Development of Readout Electronics for a Digital Tracking Calorimeter

Ola Grøttvik[*,a], **Johan Alme**[a], **Rene Barthel**[b], **Tea Bodova**[a], **Viatcheslav Borshchov**[c],
Anthony van den Brink[b], **Viljar Eikeland**[a], **Alf Herland**[a], **Naomi van der Kolk**[b],
Simon Voigt Nesbø[a,d], **Thomas Peitzmann**[b], **Dieter Röhrich**[a], **Ganesh Tambave**[a],
Ihor Tymchuk[c], **Kjetil Ullaland**[a], **Shiming Yang**[a]

[a] *Department of Physics and Technology, University of Bergen, Bergen, Norway*

[b] *Institute for Subatomic Physics, Utrecht University/Nikhef, Utrecht, Netherlands*

[c] *LTU, Kharkiv, Ukraine*

[d] *Western Norway University of Applied Science, Bergen, Norway*
E-mail: `Ola.Grottvik@uib.no`

On behalf of the Bergen pCT Collaboration

Highly segmented digital tracking calorimeters consist of multiple layers of high-granularity pixel detector CMOS sensors and absorption/conversion layers. Two separate prototypes are being developed: (1) an electromagnetic calorimeter for a proposed ALICE upgrade (during Long Shutdown 3) and (2) a hadronic calorimeter for medical proton computed tomography imaging. These prototypes employ the ALPIDE detector chip developed for the ALICE Inner Tracking System. The ALPIDE chips are mounted on intermediate aluminum/polyimide flexible circuits with ultrasonic welding. This contribution presents findings and solutions to the challenging design of high-speed readout electronics with efficient use of FPGA resources for these prototypes.

---

[*]Speaker.

1. Introduction

A Digital Tracking Calorimeter (DTC) aims to simultaneously track and measure the range, and inherently the energy, of individual charged particles. A particle moving through the layers of a DTC will produce digital hit map data in each layer by ionization. By applying a path-finding algorithm to these data, one can find the length of the particle's track through the detector. With this length, together with the cluster-size information from each layer, one can calculate the residual energy of the incoming particle.

The University of Bergen (UiB) is involved in the development of two DTC prototypes, where one is a proposed upgrade for the ALICE Long Shutdown 3, the Forward Calorimeter (FoCal) [1], and the other is for medical imaging, the proton CT (pCT) scanner [2]. The pCT is designed as a hadronic calorimeter and aims to increase the accuracy of dose planning for ion-therapy treatment. The pCT will directly measure the relative stopping power map of tissue and bone, and thus will avoid the conversion from conventional CT-units that is associated with uncertainties up to several percents [2].

This contribution covers solutions to three main challenges to the readout electronics design: (1) keeping the material budget low and maintaining a near-homogenous structure, (2) dealing with a high number of high-speed data links while restricting the number of FPGAs, and (3) avoiding back-pressure in spite of limited data buffers on the FPGAs.

2. pCT System Overview

The pCT prototype consists of 41 layers of multiple high-granularity pixel-sensors. The first two layers are intended to capture the incoming position and angle of particles, and thus have no absorber layer in between them to avoid multiple scattering whereas the following 39 layers have a 3.5 mm aluminum absorber layer in-between each layer. The absorber layer causes the incoming particles to lose energy and stop in the detector and ensures that the particle's range can be measured. Pettersen et al. [3] describes the optimization of the detector.

A single pCT layer has a total of $\sim$ 56 Mega-pixels and is composed of 9×12 ALPIDE chips (developed for ALICE Inner Tracking System) and covers an area of $\sim 27 \times 18$ cm. Each ALPIDE chip is set in high-speed data transfer mode to avoid any data loss caused by temporarily high occupancy in a small area. Monte Carlo simulations of a low intensity, fast scanning beam, along with a SystemC-model of the sensor, show that there is no pile-up of data in the sensors. As shown in Figure 1a, with a trigger rate of 5 μs, a single sensor will produce an actual data rate of up to 900 Mb/s. Figure 1b shows how the accumulated data rate of each layer vary with strobe length, the data taking window, and will peak at roughly 1.4 Gb/s. As discussed in Section 4, these rates are manageable by a single readout unit (pRU), which is dedicated to handling the data flow from a complete layer.

The sensors chips are bonded to thin, flexible printed circuits (FPC) made of aluminum and polyimide (30 μm/20 μm) with ultrasonic welding using single-point tape automated bonding. This technique reduces the overall material budget and provides high mechanical reliability [4]. Most importantly, it allows for a homogenous structure, compared to other bonding techniques. Nine chip-cables are bonded to flexible PCBs called strings, which are further mounted to carrier modules with the absorber. A complete layer is made of twelve 9-chip strings.
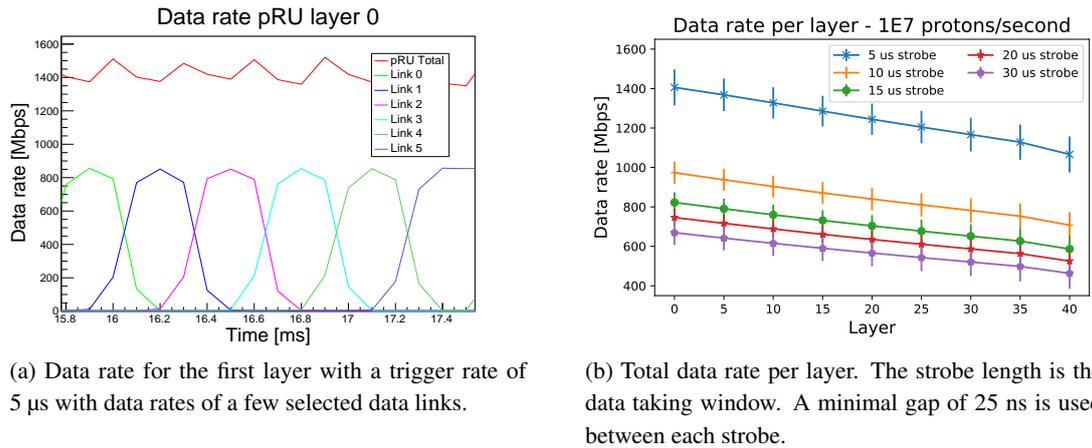
(a) Data rate for the first layer with a trigger rate of 5 µs with data rates of a few selected data links.

(b) Total data rate per layer. The strobe length is the data taking window. A minimal gap of 25 ns is used between each strobe.

Figure 1: Monte Carlo simulation of data rates of a 230 MeV proton scanning beam with an intensity of 10^7 s^{-1}. The beam scans over the detector plane in 65 ms.

A transition card is placed between the FEE and the pRU. This position allows for placing the pRU further from the ion beam axis which reduces the need for single-event upset mitigation. Furthermore, power regulators can be placed on the transition card and thus closer to the detector chips, providing better power integrity. Most importantly a transition card can be made with a low enough pitch that allows for the relatively thin spacing of the layers. The pRU will consist of a single Xilinx Kintex Ultrascale FPGA interfacing all the sensors in a layer. As data-taking only goes on for a few seconds, no high-level trigger system is applied. One pRU acts as a master and will handle synchronization and initiate triggering with the other pRUs.
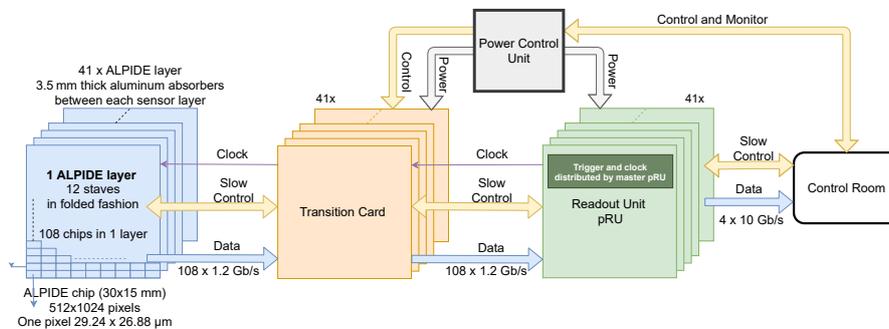

Figure 2: pCT Readout System Overview.

3. Data Recovery

Each ALPIDE chip has an 8B10B-encoded 1.2 Gb/s differential serial link. Thus, one of the challenges of the readout system is to handle all 108 high-speed links for each layer. Multi-Gigabit Transceiver pins (MGT) are the most common method for clock data recovery on FPGAs. With a high number of high-speed links, it is infeasible to use MGTs as the total cost becomes too high, either because of a high number of FPGAs or because of very expensive FPGAs. However, newer families of FPGAs have increased regular I/O pin performance beyond the gigabit range. The

Xilinx Ultrascale I/O pins have a maximum bandwidth of 1250 Mb/s, just within the requirement of the ALPIDE data interface. Xilinx does also provide the fabric logic to automatically do phase-tracking of the incoming data, as explained in [5].

4. High-Speed Data Offload

A high number of data-links restricts the buffer size for each channel as the total FPGA resources are shared between the links. Thus a high data throughput is necessary to avoid back-pressure and data loss in the pRU FPGAs. A Quad Small Form-Factor Pluggable with up to four independent 10 Gb/s Ethernet links was chosen for this task. An offload rate of 10 Gb/s is safely within the simulated data rates, but optional links may be added if e.g. luminosity is increased. The independence of each link allows for full parallelization of the offload process. The incoming data streams are grouped and routed to a priority encoder that selects data appropriately based on buffer usage. The data are formatted in a way that enables the priority encoder to be completely agnostic to data frames.
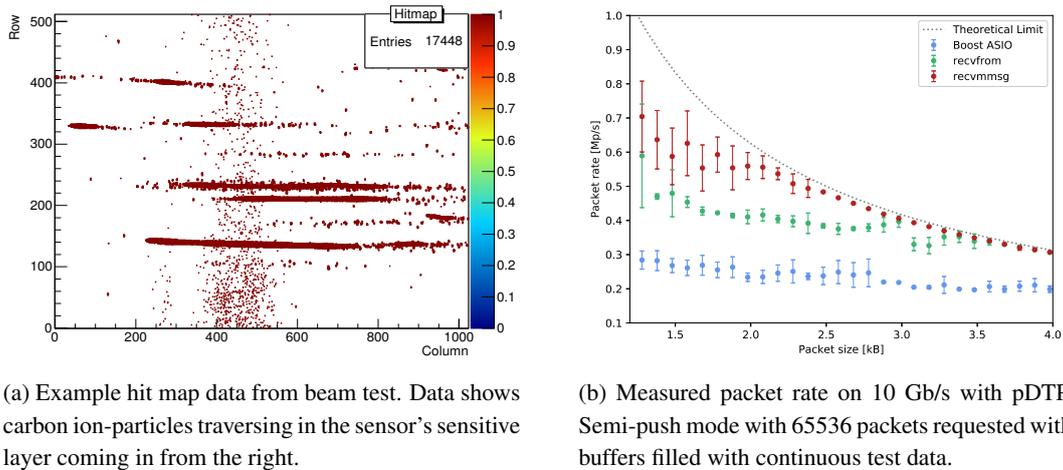
A full User Datagram Protocol (UDP) stack written in Verilog is obtained from an open-source library [6]. Because UDP does not ensure that data will be transmitted safely, a custom protocol is developed on top of UDP to complete this task: the pCT Data Transfer Protocol (pDTP). pDTP behaves as a server and is designed to offload data stored in a given FIFO. The server operates in three different modes: pull, semi-push, and full-push. Each mode provides a varying degree of control. Pull-mode enables the client to request packets of a certain size, and the option to retransmit each packet. The latency in the system limits the throughput in pull mode. In semi-push and full-push, the option of retransmitting packets is lost, but the client may throttle the transmission to avoid packet drops. Semi-push requires the client to periodically request a stream of packets, while full-push is always transmitting when data is available. The client software can mix the use of modes to optimize for various scenarios. E.g., in periods with higher data generation, the client may prefer to risk data loss by using the push-modes to avoid buffer overflows in the pRU.

5. Results

Both laboratory and beam tests have been performed. All tests were done with Xilinx Evaluation Kit VCU118 with a Virtex Ultrascale+ FPGA. The sensor chips are connected via a 2 m Samtec Firefly cable via a custom FMC adapter. A simple desktop computer with a Intel X710 network card is used for control and as an end-point for the high-speed offload.

Regular I/O pin-performance is comparable to MGT-performance when interfacing 1.2 Gb/s links. Simple Pseudo-Random Binary Sequence testing gives a bit-error rate $< 9 \times 10^{-15}$. Figure 3a shows data taken at beam test at the Heidelberg HIT facility. This test shows that the data recovery approach works successfully. Some errors are observed during data collection, but these errors have been identified as a power-integrity issue caused by activity on the sensor chip leading to jitter. These kinds of errors are also observed with the MGT-approach and are being addressed in a future FEE-design.

Data offload performance with pDTP and UDP is satisfactory. Figure 3b shows the result of transmitting over 65k packets for each packet size with three different software approaches. A larger packet size gives the host computer more time to perform checksum-calculations for each packet and thus is the most efficient. Figure 3b also shows that the approaches that are closer

(a) Example hit map data from beam test. Data shows carbon ion-particles traversing in the sensor's sensitive layer coming in from the right.

(b) Measured packet rate on 10 Gb/s with pDTP. Semi-push mode with 65536 packets requested with buffers filled with continuous test data.

Figure 3: Test results.

Table 1: pDTP and UDP Stack Resource Utilization on Xilinx Kintex KU085.

	Slice LUTs	Slice Registers	Block RAM Tiles
pDTP Core	989 (0.2 %)	570 (0.06%)	3.5 (0.06%)
UDP Stack	4924 (1%)	4914 (0.5%)	11 (0.5%)

to the kernel (recvfrom and recvmmsg) have higher efficiency, and that the efficiency is close to the theoretical limit when reading large packets. No packet loss is observed with packets larger than 1 kB and with a tuned Linux-kernel on the receiving computer. The pDTP round-trip time is measured to be ~ 30 µs. This limits the throughput of pull-mode but does not significantly affect either push-mode. Table 1 shows the resources used by pDTP and UDP stack for one independent offload link.

References

[1] A.P. de Haas et al., *The FoCal prototype - an extremely fine-grained electromagnetic calorimeter using CMOS pixel sensors*, *Journal of Instrumentation* **13** *(2018) P01014*.

[2] H. Pettersen, J. Alme, A. Biegun, A. van den Brink, M. Chaar, D. Fehlker et al., *Proton tracking in a high-granularity digital tracking calorimeter for proton ct purposes*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **860** *(2017) 51-67*.

[3] H. E. S. Pettersen, J. Alme, G. G. Barnaföldi, R. Barthel, A. van den Brink, et al., *Design optimization of a pixel-based range telescope for proton computed tomography*, *Physica Medica* **63** *(2019) 87-97*.

[4] V. Borshchov, O. Listratenko, M. Protsenko, I. Tymchuk and O. Fomin, *Innovative microelectronic technologies for high-energy physics experiments*, *Functional Materials* **24** *(2017) 143-153*.

[5] Xilinx Inc., *Native High-Speed I/O Interfaces Application Note (XAPP1274)*, 2017. *URL: xilinx.com/support/documentation/application_notes/xapp1274-native-high-speed-io-interfaces.pdf (Last accessed: 7 October 2019)*

[6] A. Forencich, *Verilog Ethernet components for FPGA implementation*, *URL: github.com/alexforencich/verilog-ethernet (Last accessed: 7 October 2019)*