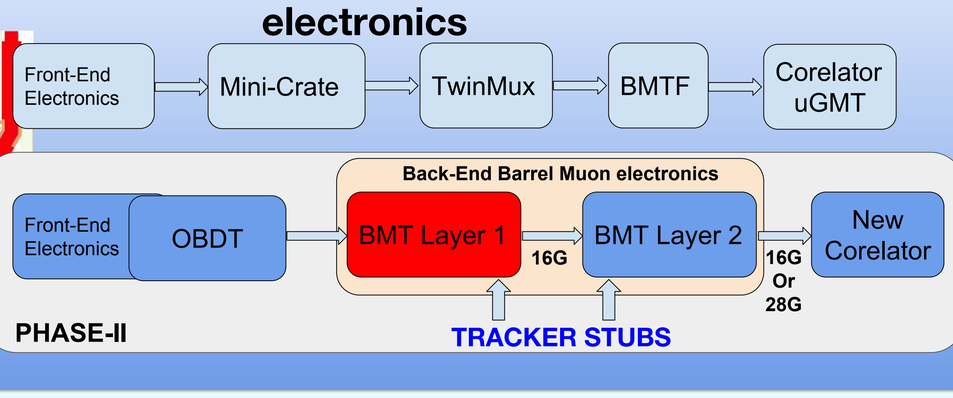

The new BMT Layer-1 processor

A two layer architecture for the entire BMT will provide standalone Barrel Muon Trigger Tracks, which are the result of Barrel Muon stubs and tracks with Track Trigger tracks The track trigger matching will increase the BMT efficiency for tracks which are currently lost because they have only one stub in the muon system.

At Layer-1 BMT Primitives are computed based on the TDC hit-data and using information from multiple DTs to achieve stubs of best resolution and BX identification. The data will be coming from the detector at 10 Gbps optical links. Preliminary tracking algorithms are executed and the results are transmitted to Layer-2 at 16 Gbps optical links. Optionally Track Trigger Stubs can be received from the Tracker two outer layers which gives another way to search for displaced muons other than the Kalman. The data will be further processed and combined with Track trigger tracks to provide BMT standalone tracks. After this a layer of 12 track finder cards would complete the track finding in the Barrel region. The new Barrel Muon Trigger includes a large part of the Mini-crate logic, the TwinMux and the BMTF logic.

What is a trigger?

A "trigger" is a filter which allows an event to pass only if a condition is satisfied.

Why we need trigger?

The detector data rate is enormous (10^9 events per second) x (typical event size of 1-2 Mbytes) ⇒ more than 1 Petabyte/s!
L1 trigger uses super-fast algorithms to select interesting events while suppressing less interesting events.

PCB design techniques

PCB Top Layer

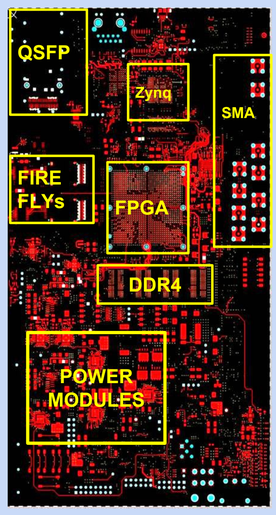

Figure 4. Top Layer

PCB Bottom Layer

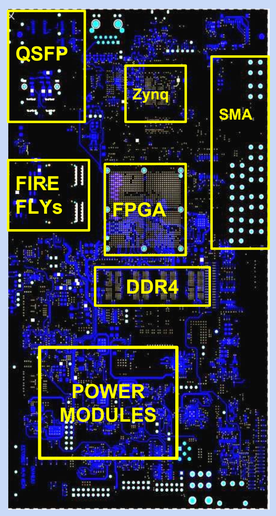

Figure 5. Bottom Layer

Backdrilling

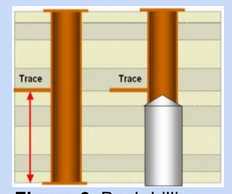

Figure 6. Backdrilling

Plated-through hole (PTH) via structures in high layer-count printed circuit boards (PCBs) can significantly distort high speed digital signals that pass through them. Often the distortion is severe enough that digital receivers can no longer ascertain whether a logical one or a logical zero was originally transmitted.
A dominant structure within a PTH via is the via stub. Since a via stub serves no useful function in the circuit, it can be removed using a technique known as backdrilling. Essentially, a drill bit slightly larger in diameter than the one used to create the original via hole is used to remove the undesired conductive plating in the via stub region. Decreasing via stub length by backdrilling significantly reduces a particularly problematic form of signal distortion called deterministic jitter.

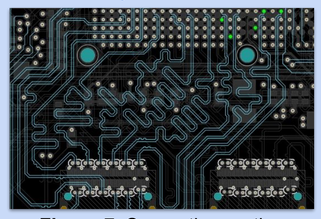

Figure 7. Serpentine routing

High-Speed Signal Trace Length Matching

Use of serpentine routing to match route lengths of high speed differential pairs route. We benefit from electromagnetic field cancellation only when both the signals in your differential pair have the same length and gap..

Firmware

Why 16Gbps links?

The benefits of migrating to 16Gbps links start with performance. The higher-speed interconnects create faster data transfers that will ensure that highly demanding, data-intensive trigger applications will have a simple migration path. The latest FPGA transceivers allow data rates up to 32 Gbps. However the cost of these high-end FPGAs is very high. The 16Gbps links, can be implemented using mid-range FPGAs at a much lower cost, while they are sufficient for most of the high speed interconnects of the upgraded Level-1 trigger system.

16 Gbps Ultrascale Asynchronous Links

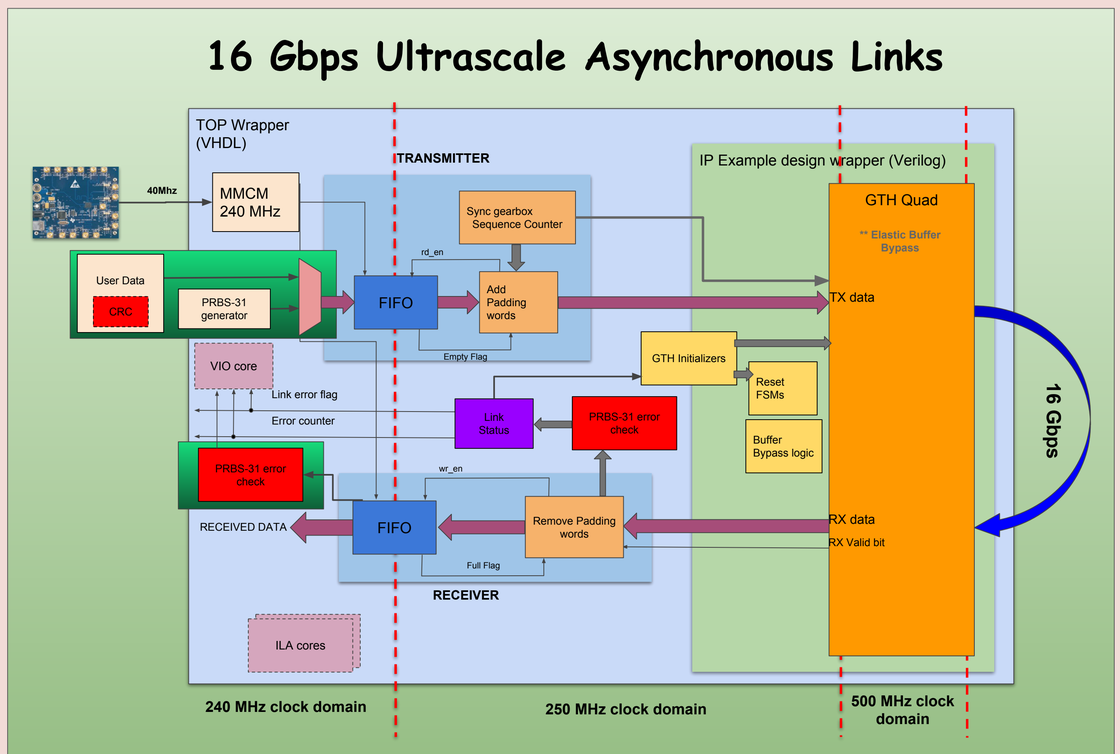

Figure 8. Asynchronous links Block Diagram

A simple protocol

The protocol is using the 64b/66b encoding. The links are Asynchronous meaning that the main algorithmic logic is clocked with a lower frequency than the link clock. It allows more flexibility when choosing the logic clock. This is achieved using asynchronous FIFOs in the receiving and the transmitting side. To compensate for the difference of the frequency a special word is being injected when the FIFO is empty. The alignment of the links are also based on the insertion and check of this padding word. For testing purposes the local clock is running at 240 MHz and the link clock at 250 MHz.

The 64b/66b encoding

Transforms 64-bit data to 66-bit line code to provide enough state changes to allow reasonable clock recovery and alignment of the data stream at the receiver. If the 2 header bits are :

- 01, the 64 payload bits are data.
- 10, the 64 payload bits hold control information
- 00 and 11 are not used and indicate an error

The overhead of 64b/66b encoding is 2 coding bits for every 64 payload bits or 3.125%.

Links Testing and debugging

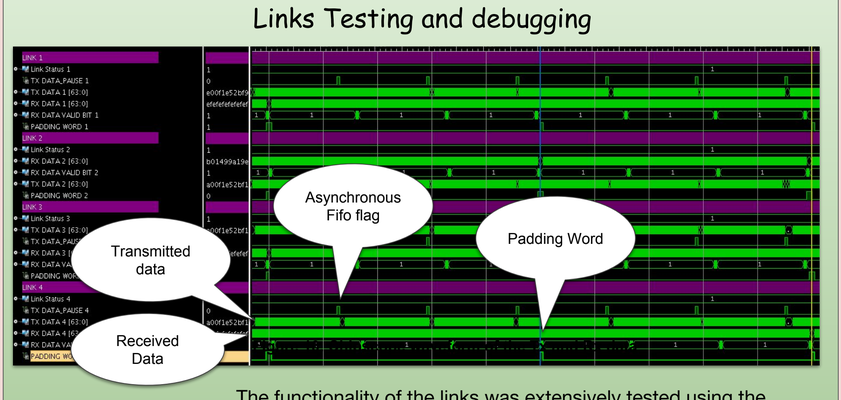

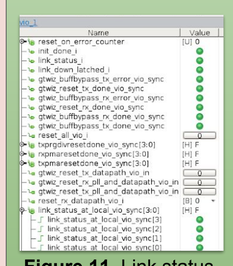

Figure 11. Link status indicator

The functionality of the links was extensively tested using the KCU-105 ultrascale developement board. For the tests an FMC loopback card was used to implement a copper loopback quad link.

Latency :
The latency of the GTH was measured at 9 CLKs adding the 2 FIFOs latency to cross between clock domains and the error checking code the link latency add up to a total latency of 23 CLKs.
Bit Error Rate Tests :
BER test were done by sending PRBS-31 data over an FMC copper loopback card. The links run for more than 72 hours without errors resulting in a **BER < 10^{-16}**

Bring-up and Error Check

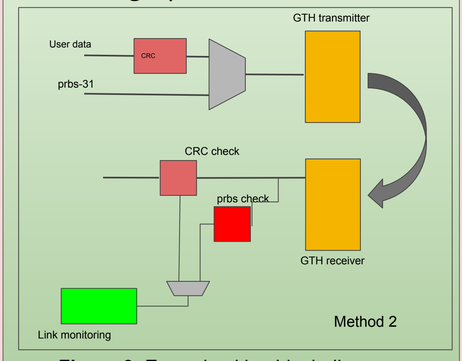

Figure 9. Error checking block diagram

The link bring-up and error detection is based on the generic 2-bit 64b/66b encoding header combined with the periodically sending of a padding word and CRC blocks.

- Single errors are considered **soft errors** and are monitored with a soft error counter.
- Continues errors are considered **hard errors** and result in auto reset and re-alignment of the links.

The overhead of the 64b/66b encoding is 3.125% and the CRC/padding blocks are injected every 100 blocks resulting at a **total overhead of 4.125%.**

The maximum time for the link (re)allignment is ~200us

References :

1. CMS Collaboration, The Phase-2 Upgrade of the CMS L1 Trigger, Interim Technical Design Report, CERN-LHCC-2017-013, CMS-TDR-017 , https://cds.cern.ch/record/2283192/files/CMS-TDR-017.pdf
2. CMS Collaboration, The Phase-2 Upgrade of the CMS Muon Detectors, Technical Design Report, CERN-LHCC-2017-012 CMS-TDR-016, https://cds.cern.ch/record/2283189/files/CMS-TDR-016.pdf

contact : stavros.mallios@cern.ch