

The Development of a General Purpose Processing Unit for the Upgraded Electronics of the ATLAS Tile Calorimeter



Abstract. The Large Hadron Collider at CERN generates enormous amounts of raw data which presents a serious computing challenge. After planned upgrades in 2022, the data output from the ATLAS Tile Calorimeter will increase by 200 times to over 40 Tb/s! System on Chips (SoC) such as ARM and Intel Atom are common in mobile devices due to their low cost, low energy consumption and high performance. It is proposed that a cost-effective, high data throughput Processing Unit (PU) can be developed by using several consumer SoCs in a cluster configuration to allow aggregated processing performance and data throughput while maintaining minimal software design difficulty for the end-user. This PU could be used for a variety of high-level functions on the high-throughput raw data such as spectral analysis and histograms to detect possible issues in the detector at a low level. High-throughput I/O interfaces are not typical in consumer SoCs but high data throughput capabilities of greater than 20 Gb/s per PU is feasible via the novel use of PCI-Express as the I/O interface to the SoCs. An overview of the PU is given and the results of throughput testing of Freescale Semiconductor i.MX6 quad-core ARM Cortex-A9 processors are presented.

1. Introduction

Projects such as the Large Hadron Collider (LHC) generate enormous amounts of raw data which presents a serious computing challenge. After planned upgrades in 2022, the data output from the ATLAS Tile Calorimeter (TileCal) will increase by 200 times to over 40 Tb/s (Terabits/s) [1]. It is infeasible to store this data for offline computation.

A paradigm shift is necessary to deal with these future workloads and the cost, energy efficiency, processing performance and I/O throughput of the computing system to achieve this task are vitally important to the success of future big science projects. Current x86-based microprocessors such as those commonly found in personal computers and servers are biased towards processing performance and not I/O throughput and are therefore less-suitable for high data throughput applications otherwise known as Data Stream Computing [2].

ARM System on Chips (SoCs) are found in almost all mobile devices due to their low energy consumption, high performance and low cost [3]. One of the first steps to a true Data Stream Computing system is a high data throughput Processing Unit (PU). The author is developing an ARM-based PU for use by ATLAS TileCal as a high throughput, general purpose co-processor to the read-out system Super Read Out Driver (sROD) which will be used to combat the issue of

pile-up. A general purpose co-processor is able to run more sophisticated and memory intensive algorithms than the FPGA-based sROD, although the latency is inferior which is why the sROD is used in the data path.

A brief discussion of the ATLAS TileCal read out architecture and the sROD is given in Section 2. An overview of the energy reconstruction methodology for Photo Multiplier Tubes (PMTs) is provided in Section 3. Possible issues with the current methodology is also given. The Processing Unit (PU) is described in Section 4 along with how it could be used to help solve the issues with the PMT energy reconstruction algorithms. Section 5 concludes with a brief discussion of future work.

2. TileCal Read Out Architecture

The TileCal read out architecture is required to digitize the analog signals produced by the PMTs located on the Tile Calorimeter. In the existing hardware, shown in Fig. 1, analog circuitry is used extensively. At the time of design and construction, digital electronics were not fast enough to satisfy the requirements. In the upgraded system, shown in Fig. 2, the digital electronics have been superseded by the sROD which is based on high-end Field Programmable Gate Arrays (FPGAs). There is also an upgraded front-end which also contains FGPA for digitising the analog signals.

The sROD is located in the back-end, off the detector to avoid the requirement for expensive radiation-hard electronics. A photo of the prototype sROD circuit board before assembly is shown in Fig. 3. The sROD will be located in an industry standard AdvancedTCA (ATCA) chassis (seen in Fig. 3) which enables comprehensive redundancy and monitoring to ensure maximum uptime.

In both the existing and the upgraded systems, a pipeline is used to store events until the level one trigger provides an accept signal. This short delay is required while the level one trigger performs computations. In the upgraded system the sROD is able to perform some calculations before sending data to the rest of the triggering and data acquisition system. Some of these sROD calculations will be explained in Section 3. A general purpose Processing Unit can be used to enhance this functionality.

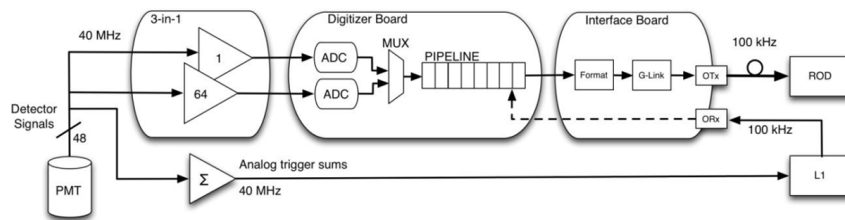


Figure 1: ATLAS TileCal current read out architecture [1].

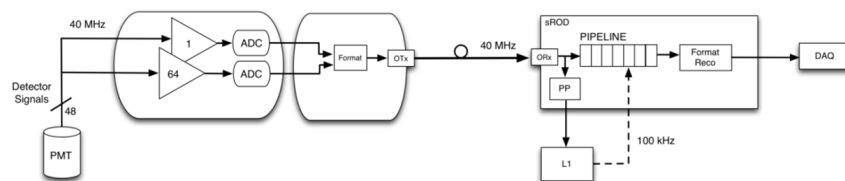


Figure 2: ATLAS TileCal upgraded read out architecture showing the sROD [1].




Figure 3: Photo of the sROD (left) and AdvancedTCA chassis (right).

3. TileCal Energy Reconstruction

Bunches Crossings (BC) in the ATLAS detector happen at a rate of 40 MHz. In reality there can be over 20 separate collisions in a single bunch crossing which leads to pile-up or Out Of Time (OOT) signals in the detector. The Tile Calorimeter is made up by layers of scintillator and steel. When a particle interacts with the scintillator, a pulse of light is produced which is converted to an electrical signal by PMTs.

This electrical signal is conditioned and spread into a pulse with a length of 150 ns and full width at half maximum of 50 ns. An Analog to Digital Converter (ADC) samples this pulse every 25 ns resulting in seven samples per pulse. A reference pulse showing example sampling points is visible in Fig. 4 with three main parameters of the pulse also illustrated.

Optimal Filtering (OF) or a Matched Filter (MF) are two methods by which the amplitude, A , phase, τ , and base-line pedestal, p , parameters can be calculated from the seven ADC samples of a pulse [4]. The pulse shape and therefore the energy can be reconstructed, when required, from these three parameters [5].

For both the OF and MF algorithms, each parameter (A , τ or p) can be found by multiplying the ADC samples by a specific set of weights which are calculated ahead of time. Both algorithms work well in low luminosity operation where the background noise of the PMT signals is gaussian and uncorrelated. This assumption fails for high luminosity operation (above about $\sqrt{s} = 8$ TeV) where the background noise is no longer uncorrelated due to pile-up [4]. This effect is visible in Fig. 4 where the algorithm energy calculations are not correlated with multiple OOT signals as well as the negative energies produced by the OF algorithm.

3.1. Energy Reconstruction Supervision

As discussed in Section 3, the energy reconstruction algorithms fail to work as expected under high pile-up conditions caused by higher luminosity operation of the LHC. A general purpose PU, as described in Section 4, can be used to alleviate this issue through the use of higher level programming of more sophisticated algorithms and the availability of more memory.

The PU is unlikely to operate within the latency constraints of the level one trigger and sROD and therefore cannot be in the critical data path. It can, however, be used to plot a histogram of the raw data and run its own analysis to determine if the algorithms on the sROD are functioning correctly. If they are not, the operators or engineers will be notified. This analysis is very difficult without the use of a PU. It is feasible that the PU can also be used to recompute the weights used by the OF or MF algorithms when an issue is detected which would greatly enhance the precision and reliability of the TileCal.

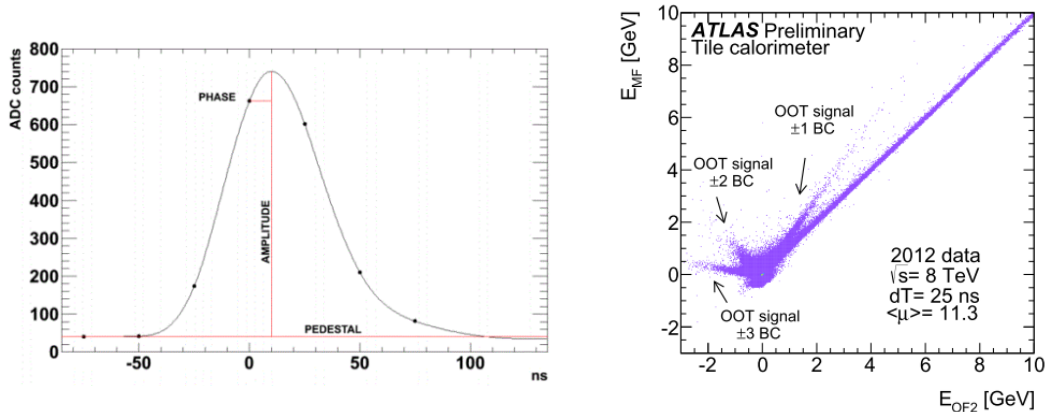


Figure 4: TileCal reference pulse from the PMT showing 7 sample points (left) and the energy reconstruction mismatch between the OF and MF algorithms under higher luminosity (right) [4].

4. Processing Unit

ARM System on Chips (SoCs) are low cost, energy efficient and high performance which has led to their extensive use in mobile devices. ARM performance and energy efficiency results have been published previously [2]. An ARM-based Processing Unit (PU) is under development to complement the sROD with higher-level computational tasks such as those discussed in Section 3.1.

The completed PU will be located in the ATCA chassis on an Advanced Mezzanine Card next to the sROD or as a separate board connected to the back-plane. The PU will be able to process up to 40 Gb/s raw data, fed via the ATCA from the sROD. A PCI-Express I/O interface will be used to link the FPGA on the sROD to a cluster of ARM SoCs on the PU.

PCI-Express throughput tests have been performed on a pair of Freescale i.MX6 quad-core ARM Cortex-A9 SoCs clocked at 1 GHz, located on Wandboard development boards [6]. The results are presented in Tab. 1 and a photo of the custom test setup designed by the author is in Fig. 5. Three tests were run to ascertain the maximum data throughput that can be obtained from the i.MX6 SoC: a simple CPU based memcopy command and two Direct Memory Access (DMA) transfers, initiated by the Endpoint (EP) or slave and the Root Complex (RC) which is the host.

Table 1: PCI-Express throughput results of a i.MX6 pair.

	CPU memcpy	DMA (EP)	DMA (RC)
Read (MB/s)	$94.8 \pm 1.1\%$	$174.1 \pm 0.3\%$	$236.4 \pm 0.2\%$
Write (MB/s)	$283.3 \pm 0.3\%$	$352.2 \pm 0.3\%$	$357.9 \pm 0.4\%$

The theoretical maximum throughput for the PCI-Express Gen 2 x1 link that was used is 500 MB/s. The best result is using DMA initiated by the RC but it is only 72% of the theoretical maximum. The RC-mode drivers are more optimized than the EP-mode drivers due to limited manufacturer support for EP-mode. The read results are lower than write because of overheads to initiate the read. The PU architecture will take these differences into account and use a data push rather than a pull based approach.




Figure 5: PCI-Express test setup for a pair of i.MX6 SoCs.

5. Discussion, Conclusions and Future Work

High data throughput computing, or Data Stream Computing, is required for projects such as the LHC which produce enormous amounts of raw data. A general purpose ARM System on Chip based processing unit is being developed which will be used as a co-processor to the sROD to help mitigate the energy reconstruction issues caused by pile-up under higher luminosity (over $\sqrt{s} = 8$ TeV) operation of the LHC.

A PCI-Express interface will be used for the raw data transfer between the sROD and the PU. Initial throughput measurements presented for a pair of Freescale i.MX6 quad-core Cortex-A9 SoCs are 72% of the theoretical maximum 500 MB/s for the available x1 link. Twelve of these SoCs will therefore be connected in parallel to provide the required 40 Gb/s throughput.

The next stage of research by the author is to measure the data throughput through a PCI-Express switch with up to eight i.MX6 SoCs connected. This will enable a close prototype of the final PU design to be tested.

Acknowledgements

The financial assistance of the National Research Foundation (NRF) towards this research is hereby acknowledged. We would also like to acknowledge the School of Physics, the Faculty of Science and the Research Office at the University of the Witwatersrand, Johannesburg.

References

- [1] F Carrió *et al.*, *The sROD Module for the ATLAS Tile Calorimeter Phase-II Upgrade Demonstrator*, Meyrin, Switzerland, 2013. [Online]. Available: <http://cds.cern.ch/record/1628753?ln=en>.
- [2] M. A. Cox, R. Reed, T. Wrigley, G. Harmsen, and B. Mellado, "Performance Characterisation of ARM Cortex-A7, A9 and A15 System on Chips for Data Stream Computing," *In Review*, 2014.
- [3] T. Krazit, *ARMed for the living room*, 2006. [Online]. Available: http://news.cnet.com/ARMed-for-the-living-room/2100-1006_3-6056729.html.
- [4] B. S. Peralva, "The Tilecal Energy Reconstruction for Collision Data Using the Matched Filter," ATLAS TileCal, 2013. [Online]. Available: <https://cds.cern.ch/record/1629575/files/ATL-TILECAL-PROC-2013-023.pdf>.
- [5] E. Fullana *et al.*, "Digital Signal Reconstruction in the ATLAS Hadronic Tile Calorimeter," *IEEE Transactions on Nuclear Science*, vol. 53, no. 4, pp. 2139–2143, Aug. 2006, ISSN: 0018-9499. DOI: 10.1109/TNS.2006.877267. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=1684077>.
- [6] Wandboard.org, *Wandboard - Freescale i.MX6 ARM Cortex-A9 Opensource Community Development Board*, 2012. [Online]. Available: <http://www.wandboard.org/>.