Architecture and Implementation of the Front-End Electronics of the Time Projection Chambers in the T2K Experiment

P. Baron, D. Besin, D. Calvet, C. Coquelet, X. De La Broïse, E. Delagnes, F. Druillole, A. Le Coguie, E. Monmarthe, E. Zonca

***Abstract*—The tracker of the near detector in the T2K neutrino oscillation experiment comprises three time projection chambers based on micro-pattern gaseous detectors. A new readout system is being developed to amplify condition and acquire in real time the data produced by the 124,000 detector channels. The cornerstone of the system is a 72-channel application specific integrated circuit which is based on a switched capacitor array. Using analog memories combined to digitization deferred in time enables reducing the initial burstiness of traffic from 50 Tbps to 400 Gbps in a practical manner and with a very low power budget. Modern field programmable gate arrays coupled to commercial digital memories are the next elements in the chain. Multi-gigabit optical links provide 140 Gbps of aggregate bandwidth to carry data outside of the magnet surrounding the detector to concentrator cards that pack data and provide the interface to commercial PCs via a standard Gigabit Ethernet network. We describe the requirements and constraints for this application and justify our technical choices. We detail the design and the performance of several key elements and show the deployment of the front-end electronics on the first time projection chamber where the final tests before installation on-site are being conducted.**

***Index Terms*—front-end electronics, FPGA, fast networks**

I. Introduction

THE Tokai To Kamioka (T2K) experiment aims at the precise determination of several neutrino oscillation parameters. An artificial high intensity neutrino beam formed at the Japan Proton Accelerator Research Center (J-Parc) is directed through the earth to a massive underground water Cherenkov detector located in Kamioka, 295 km away. By measuring the precise beam characteristics at a near station located 280 m from the source and at the far site, evidence of neutrino oscillations such as ν_μ disappearance and ν_e appearance will be searched. The far detector, Super Kamiokande [1] has been in operation for several years and the near detector, Nd280m [2], is presently under construction for a scheduled completion by the end of 2009.

Manuscript received May 20th 2009.

All authors are with Institut de Recherche sur les lois Fondamentales de l'Univers, CEA Saclay F-91191 Gif sur Yvette France. Corresponding author: calvet@hep.saclay.cea.fr.

The Nd280m detector comprises an on-axis detector and off axis detectors mounted inside a magnet re-used from the UA1 and Nomad experiments. The ~50 m^3 space available in the magnet is instrumented with a pi-zero detector, 3 Time Projection Chambers (TPCs), two Fine-Grain Detectors (FGDs), an electro-magnetic calorimeter and muon detectors. A schematic view of one TPC is shown in Fig. 1.

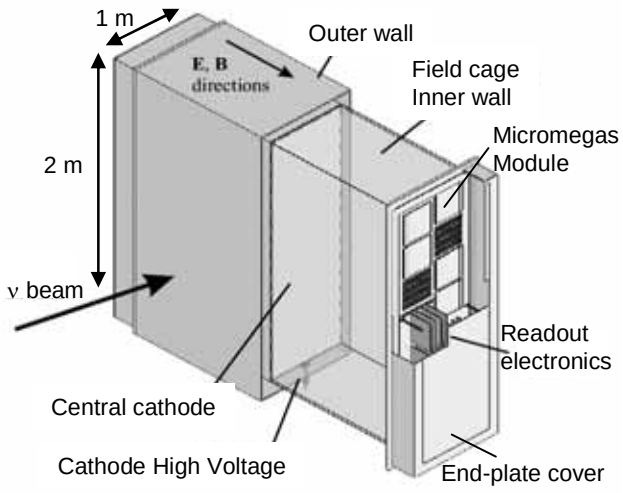

Fig. 1. Schematic view of one of the 3 TPCs.

Each of the 6 TPC end-plates uses 12 pixellated Micromegas modules [3] that cover an effective area of ~1.5 m^2. We present in the design and the implementation of the electronic system to readout these TPCs.

II. System Overview

A. Requirements and Constraints

The required precision for the reconstruction of tracks calls for a segmentation of each of the 72 Micromegas detector modules in 36 × 48 = 1728 pads of 9.8 mm × 7 mm. The typical signal delivered by a pad is a charge of a few tens of femto-coulombs. We require a dynamic range of 10 Minimum Ionizing Particle (MIP) and a signal to noise ratio of 100. The integral non linearity should be less than 1% in the [0; 3 MIP] range and 5% in the [3 MIP; 10 MIP] interval.

The X and Y coordinates of a track are reconstructed by computing the centroïd of the charges collected by the pads hit

along the track, while the Z coordinate is determined by the drift time of electrons in the gas volume of the TPC. We require 500 points in the Z direction for a maximum drift time of 10 µs to 500 µs. The system should therefore support a sampling frequency from 1 MHz to 50 MHz. The sampling of all pads must be synchronous.

The neutrino beam is pulsed; there is one spill every ~3.5 s. The TPCs require an external trigger signal and must be able to capture all beam spills and calibration events (cosmic rays and internal illumination by a laser) at up to 20 Hz. The maximum allowable dead-time for acquiring an event is 50 ms.

In addition to these functional requirements, the front-end part operates in a modest magnetic field (0.2 T) with limited space available, a low power budget and no access during operation. There is no constraint of radiation. The rules for underground experiments in a Japanese nuclear facility apply.

B. The challenge

The large number of channels (124,416) sampled at high rate (typically 33 MHz / 12 bits) leads to a peak dataflow of ~50 Tbps. Data of potential interest are produced in the time interval from a trigger to the maximum drift time of electrons in the gas volume of the TPC, i.e. ~16 µs. The volume of raw data per event is ~90 MB while the allowable size to be kept for off-line analysis is ~250 KB. Data reduction by a factor ~400 must be performed in less than 50 ms.

C. Design methodology

We propose a bottom-up system approach to the design of the TPC readout system. Starting from the detector, we select at each stage the best suited strategy among 3 possible actions: a) process data, b) transport data, c) buffer data. The decision between these 3 actions is made taking into account the technical feasibility and external constraints. The complete data acquisition chain is established by the direct translation into physical elements of this series of actions.

III. DESIGN OF THE ARCHITECTURE

The logical diagram of the complete TPC readout architecture is shown in Fig. 2.

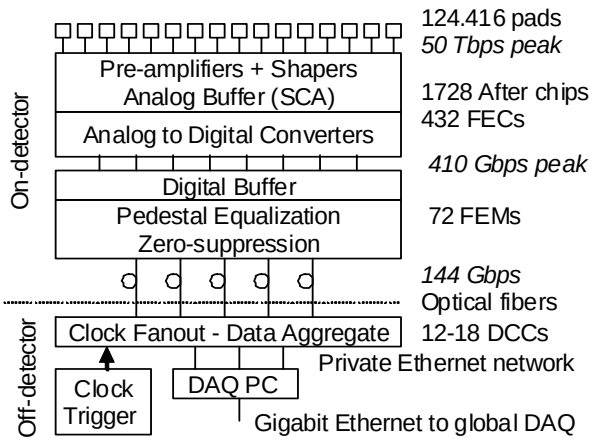

Fig. 2. TPC readout logical flow.

A. Interface to the detector

Each detector pad delivers an analog signal which is extremely sensitive to external perturbations. Although mounting electronics directly at the back of the printed circuit board that makes each detector was thought, it was not found practical for several reasons: i) fabrication would require a complex process ii) the design of the detector could not proceed in parallel with that of electronics, iii) the surface available for mounting components would be constrained by the size of the detector, iv) the electronics could not be replaced without opening the inner gas volume of the TPC. We decided to transport signals over the shortest possible distance, excluding the use of flexible circuits (e.g. Kapton) to reduce cost and minimize signal degradation. The Front-End Cards (FECs) are directly inserted at the back of each detector using fine pitch connectors.

B. Analog front-end

At this stage, only a massive replication of the same analog electronic chain is practical. Performing a direct digitization of all channels after the pre-amplifiers/shapers like in [3] would not fulfill our goals in terms of cost, density and power budget. Decision was made to buffer data in the analog domain using a Switched Capacitor Array (SCA). By de-coupling the write phase into the SCA from the read phase and by multiplexing in time multiple channels (72 in our case) towards the same Analog to Digital Converter (ADC), substantial savings are made. A new multi-channel ASIC, called "AFTER" was devised [5]. To cut design time, the ADC was not integrated in the ASIC. Each FEC comprises four 72-channel AFTER chips. Six FECs are needed per Micromegas detector module.

The burst of data at the SCA stage is ~50 Tbps during ~16 µs. At the output of the ADC stage, it is smoothed: traffic peaks at ~410 Gbps and lasts ~2 ms. When the SCAs are being readout, the front-end is in dead time. No data reduction is performed by this stage, but the initial burstiness is brought to a level that becomes manageable for digital circuits.

C. Digital front-end

Although it would be technically feasible to transport the ~410 Gbps flow of ADC data outside of the detector magnet (where space and electrical power are less constrained), it was judged more advantageous to further reduce the burstiness of data using a digital buffer. For each Micromegas detector module, the volume of raw data per event is 1:72th of a full event, i.e. 1.25 MB. The bandwidth at the output of the ADCs of one detector module is 5.7 Gbps. The receiving logic (de-serializer and digital buffer) is mounted close to the ADCs. After the digital buffer, two strategies can be applied: reduce data locally, or send raw data to external hardware. We only require raw event data for a fraction of calibration events and for diagnosis purposes. For beam events, zero-suppression can be applied. The minimum aggregate bandwidth required to transport raw event data is only 15 Gbps (90 MB per event at 20 Hz), i.e. 200 Mbps per detector module. To avoid electromagnetic interferences with the extremely sensitive

front-end, we decided to use optical links to transport data outside of the magnet. Although 622 Mbps (OC-12) or 1 Gbps (Gigabit Ethernet) optical links would be sufficient, we selected 2 Gbps links (Fiber Channel). This choice allows to fanout in a simple way the 100 MHz experiment-wide reference clock and trigger signals over the return path of the links used for data readout. Field Programmable Gate Arrays (FPGA's) are today's technology of choice to build the bridge between ADCs, memories, high speed links and perform various other functions. We designed the Front-End Mezzanine (FEM) card to perform the digital part of the readout of one detector module. The physical implementation of the readout of one detector module is pictured in Fig. 3. All 72 front-end modules are identical.

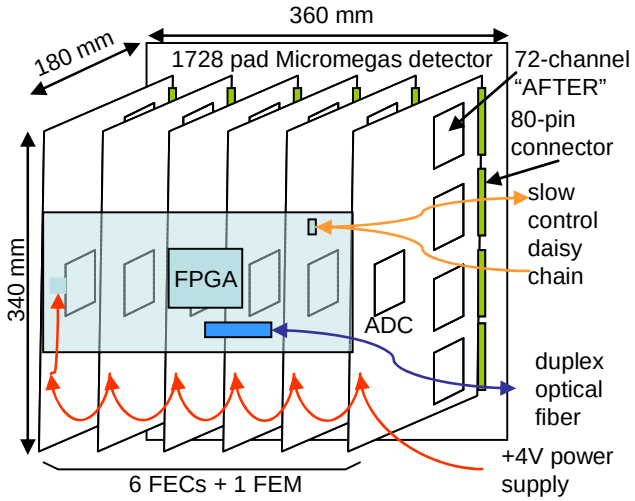

Fig. 3. Front-end electronics of one detector module.

D. *Back-end electronics*

The main roles of the back-end electronics are to aggregate data received from the front-end over seventy-two 2 Gbps optical links (1 per FEM), pack and send formatted data to the shared DAQ system of the experiment. In addition, the back-end electronics is used to distribute the reference clock and trigger signal to the front-end. We split the back-end into 2 components: the Data Concentrator Cards (DCCs) take on one side several optical links from the front-end (4 to 12 depending on design), and interface to a PC via a private Ethernet network. The PC is used to aggregate the flow from all DCCs (18 down to 6 depending on the number of optical links per DCC) and performs the final data encapsulation for the DAQ. The required data throughput to the DAQ is ~5 MB/s, well within the capacity of a single Gigabit Ethernet link.

IV. Implementation and Performance

A. *The "AFTER" chip*

A detailed description of the 72-channel AFTER chip is presented in [5]. Each channel comprises: a Charge Sensitive Amplifier (CSA), a Pole-Zero Cancellation (PZC), a second order low pass filter and a gain-2 amplifier. This is followed by a 511-bucket SCA. Four gains are available (120 fC, 240 fC, 360 fC and 600 fC) and 16 values of peaking time from 100 ns to 2 µs. Selecting external polarization resistors allows operation with positive or negative input signals.

The chip is designed in 0.35 µm CMOS technology. The 500.000 transistor, 7.8 mm x 7.4 mm die is housed in a 0.65 mm pitch 160-pin low-profile quad flat package (LQFP). Five thousands chips have been produced with a yield of ~89%.

The AFTER chip operates at up to 50 MHz write rate; readout is done at 20 MHz. The integral non linearity for 100 ns peaking time in the 120 fC range is less than 1.2%. The Equivalent Noise Charge in operating conditions (120 fC range, 200 ns peaking time, 22 pF detector pad capacitance) is ~850 e^- r.m.s. The internal cross-talk on adjacent channels is ~1%. The gain across different chips is homogeneous to ~1.7% r.m.s. The peak-to-peak spread of the baseline across chips is sufficiently low (~360 LSBs) to avoid the need of analog compensation. However, this spread causes a 10% reduction in the exploitable dynamic range of the ADC. The chip consumes ~8.5 mW per channel. A comparable device, the Altro [4], which comprises one 10-bit ADC per channel, requires two packages for 16-channels and draws ~20 mW per channel. Using the SCA approach allows more than a 4-fold increase in density for less than half of the power budget. Contrary to the Altro, our solution is not dead time free: digitizing the SCA introduces an incompressible dead time of ~4 µs per time bucket (i.e. 2 ms for the whole array).

B. *The Front-End Card (FEC)*

The FEC comprises 4 AFTER chips with external passive protection circuits that connect to the detector via 4 dual row right angled 80-point 1.27 mm pitch surface mount connectors. Precise mechanics is needed for insertion and extraction. The output of the AFTER chips is digitized by a 4-channel 12-bit ADC clocked at 20 MHz. We selected the AD9229 from Analog Devices for compactness and because it features an extremely fast wakeup time (typically ~35 µs) after the sampling clock is applied. To avoid interference with the charge sensitive pre-amplifiers, the clock of the ADCs is turned off when the font-end is sampling detector signals. The complete digitization of SCAs takes ~2 ms, hence the device places an upper limit on the event rate of ~500 Hz.

The FEC includes an on-board pulser for calibration and a serial identification chip (Maxim DS2438) that is also used to measure the current, supply voltage and temperature of the card. The FEC is essentially a slave device that is driven by the FEM board. Each FEC consumes 1 A under 4 V. The on-board 3.3V voltage regulator and the ADC dissipate 0.7 W and 0.85 W respectively.

The average r.m.s noise measured for the complete chain in operating conditions on the TPC field cage is less than 5 ADC counts. We show on Fig. 4 a typical map of the r.m.s noise of the 1728 channels (48 × 36) of one detector module. A slight noise increase is observed for pads that have a longer trace – hence a higher parasitic capacitance – between the detector and the corresponding input pin on an AFTER chip.

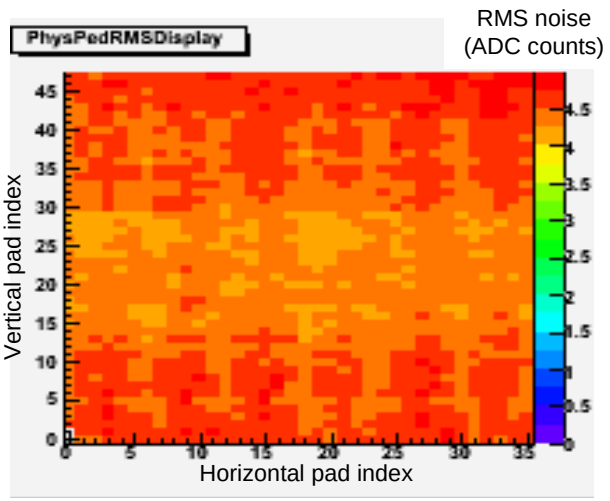

Fig. 4. Noise map of a typical detector module of a TPC.

C. *The Front-End Mezzanine card (FEM)*

The FEM card aggregates data from the 6 FECs of a detector module. Initially, each detector module was meant to be segmented into 1024 channels readout by 4 FECs comprising four 64-channel AFTER chips each. Studies showed that a finer segmentation of the detector is required and we decided to increase the number of channel to 72 per AFTER chip and use 6 FECs instead of 4 per module. This allowed increasing by ~70% the number of channels for only ~40% extra cost.

1) *Interface to the FECs*

Each FEM receives data from 6 quad-channel ADCs (data throughput: 5.76 Gbps total) and stores them in a 18-Mbit static RAM as detailed in [6]. The central element of the FEM is a Xilinx Virtex 2 Pro FPGA (XC2VP4) coupled to two 256K x 36 bit Zero Bus Turnaround (ZBT) static RAM clocked at 120 MHz.

2) *Data processing*

During SCA digitization, the FEM stores in the SRAM buffer event data in raw format, i.e. 1728 × 511 twelve bit ADC samples. During the data acquisition phase by the DCC, all the samples of each given channel can be acquired, or only those above a per-channel threshold. To avoid information loss on the edges of waveforms, 10 time buckets preceding a threshold cross and 4 trailing time buckets are kept. The FEM can subtract a per channel programmable pedestal. No further data reduction or processing is performed at the FEM level. The pedestal correction and zero-suppression are implemented in pipe lined logic clocked at 60 MHz. It takes ~520 clock cycles (17 ns) to process all the time bins of one channel and it is followed by up to ~520 cycles (10 ns) to serialize data over the optical link. The upper limit on the event rate at the level of the FEMs is ~58 Hz if no data are above threshold and ~37 Hz for full readout without zero-suppression.

3) *Interface to back-end electronics*

Communication with the DCC uses one of the RocketIO gigabit transceivers embedded in the FPGA of the FEM. It is coupled to a 2 Gbps optical transceiver. The data path from the FEM to its DCC is logically a 16-bit bus clocked at 100 MHz. We use 8B/10B encoding: K30.7 and K23.7 comma characters are the start-of-frame and end-of-frame delimiters respectively. Standard CRC32 is used for error detection. The DCC to FEM path is logically split into two 8-bit wide busses clocked at 100 MHz. The lower byte contains time synchronous information updated at each clock cycle: trigger bit, trigger type, event count and time stamp reset signals, and a parity bit computed over the 2 bytes. The upper byte is used to send asynchronous commands from the DCC to the frontend. Five consecutive bytes are accumulated to build one transaction: the first byte specifies the direction of the transfer (read/write), the next 2 bytes give an address and the last 2 bytes are the operand. This implements a 64K address space / 16-bit data virtual bus.

In addition to the fast optical path, the FEM also comprises a robust path (CANbus) for monitoring and in-situ firmware upgrade. The FEM uses a Silicon Laboratories C8051F040 8-bit micro-controller to power On/Off each FEC and the FPGA of the FEM independently and monitor the voltage, current, and temperature of all boards attached to a detector module. This low cost device features 64 digital I/Os, several embedded ADCs, a temperature sensor and a CANbus controller. Opto-isolators are placed on the CANbus drivers. Twenty-four FEMs (i.e. the complete readout of one TPC) are chained in the same CANbus segment. Tests showed stable operation at 500 kbps over 20 m of cable. Upgrading the FPGA firmware in the PROM of a FEM takes ~3 minutes over the CANbus link.

The FEM draws 0.75 A from a single 4 V input. To minimize switching noise, only Low voltage DropOut (LDOs) are used to produce the voltages required on board (3.3 V, 2.5 V and 1.5 V). About one third of the 3 W dissipated by the FEM is wasted in the LDOs. The global power dissipation of the front-end electronics for the 3 TPCs is ~2 kW, i.e. 16 mW in average per channel. Because this system is mounted in confined space, a water cooling system is needed.

D. *The back-end electronics*

The back-end electronics is essentially a 72-input, 144 Gbps aggregate bandwidth data concentrator in the FEM to DCC direction and a 1-to-72 trigger and reference clock fanout in the opposite direction. At present, no data reduction (beyond lossless compression) is thought to be made by the back-end electronics. The desired 1:400 data reduction will simply be achieved by the thresholds applied by the FEMs given the expected low channel occupancy. Evolutions of the back-end electronics will consider more sophisticated data reduction methods, and the ability to acquire the data of pads under threshold that are neighboring pads over threshold.

1) *Clock and trigger fanout*

The sampling of all detector signals should be synchronous to the 100 MHz global clock of the experiment. For simplicity, we decided to use in the front end the clock recovered from optical link driven by the DCC. The FEM uses a synchronous divider on the clock recovered from the RocketIO transceiver to make the sampling clock of the SCAs. The drawback of this

approach comes from the fact that the latency of the RocketIO path in a Xilinx Virtex 2 Pro cannot be precisely controlled: the recovered clock on the receiver side exhibits an unknown phase shift with respect to the clock of the transmitter. Skew remains constant during operation, but each time the link is reset, a new value, in [0; 10 ns] (100 MHz primary clock), is reached. Given that the sampling frequency of the front-end is 33 MHz, an imprecision of 10 ns (i.e. 1:3rd of a sampling clock period) was found acceptable. The phase offset will be extracted by off-line data analysis. The jitter of the sampling clock of an AFTER chip is ~50 ps r.m.s. The jitter histogram exhibits two narrow peaks ~200 ps apart in some cases. It is suspected that a very low frequency component is present. This aspect is still being investigated although it does not seem to cause any performance degradation.

2) Implementation of the DCC

The DCC is under the responsibility of collaborators and final boards are still under development. At present, commercial Xilinx ML405 [7] evaluation boards are used as an interim solution.

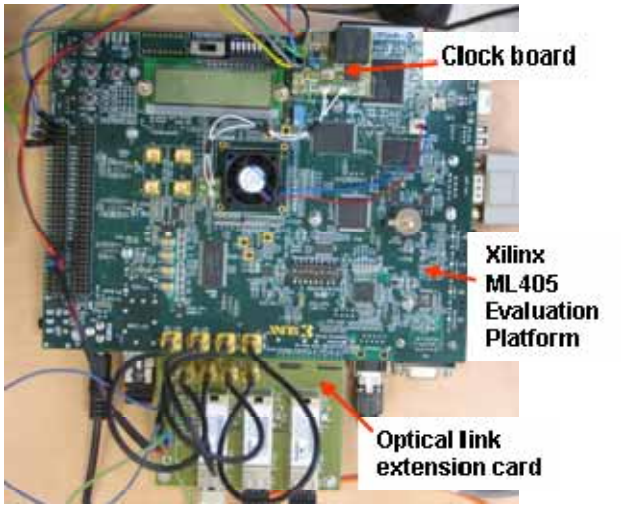

Fig. 5. DCC based on Xilinx ML405 evaluation board.

We built a small add-on card to bring the 100 MHz external clock and trigger to a ML405 board. After jitter reduction with a PLL (National Semiconductor LMK03000) the reference clock is fed to both sides of the Virtex 4 FPGA of the ML405. An off-the-shelf ML405 board features only one optical transceiver. We designed an adapter board with 3 optical transceivers to exploit the RocketIO routed to SMA and SATA connectors. Although the setup shown in Fig. 5 is not ideal from the engineering point of view, it provides the 8 Gbps I/O bandwidth required to readout 4 FEMs at very low cost and with minimal design effort. Scaling up the ML405 scheme to the readout of the 3 TPCs (i.e. 18 quad optical link ML405) is planned if no better solution is made available in time.

3) Firmware and embedded software

Data transfers from the FEMs to the DCC are controlled by a finite state machine implemented in hardware. Each channel is requested in turn but all FEMs are questioned in parallel. To collect a complete event, the DCC multicasts 1728 requests to the FEMs it controls. A FEM is allowed not to respond to its DCC if no data above threshold is found for the current channel. At the end of the data gathering phase, the DCC needs to check that each FEM received the expected number of requests to make sure that missing responses really correspond to empty channels. The data received from each FEM are buffered in a 1K deep 32-bit wide FIFO accessible by the embedded PowerPC 405 processor via the Processor Local Bus (PLB). The processor checks the consistency of the received data, strips duplicated headers and encapsulate meaningful data in UDP/IP frames before they are sent to the data acquisition PC via the Gigabit Ethernet Media Access Controller (MAC) embedded in the Virtex 4 FPGA. The logic requesting and receiving data from the FEMs runs concurrently with the software task that unloads the received data and sends them to the Ethernet MAC.

4) Performance of the DCC

The event readout time achievable by one DCC versus event size is shown in Fig. 6. The DAQ PC does not process data and simply acts as a data sink.

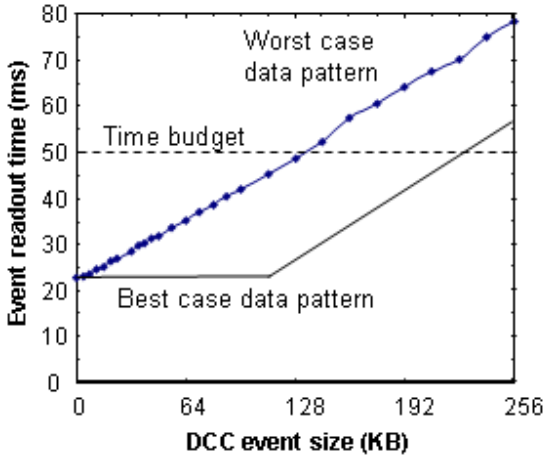

Fig. 6. Event readout time of a ML405 based quad-link DCC.

The overhead of the data collection protocol reduces the sustainable event rate for minimum size events from ~58 Hz (FEM limit) to ~44 Hz (i.e. 23 ms handling time). The event acquisition time is independent of the number of FEMs and it increases linearly with a ~5 MB/s slope for the fraction of data where transmission to the DAQ PC does not completely overlap with concurrent data transfers from the FEMs. When consecutive channels are hit, a worst case traffic pattern is induced because the task requesting data from the FEMs often has to wait that room is made available by the processor in the receiving FIFOs. When hit channels are evenly spread, the processor is sufficiently fast to avoid that the FIFOs fill-up and the state machine gathering data from the FEMs never gets blocked. At best, the theoretical achievable readout time is 23 ms for events up to ~120 kB. At worst, the target rate of 20 Hz is achieved for events up to ~128 kB per DCC. In T2K, the expected event fragment size is up to ~64 kB per DCC.

The limited I/O performance of the PPC405 processor over the PLB is an important factor: in the current version (PLB v4.6), only single beat 32-bit transfers are supported leading to a typical throughput of ~20 MB/s for a bus clocked at 100

MHz. Although the current performance is sufficient for our application, using a DMA engine to unload the FIFOs storing data received from the FEMs and copy relevant data from the external memory to the Ethernet MAC is being considered.

5) DAQ PC and interface to the experiment wide DAQ

The DAQ PC aggregates the data received from all DCCs (~5 MB/s total) and forwards them to the global DAQ which is based on the MIDAS framework [8]. The I/O requirements are modest at that level for modern computers. The interface software between the DCCs and the MIDAS based DAQ is developed by collaborators.

V. System Deployment

At present, two thirds of the front-end electronics have been produced and one TPC is fully equipped as pictured in Fig. 7.

Fig. 7. One of the two end plates of a fully equipped TPC.

Extensive studies with cosmic rays, the calibration laser and a test beam are being made at Triumf. One of the first events seen is reproduced in Fig. 8.

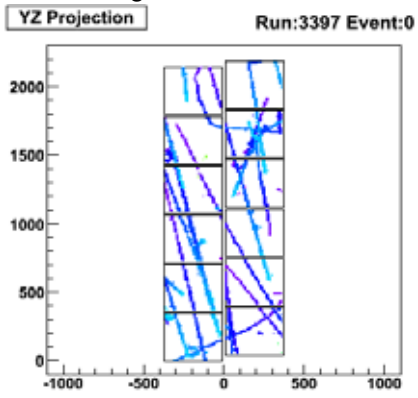

Fig. 8. A cosmic event with many tracks.

The two remaining TPCs will be equipped in summer 2009 and the complete system is expected to be ready for commissioning at Tokai by the end of the year.

VI. Conclusion

We presented the design of the readout electronics for the TPCs of the T2K experiment. We designed a dedicated ASIC for the readout of pixellated Micromegas detectors. This 72-channel chip is based on a SCA. We claim that buffering data in the analog domain provides a two fold decrease in power consumption and more than a four fold increase in channel density compared to an equivalent digital solution. However, our device is only suited to applications that tolerate the ~2 ms dead time needed to digitize the 511 time bins of the SCAs. The digital electronics mounted close to the detector comprises a memory capable of storing one event and the logic resources (FPGA) to apply pedestal equalization and zero-suppression. The global power consumption of the front-end electronics is 2 kW for 124.000 channels, i.e. 16 mW/channel. The readout of 1728 channels is multiplexed over a 2 Gbps optical fiber leading to a total of 72 fibers for the full system. Commercial Xilinx ML405 Virtex 4 evaluation cards with some customization are used as a cost effective temporary solution for the back-end electronics. Despite the limited I/O capability of the PowerPC processor embedded in the FPGA of the ML405 board, the target event acquisition rate of 20 Hz for event fragments of ~64 kB is comfortably reached.

References

[1] The Super-Kamiokande Collaboration, “The Super-Kamiokande Detector”, Nucl. Instrum. Meth. A501, 2003, pp. 418-462.

[2] Y. Kudenko, “The near neutrino detector for the T2K experiment”, in Proc. INSTR08, Novosibirsk, 5 March 2008. online: http://www.nd280.org

[3] P. Baron et al., “Large bulk-MicroMegas detectors for TPC applications in HEP”, in IEEE Nuclear Science Symposium Conference Record vol. 6, Oct. 26- Nov. 3 2007, pp. 4640 – 4644.

[4] L. Musa et al., “The ALICE TPC front end electronics”, IEEE Nuclear Science Symposium Conference Record, vol. 5, 19-25 Oct. 2003, pp. 3647 – 3651.

[5] P. Baron et al., “AFTER, an ASIC for the Readout of the Large T2K Time Projection Chambers”, IEEE TNS vol. 55, Issue 3, Part 3, June 2008, pp. 1744 – 1752.

[6] D. Calvet, “A New Interface Technique for the Acquisition of Multiple Multi-Channel High Speed ADCs”, IEEE TNS vol. 55, Issue 5, Part 2, Oct. 2008, pp. 2592 – 2597.

[7] “ML405 Evaluation Platform”, Xilinx User Guide UG210, March 21, 2007. online: http://www.xilinx.com

[8] S. Ritt and P. Amaudruz, "The MIDAS Data Acquisition System", in Proc. IEEE 10th Real Time Conference, 1997, pp. 309-312. Midas online: https://midas.psi.ch/