Upgrades of the ALICE TPC Front-End Electronics for Long Shutdown 1 and 2

Arild Velure

Abstract—This paper presents the front-end electronics upgrades of the ALICE Time Projection Chamber detector for the coming years, with a focus on the upgrades for Long Shutdown 2. The Large Hadron Collider is currently in Long Shutdown 1 following a successful first run, and upgrades of the detectors are underway to support the higher particle interaction rates planned for the next run. For the Time Projection Chamber, the increase in data due to the higher interaction rate and higher energy will cause a bottleneck in the Readout Control Unit and a new board is in development which increases the data-link speed to the back-end. Another more general upgrade of the ALICE experiment is planned for Long Shutdown 2, foreseen to start in 2018. In this case the goal is to cope with an even higher interaction rate of 50 kHz for Pb–Pb collisions. The present Multi Wire Proportional Chambers of the Time Projection Chamber will then be replaced by Gas Electron Multiplier technology. At the same time, the front-end electronics and readout system will also be replaced.

Index Terms—Detector instrumentation, front-end electronics, frontend systems, GEM detectors, mixed signal circuits, particle detectors, readout electronics, real time systems, time projection chambers.

I. INTRODUCTION

THE Large Hadron Collider (LHC) accelerates protons and lead (Pb) ions and collides them at the four experimental areas along the beamline. One of the LHC experiments is A Large Ion Collider Experiment (ALICE) [1].

The ALICE detector is comprised of several sub-detectors, of which one is the Time Projection Chamber (TPC) [2]. The TPC is the main tracking detector of ALICE. It is a 90 m³ gas-filled cylinder with a Multi Wire Proportional Chamber (MWPC) readout on both end plates, with 557568 readout pads. The signals from these pads are passed to 4356 front-end cards (FECs) mounted directly behind each end plate. Data from the FECs are further branched down to 216 Readout Control Units (RCUs) which are then connected to the rest of the ALICE readout chain.

Following a successful running period from November 2009 to January 2013 (Run 1) [3], the LHC is currently shut down for maintenance and preparation for even higher energies and luminosities. This period, named Long Shutdown 1 (LS1), lasts until end of 2014.

Manuscript received June 16, 2014; revised December 05, 2014; accepted December 11, 2014. Date of publication February 05, 2015; date of current version June 12, 2015.

The author is with the Department of Physics and Technology, University of Bergen, NO-5020 Bergen, Norway (e-mail: arild.velure@ift.uib.no).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNS.2014.2382332

Fig. 1. Block diagram of current front-end card [7].

The next running period (Run 2) is between 2015 and 2018. In this period a peak luminosity up to $1 \times 10^{27} - 4 \times 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ can be expected for Pb–Pb collisions with center-of-mass energy of 5.5 TeV per nucleon pair. As a comparison, in the heavy ion run in 2011 the luminosity reached $4 \times 10^{26} \text{ cm}^{-2} \text{ s}^{-1}$ [3]. The envisaged luminosities for Run 2 correspond to interaction rates from 8 to 30 kHz.

To get maximum benefit from the delivered luminosities, the most attractive data taking scenario for ALICE foresees data taking rates of at least a factor two higher than what can be achieved with the current TPC readout electronics. For a trigger mix containing central, semi-central, minimum bias, calorimeter and Transition Radiation Detector (TRD) triggers, an event readout rate of at least 400 Hz is envisaged for the ALICE central barrel detectors, while keeping a reasonably low dead fraction (busy time). In comparison, the maximum achieved readout rate for the same trigger mix of the current TPC readout electronics has been measured to be about 320 Hz (in this case with high dead fraction). The higher luminosities for Run 2 will also lead to a 25% increase in the event sizes since the number of tracks per interaction will be higher [4]. The maximum event size for a central event in Run 1 is 65 MB, for Run 2 a corresponding event will produce 81 MB of data.

Run 3 follows LS2 in 2018, with an expected peak luminosity of $6 \times 10^{27} \text{ cm}^{-2} \text{ s}^{-1}$ and interaction rates of about 50 kHz for Pb–Pb collisions [5]. Since the TPC drift time of about 100 μs is 5 times longer than the average time between interactions, the presently employed gating of the TPC wire chambers is not sufficient and will be replaced by Gas Electron Multipliers (GEMs). At the same time the front-end electronics needs to be replaced to match the new readout chamber technology and increased data rates.

Fig. 2. Overview of the present TPC readout electronics with the original RCU [6]. The Detector Control System (DCS) and Source Interface Unit (SIU) are add-on boards handling the specific interfaces.

This paper presents the upgrades of the readout electronics for the ALICE TPC detector for the coming years, with a focus on the upgrades for LS2. More in-depth discussion on the upgrades for LS1 can be found in [6].

II. PRESENT TPC ELECTRONICS

A single readout channel, as shown in Fig. 1, is comprised of three basic functional units:

- A charge sensitive amplifier/shaper
- A 10-bit 10 MSps low power Analog Digital Converter (ADC)
- A digital circuit that contains a tail cancellation, baseline subtraction and zero-suppression filter, and a multiple-event buffer

The amplifier stage is contained in a single chip named PASA (PreAmplifier and ShAper) [8], while the ADC and digital circuits are contained in the ALTRO (ALICE TPC ReadOut) [9] chip. Each chip handles 16 input channels for a total of 128 channels per front-end card, where each channel corresponds to a single pad on the TPC detector.

Each front-end card is connected to an RCU, each controlling between 18 and 25 front-end cards depending on the radial position of the RCU in the TPC barrel. The connectivity between the RCU and the front-end cards is implemented using two branches of a parallel, multidrop bus with a bandwidth of 1.6 Gbps per branch. The acquired data is sent to the Data Acquisition System (DAQ) over a 1.28 Gbps optical link (Detector Data Link). A sketch can be seen in Fig. 2.

The current system acquires data on a double-trigger scheme, where a trigger is sent to the front-end cards at each interaction, and a second trigger is sent within $\sim 100 \mu\text{s}$ after the first, indicating whether the data from the current interaction should be kept.

A. Motivation for the Run 2 Upgrade

The present TPC readout electronics will be a limiting factor with the foreseen readout rate for Run 2. It is able to handle an event readout rate of 250 Hz for central Pb–Pb collisions (the readout rate depends on the number of tracks per event) [10], while an event readout rate of at least 400 Hz is envisaged while keeping a reasonably low dead fraction (busy time). In addition,

Fig. 3. Overview of RCU2 highlighting the main differences from the original RCU [6]. Radmon is a radiation monitor based on SRAMs.

stability issues have been observed during Run 1, some of which can be traced to single event upsets (SEUs) in the SRAM based FPGAs. As there is a lack of spare resources on the FPGA, these problems can not be mitigated on the firmware design level.

Since the time available for development and installation was limited, a solution was chosen which involved redesigning the RCU board and splitting the two existing buses connecting the RCU to the front-end cards into four buses in addition to changing the readout-link of the RCU to a higher bandwidth. Assuming there are no other bottlenecks in the system, this by itself would double the readout rate from 250 Hz to 500 Hz and if we take into account the increased event size of 25%, gives us a final readout rate of 400 Hz which is sufficient for the envisaged Run 2 scenario. This solution also allows reuse of the existing cabling and front-end cards, minimizing the amount of installation work needed.

B. Motivation for the Run 3 Upgrade

The drift time of electrons in the Ne – CO₂ (90-10) gas currently used in the TPC is $\sim 100 \text{ s}$. The drift time of positive ions from the amplification region around the MWPC anode wires is $\sim 180 \text{ s}$. To avoid any ions drifting back from the amplification region to the drift region, a gating grid is in place which closes after the initial electron drift time. This leads to an intrinsic dead time of $\sim 280 \text{ s}$ and a limitation in the maximum interaction rate of 3.5 kHz. If the gating grid structure would be disabled and the TPC ran at the targeted interaction rate of 50 kHz (20 μs) for Run 3, space charge distortions due to the accumulated ions in the drift region would occur, which would render path reconstructions useless. Operation of the TPC at 50 kHz can thus not be accomplished with the current gating scheme. It will be replaced by a multi-stage GEM system. GEMs have been proven to operate reliably in high-rate applications and provide intrinsic ion blocking capabilities, therefore enabling the TPC to operate in a continuous, ungated readout mode.

As the current readout electronics is based around a triggered readout, a complete redesign of the whole front-end chain is needed. The new electronics must implement a continuous readout scheme and should be able to handle the resulting higher readout rate. In addition it should accommodate both the negative signal polarity of the new GEM detectors and the lower gas gain which demands a low noise design.

Fig. 4. Simulated TPC Run 2 readout performance at 5 Gbps for Pb–Pb collisions for two different modes (sparse and full readout).

III. UPGRADES FOR RUN 2

The RCU2 project was initiated only in April 2013, giving less than 20 months for completion, including design, prototyping, mass-production and commissioning.

The cables for Ethernet, Trigger, DAQ and power will be reused as is, due to the demanding time-schedule. However, each front-end card bus is split into four branches per RCU2 instead of two, see Fig. 3. This ensures at least a doubling of the data rate, requiring an upgraded Detector Data Link (DDL) as well. With use of the same fiber, this is updated to the DDL2 protocol [11] with a bandwidth of 3.125 Gbps or 5 Gbps.

Simulations in SystemC have been completed to confirm and verify the feasibility of the new readout solution. The simulation is based on real data recorded with the TPC in Run 1 and simulates the readout time per event in relation to the number of tracks in the event. The simulation result in Fig. 4 is shown for both full readout, where all channels are read out even if empty, and sparse readout, where only channels containing data are read out. Sparse readout is only beneficial for small events with only a few tracks.

When comparing the simulated readout time for the highest numbers of tracks in Fig. 4 of ~ 1.6 ms with the current readout time for central events of ~ 4 ms [10] we get an improvement factor of 2.5. If we account for the increase of 25% in event size we reach a readout rate of 500 Hz which is sufficient for the envisaged Run 2 scenario of 400 Hz.

The radiation related stability issues have been mitigated by replacing the previous SRAM based FPGAs with a Microsemi SmartFusion2 (SF2) System-on-Chip (SoC) FPGA [12], [13]. This state of the art flash based FPGA incorporates SEU immune configuration memory, a ARM Cortex-M3 [14] processor and error detection and correction (EDAC) DDR memory in addition to other radiation tolerance measures.

Several irradiation campaigns have been undertaken to ensure that the radiation tolerance of the TPC readout chain is improved compared to the existing RCU. The first campaigns

Fig. 5. Overview of the TPC readout architecture for Run 3 [4].

identified certain problem areas, which lead to a substantial delay of the project [15]. These challenges, including single event latchups in the SF2 FPGA and a non satisfactory fault tolerance of the Trigger, Timing and Control interface, has been dealt with. An irradiation campaign was conducted very recently to verify that the performed adjustments are improving the stability of the RCU2. The first analysis of the results from this campaign are positive, and the overall radiation tolerance of the RCU2 seems to now be at an acceptable level. However, the delay of the RCU2 means that Run 2 will start with the existing RCUs in spring 2015. Only one or two sectors of the TPC will have RCU2s installed in December this year. The rest is planned to be installed during the first technical stop in May 2015.

DAQ communication has been extensively tested at a speed of 4.2 Gbps without errors, Ethernet communication to the DCS is stable, and the optical interface to the Trigger, Timing and Control system is verified. Additionally, several thousand transactions have been done towards the front-end cards without any errors. Unfortunately, the delays of the project have affected the readiness of the FPGA design, and some integration work is still remaining.

IV. UPGRADES FOR RUN 3

A new front-end Application Specific Integrated Circuit (ASIC), named SAMPA, is currently in development. It has been designed to comply with the requirements for Run 3 of both the TPC as well as the ALICE Muon tracking detector.

A. Common Front-End ASIC

The new SAMPA chip combines the functionality of the previous PASA [8] and ALTRO [9] chips into one in addition to doubling the number of channels from 16 to 32, supporting bipolar input signals and adding the possibility to Run in a continuous readout mode. The ASIC will be produced in a 130 nm CMOS technology.

The acquired data from the SAMPA is transferred at 1.2 Gbps over four 320 Mbps serial links to a Gigabit Transceiver (GBTx) ASIC [16], which multiplexes the input from different SAMPA chips into one high-speed 4.8 Gbps versatile optical link component [17] to a Common Readout Unit (CRU), see Section IV-A. A sketch of the setup for Run 3 can be seen in Fig. 5.

The SAMPA ASIC consists of 32 identical channels consisting of a Charge Sensitive Amplifier (CSA), a pulse shaper,

Fig. 6. SAMPA system block diagram [5].

an ADC and a Digital Signal Processor (DSP). A sketch of the SAMPA ASIC can be seen in Fig. 6.

The signals received from the TPC pads are very small and fast, typically in the order of 10 fC and 10 ns, requiring pre-processing before they can be digitized. The CSA converts the incoming signal to a pulse with a long tail that has an amplitude proportional to the total charge of the incoming signal. In the shaper, the undershoot created by the CSA is removed in addition to any low frequency components introduced.

The design of the CSA and shaper is based on the previous PASA design, but contains some additional options:

- Selectable polarity of the input signal is provided to adapt to different detectors.
- Configurable gain for the CSA is provided for the different detector applications (TPC and Muon tracking). Moreover, this feature can be used to compensate for differences in length of the pads in the inner and outer readout chambers of the TPC (larger pads will collect more ionization due to the longer track length seen by the pad).
- In case the gas mixture in the TPC would be changed to a gas with faster electron drift velocities (to combat the problems of signal pileup) a configuration to set a lower shaping time is also implemented.
- Additionally, specific gain and shaping settings are available to satisfy the requirements of the ALICE Muon tracking detector.

The ADC has been implemented as a 10-bit capacitive successive-approximation (SAR) ADC [18] optimized for operation at 10 MSPs, but optionally up to 20 MSPs. Using this type of SAR ADC significantly reduces the power consumption compared to the pipelined ADC implemented in the previous ALTRO chip.

The digital signal processing (DSP) chain can be divided into 5 steps:

- The first, baseline correction 1, contains logic to subtract systematic perturbations of the signal baseline. These shifts could be a result of manufacturing/process variations, slow changes in the environment (e.g. temperature and voltage), low frequency (e.g. power line interference) or triggers (e.g. induced by a pulsed gating grid). The first three effects are solved with a combination of the subtraction of a constant value and a self-calibrating value. The trigger induced effects are solved by subtracting a value stored in a 1024

samples long internal memory holding a time dependent pattern for each channel. This mode is though not useful in continuous readout mode, but the internal memory can also be used to override the signal chain and issue test-patterns for calibration purposes of the readout system.

- To combat the long tail generated by the signal in some detectors, a tail cancellation filter is implemented which uses a four stage Infinite Impulse Response (IIR) filter to smooth the tail enabling a better zero suppression of the data. This feature is not necessary for the new GEM readout as there is no ion tail.
- The third, baseline correction 2, contains logic to subtract non-systematic perturbations and slow baseline drifts. These would be caused by noise and slowly varying signals, like the ion tail for the current MWPC or due to cross-over effects from capacitive coupling between the pads for the GEM. The effects are removed through the use of a moving average filter.
- As the SAMPA chip can effectively produce 3.2 Gbps of data from its 32 channels and the serial outputs can only handle 1.2 Gbps a lossy compression method is implemented. This method, referred to as zero suppression, operates on a fixed threshold pulse detection scheme, where samples of a value smaller than a constant decision level (threshold) are suppressed. Together with the storage method, a form of Run length encoding, a compression ratio of about 0.18 to 0.5 for pad occupancies¹ of 14% and 42% should be possible [4].
- The resulting data stream is transferred into a channel specific First-In, First-Out (FIFO) event memory. The data is collected into packets which are buffered until a readout command is issued to the specific channel by one of the serial output links and the data is transmitted to the CRU.

The DSP can operate in both a triggered and continuous readout mode. In triggered mode, a configurable amount of samples, defined as a time window, are acquired upon the reception of an external trigger signal, while in continuous readout a new time window is started at the end of the previous.

The four serial output links from the SAMPA have been optimized for use in conjunction with the GBTx ASIC. For use in regions of lower occupancy or with other detectors with low data volumes, two or three of the serial links can be disabled. For applications with very low data-rate, up to 16 devices can be daisy-chained on a single link.

The GBTx can handle up to 10 serial links of 320 Mbps including error correction, giving 2.5 SAMPAs per GBTx. The receiving part of the GBTx is used for supplying system level synchronization and triggering signals in case of running in triggered mode.

The SAMPA implements an I²C interface for control and monitoring. This is used in conjunction with the GBT Slow Control ASIC (GBT-SCA) [19], which has a dedicated data link on the GBT interface, acting as a board controller for configuring all the SAMPA chips on the same front-end board.

¹Pad occupancy is the fraction of samples within a given time window exceeding the zero suppression threshold.

TABLE I
SIMULATED SPECIFICATIONS FOR THE SAMPA ASIC [5]

Specification	TPC	Muon Chambers
Voltage supply	1.25 V	1.25 V
Polarity	Positive/Negative	Positive/Negative
Detector capacitance (Cd)	18.5 pF	40 pF - 80 pF
Peaking time (ts)	80 ns or 160 ns	300 ns
Shaping order	4th	4th
Noise (ENC)	<536 e@ $t_s=80$ ns* <482 e@ $t_s=160$ ns*	<950 e@ $C_d=40$ p* <1600 e@ $C_d=80$ p*
Linear Range	100 fC / 67 fC	500 fC
Sensitivity	20 mV/fC / 30 mV/fC	4 mV/fC
Return to baseline time	<164 ns@ $t_s=80$ ns <288 ns@ $t_s=160$ ns	<541 ns
Non-Linearity (CSA + Shaper)	<1 %	<1 %
Crosstalk	<0.3 %@ $t_s=80$ ns <0.2 %@ $t_s=160$ ns	<0.2 %@ $t_s=300$ ns
ADC effective input range	2 Vpp	2 Vpp
ADC resolution	10-bit	10-bit
Sampling Frequency	10 MSps or 20 MSps	10 MSps
INL (ADC)	<0.65 LSB	<0.65 LSB
DNL (ADC)	<0.6 LSB	<0.6 LSB
SFDR (ADC)**	68 dBc	68 dBc
SINAD (ADC)**	57 dB	57 dB
ENOB (ADC)	9.2 bit	9.2 bit
Power consumption (per ch)		
ADC	2 mW (4 mW)	2 mW (4 mW)
CSA + Shaper	6 mW	6 mW
Channels per chip	32	32

* $R_{\text{esd}} = 70 \Omega$

** @ 0.5 MHz, 10 MSps

Preliminary specifications for the SAMPA ASIC can be seen in Table I.

B. Common Readout Unit

The CRU is another common ALICE development, aimed to be used for most of the ALICE detectors in the future. It acts as the interface between the front-end electronics, the online and offline computing system and the central trigger processor and is based on a high performance FPGA equipped with multiple gigabit optical inputs and outputs.

The CRU handles and controls the readout, configuration and monitoring of the front-end cards. The per-channel data packets received from the SAMPA chips are transferred verbatim to the online system, but to facilitate reconstruction of the data in the online system the packets are temporarily buffered and re-ordered according to their geometrical position in the pad row on the detector end plate.

Placing the CRU in the detector hall, close to the front-end electronics, would give the benefit of reduced material costs due to the reduced number of optical links needed. The CRU would be connected to the front-end electronics through copper cables and communicate with the online system through faster 10 GbE Detector Data Link 3 (DDL3) links.

The drawback of this solution is that it restricts the design to using low performance radiation tolerant FPGAs. Additionally there is the consideration of limited access during LHC operation, difficult installation and maintenance, extra design work to make the design radiation tolerant and the relatively costly radiation-verification campaign of the electronic components.

Locating the CRU in the control room outside of the radiation area presents thus a cleaner and more robust solution and enables the use of commercial off-the-shelf hardware.

V. OUTLOOK AND CONCLUSION

A new board, the RCU2, which rectifies most of the current limitations of the original RCU in addition to satisfy the higher readout-speeds needed for Run 2 is currently in the prototyping stage aimed to be fully installed in during the first technical stop in May 2015.

For Run 3 a new front-end ASIC, the SAMPA, which satisfies the requirements for running with continuous mode and being compatible with the new GEM readout has been described. In combination with the CRU this creates a full new readout system for the TPC detector.

REFERENCES

- [1] The Alice Collaboration, "The ALICE Experiment at the CERN LHC," *J. Instrum.* vol. 3, no. 08, p. S08002, 2008.
- [2] J. Alme *et al.*, "The ALICE TPC, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events," *Nucl. Instrum. Methods Phys. Res. A*, vol. 622, no. 1, pp. 316–367, 2010.
- [3] K. Šafařík, "Overview of recent ALICE results," *Nucl. Phys. A*, vol. 904–905, pp. 27c–34c, 2013.
- [4] Upgrade of the ALICE Time Projection Chamber, Tech. Des. Rep., CERN-LHCC-2013-020, ALICE-TDR-016, 2013, The ALICE Collaboration.
- [5] Upgrade of the ALICE Read-Out and Trigger System, Tech. Des. Rep. CERN-LHCC-2013-019, ALICE-TDR-015, 2013, The ALICE Collaboration.
- [6] J. Alme *et al.*, "RCU2 - The ALICE TPC readout electronics consolidation for Run2," *J. Instrum.*, vol. 8, no. 12, p. C12032, 2013.
- [7] L. Musa *et al.*, "The ALICE TPC front end electronics," in *Proc. IEEE Nuclear Science Symp. Conf. Rec.*, 2003, vol. 5, pp. 3647–3651.
- [8] H. Soltveit *et al.*, "The preamplifier shaper for the ALICE TPC detector," *Nucl. Instrum. Methods Phys. Res. A*, vol. 676, pp. 106–119, 2012.
- [9] R. Bosch *et al.*, "The ALTRO chip: A 16-channel A/D converter and digital processor for gas detectors," *IEEE Trans. Nucl. Sci.*, vol. 50, no. 6, pp. 2460–2469, Dec. 2003.
- [10] A. U. Rehman, "The ALICE TPC readout electronics - design, performance optimization and verification of the DAQ circuit," Ph.D. dissertation, Univ. of Bergen, Bergen, Norway, 2012.
- [11] F. Costa, "DAQ readout after LS1 and LS2," ALICE Week Presentation 2012 [Online]. Available: <https://indico.cern.ch/getFile.py/access?contribId=11&resId=0&materialId=slides&confId=232150>
- [12] IGLOO2 FPGA and SmartFusion2 SoC FPGA Datasheet, 4th ed., Microsemi, Inc., Nov. 2014.
- [13] SmartFusion2 System-on-Chip FPGAs Product Brief, 18th ed., Microsemi, Inc., Oct. 2014.
- [14] SmartFusion2 Microcontroller Subsystem User Guide, 7th ed., Microsemi, Inc., Sep. 2014.
- [15] C. Zhao *et al.*, "First irradiation tests results of the ALICE TPC Readout Control Unit 2," *J. Instrum.*, 2014, accepted for publication.
- [16] P. Moreira, "The GBT project," in *Proc. TWEPP Conf., CERN-2009-006*, 2009, pp. 342–346.
- [17] J. Troska, "The versatile transceiver proof of concept," in *Proc. TWEPP Conf., CERN-2009-006*, 2009, pp. 347–351.
- [18] C.-C. Liu *et al.*, "A 10-bit 50-MS/s SAR ADC with a monotonic capacitor switching procedure," *IEEE J. Solid-State Circuits*, vol. 45, no. 4, pp. 731–740, 2010.
- [19] A. Gabrielli, "The GBT-SCA, a radiation tolerant ASIC for detector control applications in SLHC experiments," in *Proc. TWEPP Conf., CERN-2009-006*, 2009, pp. 557–560.