

A full-function Global Common Module prototype for ATLAS Phase-II upgrade

W. Qian¹

*STFC Rutherford Appleton Laboratory,
Harwell Campus, Chilton, Didcot, Oxfordshire OX11 0QX, UK
E-mail: Weiming.Qian@stfc.ac.uk*

ABSTRACT: The High Luminosity Large Hadron Collider (HL-LHC[1]), an upgrade of the LHC, is set to become operational in 2029, aiming to achieve instantaneous luminosities 5-7.5 times larger than the nominal value of the LHC. However, unlocking the full physics potential at this much higher luminosity level necessitates a tenfold increase in the data bandwidth processed by ATLAS. This poses significant challenges to the design of the Trigger and Data Acquisition systems. To address these challenges, a baseline architecture has been chosen for the ATLAS Phase-II upgrade, relying on a single-level hardware trigger known as the Level-0 Trigger. This trigger has a maximum rate of 1 MHz and a latency of 10 μ s. Central to this upgrade is the inclusion of a new subsystem - the Global Trigger [2]. This component performs complex algorithms, akin to those currently used in Phase-I high-level trigger software (such as Topoclustering), on full-granularity calorimeter data. The Global Trigger is divided into three sublayers: the Multiplexer Processor (MUX) layer, the Global Event Processor (GEP) layer, and the Global to Central Trigger Processor [3] interface (gCTPi). A full-function Global Common Module (GCM) hardware prototype has been designed to fulfill the requirements of all three sublayers of the Global Trigger, featuring different firmware loads. This GCM prototype, based on the ATCA [4] front board form factor, incorporates two of the latest AMD (Xilinx) Versal Premium devices VP1802 [5]. These devices boast double the density of the Virtex UltraScale+ FPGA VU13P used in the previous design [6] and include an integrated SoC with a completely new architecture. To handle high-speed I/Os, this GCM prototype employs twenty 12-channel 25.7 Gb/s FireFly [7] optical engines. The estimated maximum power consumption of this GCM prototype is 400 W, which falls within the cooling capabilities of the ATLAS ATCA shelf. To ensure power integrity, signal integrity, and thermal performance, extensive PCB simulations and thermal simulations have been done to guide the layout design of the GCM prototype. This paper provides an in-depth overview of the design process for this full-function GCM prototype hardware, with a particular focus on technology choices and simulation results.

KEYWORDS: ATLAS; Trigger; Power Integrity; Signal Integrity.

¹On behalf of the ATLAS TDAQ Collaboration.

Copyright 2023 CERN for the benefit of the ATLAS Collaboration. CC-BY-4.0 license



Contents

1. Global Trigger Archietectre	1
2. Full-function GCM hardware prototype	2
2.1 GCM prototype hardware implementation	2
2.2 GCM prototype design methodology	2
3. GCM prototype PCB post-layout simulation	3
3.1 Power Integrity simulation	3
3.2 Signal Integrity simulation	3
3.2.1 VP1802 breakout optimization	4
3.2.2 Firefly optical engine breakout optimization	4
3.2.3 Typical 25 Gbps channel performance	4
3.2.4 Crosstalk optimization	5
3.2.5 Special test launch points	5
4. Summary	5

1. Global Trigger Archietectre

Figure 1 illustrates the architecture of the Global Trigger subsystem, specifically designed to provide Event Filter-like capabilities to the Level-0 trigger system for the HL-LHC. The Global Trigger architecture is structured into three layers: the Multiplexer Processor (MUX) layer, the Global Event Processor (GEP) layer, and the Global-to-Central Trigger Processor interface (gCTPi). The MUX layer consists of up to 56 nodes, managing a total throughput of around 60 Tb/s. These nodes gather data from detectors (calorimeter and muon) and legacy Feature Extractor modules through more than 2300 fibres. The MUX layer employs time-multiplexing on a bunch-crossing basis, sending complete events to 49 nodes in the GEP layer in a round-robin manner. Within the GEP layer, each GEP node processes the entire event data related to a specific bunch-crossing. These nodes execute complex algorithms and transmit the outcomes to the gCTPi. The gCTPi, comprising a single node, selects and resynchronizes results from all GEP nodes before forwarding them to the Central Trigger Processor (CTP). Notably, this time-multiplexed architecture of Global Trigger is highly scalable and enables the implementation of asynchronous and iterative high-level algorithms.

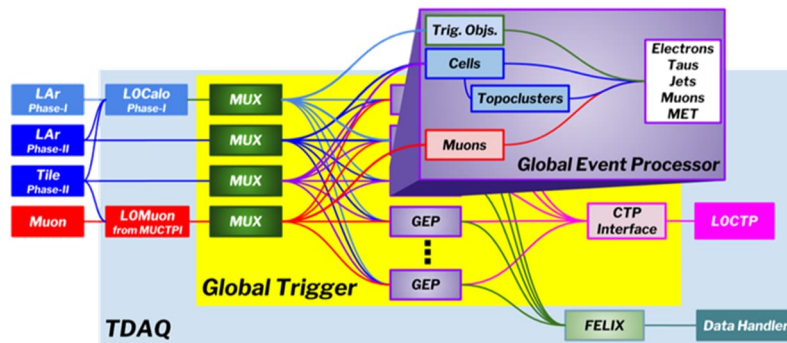


Figure 1. Global Trigger Architecture for ATLAS Phase-II TDAQ Upgrade

2. Full-function GCM hardware prototype

2.1 GCM prototype hardware implementation

Having proven successful in the ATLAS Phase-I system, the ATCA platform is chosen for the implementation of this GCM prototype. Figure 2 depicts the block diagram of its design. The top right Versal Premium VP1802 is designated for a MUX node, while the bottom left VP1802 is allocated for a GEP or gCTPi node. Surrounding both VP1802 units are twenty Firefly 25Gbps parallel optical engines, resulting in a total of 240 high-speed links.

The maximum estimated power consumption for each VP1802 device, using AMD’s Power Design Manager software, is 130W with 50% resource utilization running at 320MHz. To provide an adequate margin, the hardware power design on the board is capable of supplying 165W to each VP1802 device, equivalent to 70% resource utilization running at 320MHz, with the current on the VP1802 core voltage rail VCCINT reaching 170A.

Considering the vertical airflow configuration of the ATLAS standard ATCA shelf, the placement of the two VP1802 devices on the GCM is staggered vertically to effectively manage the cooling design challenge of the board. It is anticipated that the modules will operate for more than twenty years so it is crucial to maintain critical devices at temperatures significantly below their maximum ratings. The design and simulation of heatsinks for both the VP1802 and Firefly optical engines have been outsourced to specialized companies. Simulations suggest achievable temperatures of 70 °C for the VP1802 and 50 °C for the Firefly optical engine, following the board layout as depicted in Figure 2.

With different firmware loads, this GCM board fulfils the requirements of all three sublayers of the Global Trigger. This common hardware approach significantly simplifies system design and long-term maintenance, minimizing the complexity of firmware and software development by leveraging a shared infrastructure.

2.2 GCM prototype design methodology

This GCM prototype stands out as a high-speed, high-density, and high-power ATCA front board. To ensure the success of this project, a systematic PCB design methodology is adopted, as illustrated in Figure 3.

PCB simulation is integral to such a complex board design and is seamlessly integrated into our design flow. Pre-layout simulation has played a crucial role in determining various aspects of PCB technology, including laminate material, layer count, via technology, copper thickness, BGA breakout pattern, and more. Consequently, this GCM prototype utilizes a 26-layer PCB with via-in-pad and backdrill technology. Via-in-pad technology is particularly vital for high-speed large BGA breakout, as it reduces 3D impedance discontinuities and improves routability simultaneously. Backdrill technology is essential

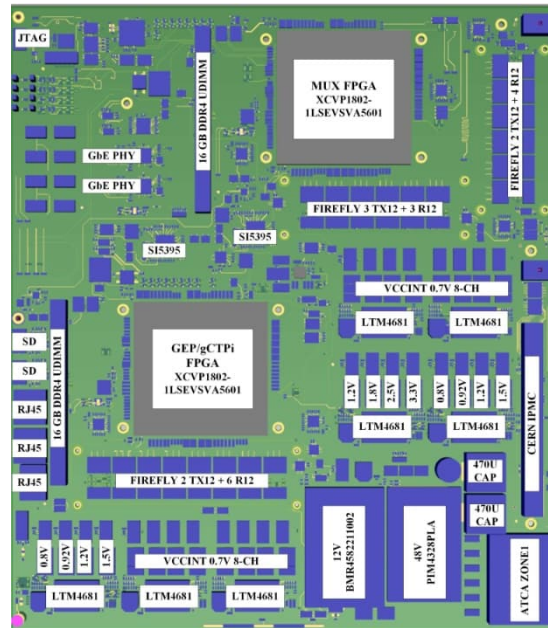


Figure 2. Full-function GCM prototype block diagram

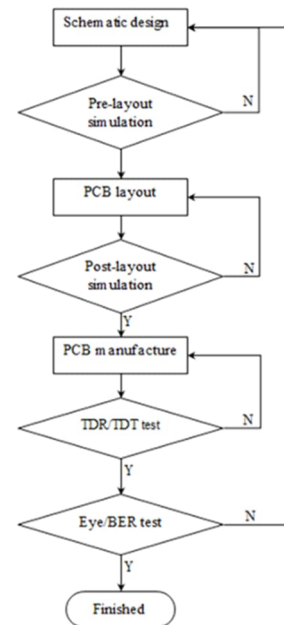


Figure 3. GCM PCB design methodology

for the performance of high-speed vias. The total copper weight in the GCM PCB stackup is 6 oz for four power layers and 7 oz for twelve ground layers. Only ground layers are used as the reference planes for high-speed signal layers. This PCB configuration represents the state of the art in the PCB industry, considering the ATCA board thickness constraint. Post-layout simulation has been instrumental in optimizing and verifying the PCB performance, with detailed results presented in the next section.

3. GCM prototype PCB post-layout simulation

3.1 Power Integrity simulation

With the AC performance (e.g., ripple noise) of DC-DC components (e.g., LTM4681) already tested and verified on evaluation boards beforehand, the primary challenge in the GCM PCB power distribution lies in addressing the issue of on-board DC drop on large current power rails.

Figure 4 displays the simulation results of the DC drop before optimization on the VP1802 core voltage power rail VCCINT at a maximum current of 170 A. Two issues are identified. Firstly, the DC drop on VCCINT is excessively high, resulting in a total of 14 W dissipated in the copper. This impacts the GCM power budget and leads to cooling problems of the PCB itself. Secondly, numerous power vias carry over 2 A current, posing a concern for the board's longevity.

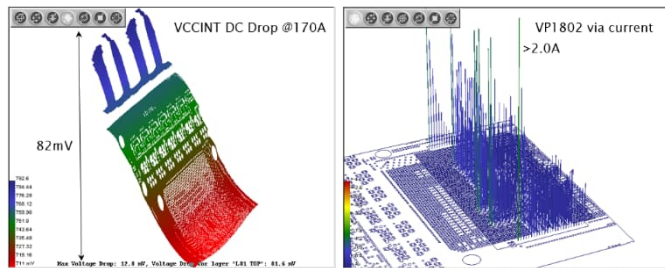


Figure 4. VP1802 core voltage VCCINT simulation before optimization

To address these problems, the VCCINT power distribution is refined through a meticulous examination of nearby signal layers. Some signals are strategically rerouted, facilitating the addition of targeted copper fills in signal layers for improved power distribution. Figure 5 demonstrates the results of this optimization after multiple iterations, highlighting a reduction of more than half in the DC drop and the effective suppression of via current spikes.

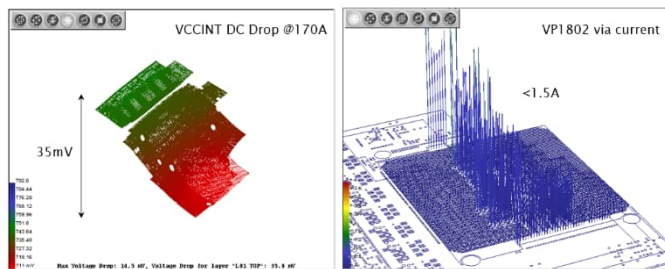


Figure 5. VP1802 core voltage VCCINT simulation after optimization

Another issue identified during the power DC simulation is the occurrence of a hotspot on a relatively low-current power distribution, as depicted in Figure 6 (left). Despite the current on this power rail being only 15 A, a hotspot develops due to the heavily perforated power plane. Once spotted, addressing such an issue is relatively straightforward, as demonstrated in Figure 6 (right).

Another issue identified during the power DC simulation is the occurrence of a hotspot on a relatively low-current power distribution, as depicted in Figure 6 (left). Despite the current on this power rail being only 15 A, a hotspot develops due to the heavily perforated power plane. Once spotted, addressing such an issue is relatively straightforward, as demonstrated in Figure 6 (right).




Figure 6. Left: Hotspot on lower current power rails on GCM Right: hotspot fixed

3.2 Signal Integrity simulation

The majority of the 240 on-board high-speed links operate at 25 Gbps. The key challenges here involve optimizing the 3D breakout area at both ends of the links and minimizing the crosstalk between the links.

3.2.1 VP1802 breakout optimization

Figure 7 (left) illustrates the use of via-in-pad technology for the VP1802 breakout area. In this approach, the vias are drilled directly into the VP1802 footprint pads, effectively merging the two 3D impedance discontinuities of the pad and via into one. Additionally, back drilling is employed to remove the via stubs. To align the impedance




Figure 7. VP1802 breakout optimization. Left: BGA routing and backdrill. Middle: Simulated differential TDR response ($T_r = 20\text{ps}$). Right: insertion loss (green) and return loss (red)

of the differential vias with the 93-ohm target of VP1802, a dog-bone-shaped anti-pad is incorporated around the differential vias. The simulation result of the final optimization is depicted in Figure 7 middle and right, where the VP1802 breakout pattern's impedance is brought within 10% of the target impedance, and the performance in frequency domain is excellent beyond the signal spectrum of 25Gbps.

3.2.2 Firefly optical engine breakout optimization

The Firefly optical engine employs a fine-pitch surface mount connector, as shown in Figure 8 (left). To align the impedance closer to the target, the ground plane immediately under the differential pads in the connector is cut out. Differential vias with four ground vias are utilized to connect the high-speed links to inner layers. A sweeping process in the simulation is employed to determine the optimal size of the anti-pads for differential vias. The simulation result of the final optimization is depicted in Figure 8 middle and right, where the Firefly optical engine breakout pattern's impedance is brought within 10% of the target impedance. The performance in the frequency domain is very good beyond the signal spectrum of 25Gbps.

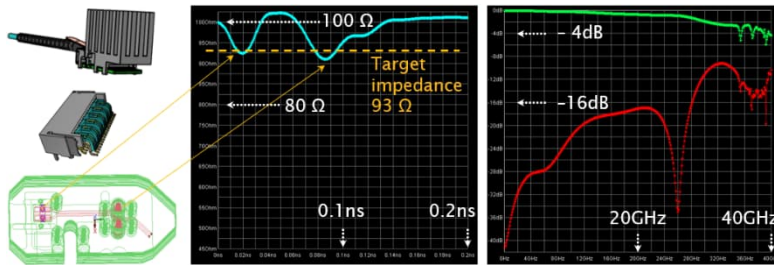


Figure 8. Firefly optical engine breakout optimization. Left: firefly component, connector, and breakout routing. Middle: Simulated differential TDR response ($T_r = 20\text{ps}$). Right: insertion loss (green) and return loss (red)

Figure 8 middle and right, where the Firefly optical engine breakout pattern's impedance is brought within 10% of the target impedance. The performance in the frequency domain is very good beyond the signal spectrum of 25Gbps.

3.2.3 Typical 25 Gbps channel performance

25Gbps links (3 to 4 inches long) on the GCM prototype are simulated with optimization applied to both ends. The impedance of the entire channel is well controlled, as depicted in Figure 9 (left). The insertion loss of this channel smoothly rolls off with frequency, providing a significant margin against the limit set in the industry standard OIF-CEI-04.0 [8].

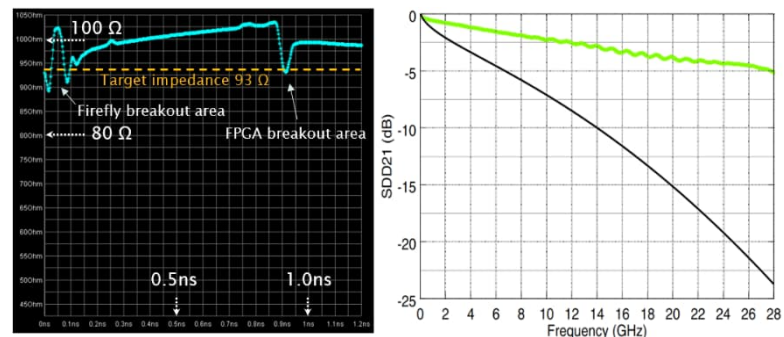


Figure 9. Performance of a typical 25 Gbps link on the GCM prototype. Left: Simulated TDR response ($T_r=20\text{ps}$) for the whole channel. Right: Insertion loss of a whole channel on GCM (green) and minimum SDD21 recommended by industry standard OIF-CEI-04.0 for VSR channel (black).

3.2.4 Crosstalk optimization

Crosstalk control on a high-density, high-speed board is of utmost importance. AMD has specified stringent crosstalk requirements between their multi-gigabit transceivers (MGT) on the VP1802 in three cases. For the Tx-Tx very short-range (VSR) case, the crosstalk limit is -35dB at signal Nyquist frequency. For the Tx-Rx VSR, the limit is -45dB. For the Rx-Rx VSR, the limit is -40dB. This GCM design opts for the larger package VSVA5601 among the two VP1802 package variants, as it offers better crosstalk performance due to the increased number of ground pins placed between MGTs. To further reduce crosstalk, the Tx and Rx links are allocated into separate MGT QUADs. Subsequently, MGT crosstalk simulations are conducted on the GCM prototype PCB layout. While Tx-Tx and Tx-Rx crosstalk meet the specifications, Rx-Rx crosstalk failed in some specific situation, as depicted in Figure 10 (left). This issue has been traced down to differential tracks passing closely by differential vias of another link. By swapping signal routing layers and implementing back drill, this condition can be avoided. Figure 10 (right) shows the Rx-Rx crosstalk meeting the requirements after optimization.

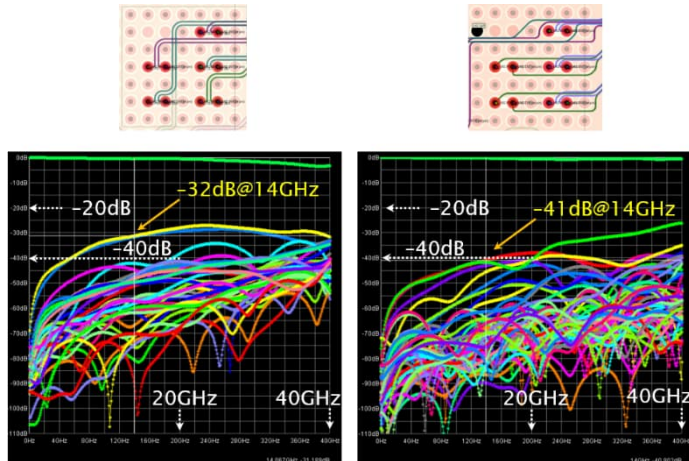


Figure 10 GCM MGT Rx-Rx crosstalk optimization. Left: before optimization. Right: after optimization

To further reduce crosstalk, the Tx and Rx links are allocated into separate MGT QUADs. Subsequently, MGT crosstalk simulations are conducted on the GCM prototype PCB layout. While Tx-Tx and Tx-Rx crosstalk meet the specifications, Rx-Rx crosstalk failed in some specific situation, as depicted in Figure 10 (left). This issue has been traced down to differential tracks passing closely by differential vias of another link. By swapping signal routing layers and implementing back drill, this condition can be avoided. Figure 10 (right) shows the Rx-Rx crosstalk meeting the requirements after optimization.

3.2.5 Special test launch points

The validation of PCB simulation results is equally important in the domain of high-speed electronics. In this context, precise and accurate testing plays a crucial role in guaranteeing the performance and reliability of intricate circuit designs. Employing specialized test launch designs becomes instrumental in achieving controlled signal transitions, thereby minimizing the impact of launch-induced distortions, and facilitating accurate measurement of signal integrity parameters. The utilization of 2.4 mm precision connectors on the GCM prototype ensures high-quality connections to high-end oscilloscopes. The breakout pattern of this connector on the GCM is also subjected to simulation and optimization using a 3D solver, as illustrated in Figure 11. This approach ensures that the test setup itself is optimized for accurate and meaningful high-speed signal measurements.

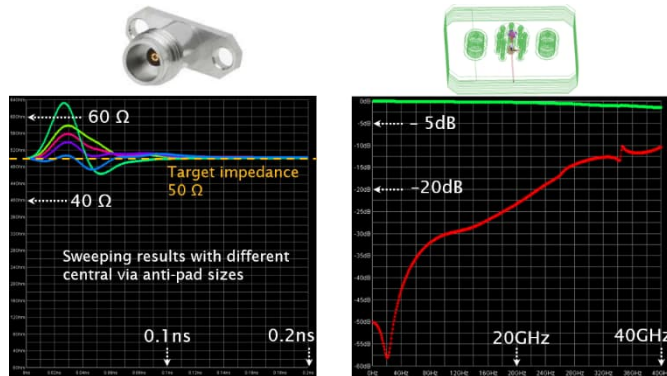


Figure 11. GCM special test launch design. Left: sweeping simulation to choose the best impedance match. Right: Insertion loss (green) and return loss (red) of the best launch design.

4. Summary

A new full-function GCM prototype has been meticulously designed and implemented for the ATLAS Phase-II Upgrade's new Global Trigger. This GCM prototype is a high-speed, high-power, and high-density ATCA front board. Throughout the design process, a systematic methodology has been employed that concurrently addresses signal integrity, power integrity, and thermal integrity. This project passed the ATLAS Preliminary Design Review in October 2023, and the first prototype board is being manufactured.

References

- [1] G. Apollinari et al., “High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report V. 0.1. CERN Yellow Reports: Monographs,” CERN, Geneva, 2017. <http://cds.cern.ch/record/2284929>
- [2] The ATLAS TDAQ Collaboration, “Technical Design Report for the Phase-II Upgrade of the ATLAS Trigger and Data Acquisition System,” CERN-LHCC-2017-020. ATLAS-TDR-029, CERN, Geneva, Jun, 2018. <https://cds.cern.ch/record/2285584/files/ATLAS-TDR-029.pdf>
- [3] G Anders et al., “The upgrade of the ATLAS Level-1 Central Trigger Processor,” Journal of Instrumentation, Volume 8, January 2013
- [4] ATCA specification. <https://www.picmg.org/openstandards/advancedtca/>
- [5] AMD Xilinx, Versal Architecture and Product Data Sheet: Overview. <https://docs.xilinx.com/v/u/en-US/ds950-versal-overview>
- [6] S. Tang et al., "Prototype Design of Global Common Module for ATLAS Experiment's Phase-II Upgrade," in IEEE Transactions on Nuclear Science, vol. 70, no. 9, pp. 2248-2255, Sept. 2023, doi: 10.1109/TNS.2023.3302158.
- [7] Samtec, “FireFly™ Active Optical Micro Flyover System™ Cable Assembly,” <https://www.samtec.com/products/ecuo>
- [8] Optical Internetworking Forum IA Title: Common Electrical I/O (CEI) - Electrical and Jitter Interoperability agreements for 6G+ bps, 11G+ bps, 25G+ bps I/O and 56G+ bps. <https://www.oiforum.com/wp-content/uploads/2019/01/OIF-CEI-04.0.pdf>