UC Irvine

UC Irvine Previously Published Works

Title

Off-detector electronics for a high-rate CSC detector

Permalink

<https://escholarship.org/uc/item/88c655m2>

Journal

IEEE Transactions on Nuclear Science, 51(3 I)

ISSN

0018-9499

Authors

Drego, N
Hawkins, D
Lankford, Aj
et al.

Publication Date

2004-06-01

DOI

10.1109/TNS.2004.828797

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Off-Detector Electronics for a High-Rate CSC Detector

N. Drego, D. Hawkins, A. J. Lankford, Y. Li, M. Medve, S. Pier, M. Schernau, and D. Stoker

Abstract—Data acquisition (DAQ) electronics are described for a system of high-rate cathode strip chambers (CSC) in the forward region of A Toroidal LHC Apparatus (ATLAS) muon spectrometer. The system provides serial streams of control signals for switched capacitor array analog memories on the chambers and accepts a total of nearly 294 Gbit/s in serial raw data streams from 64 chambers in the design configuration. Processing of the data is done in two stages, leading to an output bandwidth of 2.56 Gbit/s. The architecture of the system is described, as are some important signal processing algorithms and hardware implementation details. Although designed for a specific application, the architecture is sufficiently general to be used in other contexts.

Index Terms—Digital signal processors, particle tracking, real-time systems.

I. THE ATLAS CSC SYSTEM

A Toroidal LHC Apparatus (ATLAS) Cathode Strip Chamber (CSC) system is designed to measure high momentum muons in the high radiation environment of the forward regions. The CSC system consists of two endcaps, each containing 16 chambers in the initial configuration and 32 chambers in the design configuration. Each CSC chamber has four layers, providing precision measurements in the (radial) bend direction and coarser measurements of the transverse coordinate. Each chamber has 768 precision coordinate channels and 192 transverse coordinate channels [1], [2].

Due to severe radiation levels in the CSC environment, a minimum of the CSC electronics will be located on the detector in order to minimize development of custom rad-hard circuits. The total ionizing dose at the CSC location is approximately 3 kRad per year and single event upsets are expected from $1.52 \cdot 10^{11}$ hadrons per cm^2 per year with energies of more than 20 MeV. The on-detector electronics amplifies and shapes the cathode strip signals, and stores the pulse height information during the level 1 trigger latency. Upon receipt of a “level 1 trigger accept” (LVL1 Accept), four time samples are digitized and transmitted via high-speed fiber-optic G-Links to off-detector electronics. Many of the components of the on-detector electronics are also used by the liquid Argon calorimeter, where they have been radiation qualified. Tests for the remaining components are ongoing.

Manuscript received May 30, 2003; revised January 16, 2004. This work was supported in part by the U.S. Department of Energy by Grant DE-FG03-91ER40679, and the National Science Foundation by Grant PHY-9977770 and PHY97-22537.

The authors are with the University of California, Irvine, CA 92697 USA (e-mail: schernau@uci.edu).

Digital Object Identifier 10.1109/TNS.2004.828797

Sampling and digitization are performed on-detector but are controlled by the off-detector electronics, which consists of optical transition modules and readout drivers (RODs).

The ROD processes the received samples in two stages. The first stage, sparsification, suppresses hits below threshold and hits not associated with the current bunch crossing. The second stage, rejection, finds tracks and removes isolated neutron hits. The remaining data are sent to the ATLAS Trigger/DAQ System for further processing.

II. THE CSC ELECTRONICS SYSTEM

A. The On-Detector Electronics

The CSC on-detector electronics [3] resides on amplifier-storage module (ASM) boards. Each cathode strip is connected to a preamplifier and shaper circuit, which creates a bipolar pulse with a 140 ns shaping time to mitigate pile-up effects. The shaped pulses are sampled every 50 ns, and the pulse height information is stored in a switched capacitor array (SCA) with 144 cells for the duration of the level 1 trigger latency and digitization. Up to 7 blocks of 4 samples have to be stored during digitalization in addition to the storage cells that cover the latency period. At a sampling period of 25 ns, this would require a total of 160 cells. To accommodate the existing SCA, a sampling period of 50 ns was chosen. The SCA is only read out after a valid LVL1 Accept signal. During readout, four time samples for each channel are digitized by a commercial 12-bit analog-to-digital converter (ADC), multiplexed, and sent to the off-detector electronics.

The on-detector electronics for each CSC chamber consists of five ASM boards, each handling data collection for 192 strips. Four of these boards handle the precision strips, while the remaining board handles all transverse strips of a chamber. The digital data of each ASM board are transmitted via two fiber-optic G-Links (Agilent HDMP-1022) to the off-detector electronics. These G-Links will run at 40 Mwords/s with 16 bits/word. 16 Clock and control signals, such as the read and write addresses for the SCA, are sent to the ASM board from the off-detector electronics via an additional G-Link. This link will also run at 40 Mwords/s.

B. The Off-Detector Electronics

The off-detector electronics, shown in Fig. 1, consists primarily of Transition Modules and RODs housed in VME crates. Each ROD handles the data of two CSC chambers. The purpose of the ROD is to read out the detector, suppress data from

Fig. 1. Block diagram of the on-detector and off-detector electronics for the CSC. The Transition Modules and RODs are located off-detector.

empty channels, apply calibration constants, build event fragments, and send data to the ATLAS Trigger/DAQ System. The total digitized data collected from the 64 CSC chambers with 960 channels each is 294 Gbit/s at a trigger rate of 100 kHz when four time samples of 12 bit are read out. The RODs reduce this raw data stream by suppressing signals below a threshold cut and by rejecting signals with the wrong timing. The RODs also provide control signals for the SCAs and respond to commands from the readout crate controller (RCC).

The Transition Module plugs into the back of the ROD board. It contains optical receivers and transmitters, as well as an S-Link transmitter [4] for the readout link (ROL) to the Trigger/DAQ system. FPGAs route the data from the optical receivers to the P5 and P6 VME connectors. While the ROD was designed as a general readout and processing solution, the Transition Module is tailored toward the CSC detector.

C. The VME Configuration

The off-detector electronics is housed in two VME crates, one per endcap. A timing interface module (TIM) [5] in each VME Crate receives and delivers clock and control signals to and from

the corresponding modules. Each crate contains an RCC to control and coordinate data acquisition for the CSC system. The RCC is an off-the-shelf VME single-board computer. It communicates via Ethernet with the ATLAS Trigger/DAQ System and is responsible for starting and stopping data acquisition and for sampling events from the CSC detector for monitoring purposes.

III. ROD ARCHITECTURE

A. ROD Architecture

The DSP-based architecture of the ROD was born of the belief that software is a more manageable vehicle for solving complex problems than hardware. For some real-time problems, however, software solutions may be impractical due to cost, size, or power requirements. Texas Instruments vastly increased the range of real-time problems that can be solved in software when they introduced the TMS320C6000 family of DSPs.

Data from the detector arrives at the off-detector electronics on optical links at 40 MHz. The DSP-based architecture confines the resulting 25 ns processing imperative to the very front end of the ROD. Most of the ROD does not need to operate at

the link clock rate. The processor is free to perform some operations quickly, i.e., at its own clock rate, and others at the scale of microseconds. Monitoring and histogramming tasks can be started every millisecond. By eliminating imperatives induced by the data links, the DSP-based architecture brings the engineering challenges associated with the ROD into a realm where they are manageable. This ROD architecture can be adapted to a wide range of raw data stream formats because the bulk of the architecture is independent of the specific implementation of the data links.

The DSP-based architecture became practical when two commercial products were introduced at reasonable cost: the Texas Instruments TMS320C6000 family of DSPs and the Xilinx Spartan II FPGA. The DSP obviated the need for costly FPGAs for data processing, large FIFOs or SRAM chips for data buffering, and other hardware for histogramming and monitoring of detector performance. The Spartan II FPGA replaced many FIFO chips and transceiver chips.

Each ROD contains 10 sparsification processing units (SPUs), two rejection processing units (RPU) and one host processing unit (HPU). Identical hardware is used for all 13 processing units. The ROD processes the incoming data in two stages: in the first stage, sparsification, the SPUs remove hits below threshold and hits with the wrong timing (wrong crossing number); in the second stage, the RPUs reject unwanted hits from neutrons.

After sparsification, the SPU applies gain and offset corrections to the data, organizes spatially adjacent hits in each layer into clusters and determines their peaking time.

The output of each SPU is sent to the RPU via the data exchange (DX). The Data Exchange is a bus that connects all DSPs with the Readout Link. The RPU receives data from all four layers of the chamber. Muons traversing the chamber leave hits in the four layers, while neutrons tend to leave hits only in a single layer. The RPU organizes the clusters into tracks and rejects isolated clusters due to neutrons. The remaining data are transferred via the Data Exchange and Readout Link to the Readout Buffers in the ATLAS Trigger/DAQ system for level 2 trigger processing and for event building.

Each ROD contains one SCA Controller, implemented in an FPGA. It sends identical clock and control signals to the on-detector SCAs via G-link transmitters located on the Transition Module. The SCA Controller maintains lists of free and used cells in the SCA. These lists are necessary because readout of a single time sample for all channels takes approximately $2\ \mu\text{s}$, making continuous storage in the SCA impossible. Every 50 ns, the controller assigns a free cell to store the current voltage of the amplifier/shaper circuit. Used cells are reassigned to the list of free cells after the latency time of the level 1 trigger. For each trigger, four corresponding time samples are read out of the SCA and digitized on the ASM before being sent to the ROD.

The HPU manages the overall operation of the ROD and provides an interface to the RCC via VME. The HPU executes commands from the RCC, and the SPUs and RPUs execute commands from the HPU.

During normal data processing, the HPU commands each SPU and RPU to process data associated with a LVL1 accept. The HPU creates a header and a trailer for the current event

Fig. 2. Block diagram of a DSP module.

and starts a DMA sequence when both RPUs have finished processing their part of the event fragment. The DMA sequence transfers the header, the processed data, and the trailer onto the Data Exchange, which forward them to the Readout Link.

During datataking, the SPUs and RPUs accumulate histograms to monitor the CSC performance. The HPU has access to histograms stored in the memories of the other DSPs and makes them available to the RCC upon request. All DSPs maintain error counts, which are copied by the HPU into VME-readable memory.

Raw data entering the ROD from the Transition Module first reaches the Interconnect Subsystem. It consists of an array of FPGAs that route the data from the Transition Module to the SPUs. Alternatively, the Interconnect can route test data generated by one set of DSPs to another set of DSPs for diagnostic purposes.

The FPGAs in the Interconnect Subsystem of the ROD accept 192 data lines from the Transition Module. The data rate of each line is more than 80 Mbit/s, for an aggregate bandwidth of more than 1.92 Gbyte/s. The FPGAs of the Interconnect can be programmed to route data to subsets of the DSP modules. For example, for the ATLAS CSC, the Transition Module sends data to the ROD on 170 lines, and the Interconnect Subsystem routes the data to ten DSP modules, the SPUs.

The ROD was developed to be a general-purpose data acquisition and processing platform. The hardware (i.e., the board-level design) and the vast bulk of firmware and software are suited to a variety of applications.

B. The DSP Module

In both the SPU and the RPU processing units, the data is first buffered in an input buffer, then processed by the DSP, and finally written to an output buffer; see Fig. 2. The SPU and the RPU have in common the requirement to process large amounts of data in limited time, to respond to errors, and to accumulate histograms and monitoring information. Consequently, they share the same hardware and differ only in their firmware and software. This hardware, the DSP module or generic processing unit (GPU), is a small plug-in board containing a DSP of moderate cost (less than 100\$), off-chip memory, and two low-cost FPGAs (less than 30\$ each). The HPU also uses this hardware. The GPUs small size of 70 mm $\times$ 70 mm and low power dissipation of less than 3 W allows 13 GPUs to be housed on the 9 U $\times$ 400 mm motherboard.

We have selected the Texas Instruments TMS320C6203 DSP, which contains a large (512 k bytes) on-chip data RAM for input and output buffers, and runs at a clock rate of 300 MHz. This DSP has eight functional units that can work in parallel. Its instruction set contains bit manipulation instructions, which are ideally suited to interpreting the raw data streams. The DMA controller of the DSP can move data into or out of data memory with little or no impact on the CPU performance.

The FPGAs in the DSP module act as buffer/transceiver between the DSP and the motherboard busses. In the SPU, the Expansion Bus FPGA buffers the input stream, while the External Memory Interface (EMIF) FPGA buffers the output stream. In the RPU, however, the EMIF FPGA handles both the input and the output stream. The GPU is a fairly generic device whose functionality is entirely dependent on the environment in which it is immersed, the logic loaded into its FPGAs, and the software it is running.

IV. DIGITAL SIGNAL PROCESSING

A. SPU Processing

The sparsifier processes the incoming data in two steps. First, a fast thresholding algorithm is applied to all channels. Channels above the threshold, and their nearest neighbors, are flagged. A cut corresponding to a 75-ns timing window is also applied. The second stage of processing is performed only on these flagged channels. It consists of finding the channel with the largest signal within each contiguous group (cluster) of channels and calculating the peaking time of the signal. The four time samples retrieved from each strip provide pulse shape information. The three largest samples on the positive lobe of the bipolar waveform are used to define a parabola. The time of its maximum provides an estimate of the peaking time of the pulse. This peaking time can be determined to an accuracy of approximately 1 ns. Clusters with peaking times outside a timing window of 35 ns are rejected since they are not associated with the beam crossing of interest.

The time-critical processing described above is hand-coded in assembly language to take full advantage of the independent functional units of the DSP. Depending on whether the trigger arrived on or midway between the 50 ns sampling clock, the second or third time sample is the largest. The code performs a threshold cut on this time sample for all channels. The 75-ns timing window cut is obtained by requiring that the second or third sample is indeed larger than the first and the last time sample. In order to perform these comparisons quickly, both sides of the DSP each work in parallel on two channels at the same time. This means that four channels are processed every four clock cycles. The actual routine is 19 clock cycles long, but was designed to run in parallel with itself every four clock cycles.

Loading all the data required for this routine uses the entire memory bandwidth of the DSP, but by placing the channel-de-

pendent thresholds in a different memory block from the ADC samples, we can allow the input DMA complete access to memory on every fourth clock cycle. We must also make sure that each side of the DSP is loading from different memory banks in a given cycle, so one side starts on the first channel while the other starts on the last. As they work toward the middle, they are always loading the samples and thresholds from opposite memory banks. The code uses all eight functional units of the DSP in every clock cycle during the loop kernels.

B. RPU Processing

The goal of the RPU processing is to consolidate the output of five SPUs into one event fragment and to reject isolated hits from neutrons. Muons traverse the entire chamber and create hits in all four layers, while hits from neutrons are confined to a single layer. The RPU waits until all five parts of the event have been received from the SPUs. For each cluster it then searches for spatial overlaps with clusters in the other three layers. Clusters without overlaps are removed. Groups of overlapping clusters are saved as tracks. These tracks can be used to monitor the efficiency and resolution of the chamber. This code is written in C++ and takes advantage of the highest optimization level of the compiler.

The RPU formats the remaining data to comply with the standard ATLAS event format, while the HPU adds the necessary header and trailer for the event fragment. The RPU also flags errors due to loss of lock in the data links or buffer overflows.

V. CONCLUSION

The off-detector electronics of the ATLAS CSC system is described. Core algorithms for the SPUs and RPUs have been written and benchmarked. Data transfer from a prototype ASM to a prototype ROD has been successfully tested. Several DSP modules have been built and used in these tests. Other design entries are in development.

A full system integration test is scheduled for fall of 2003. Production of the off-detector electronics is planned to start at the end of 2003.

REFERENCES

- [1] *Atlas Muon Spectrometer Tech. Design Rep.*
- [2] N. Dreger *et al.*, "Off-detector electronics for a high-rate CSC detector," in *Proc. 6th Workshop on Electronics for LHC Experiments*, 2000.
- [3] J. Dailing *et al.*, "Performance and radiation tolerance of the ATLAS CSC on-chamber electronics," in *Proc. 6th Workshop on Electronics for LHC Experiments*, 2000.
- [4] The S-Link Interface Specification, O. Boyle, R. McLaren, and E. van der Bij. (1997, Mar.). [Online]. Available: <http://hsi.web.cern.ch/HIS/s-link/spec/s-link.pdf>
- [5] TTC Interface Module—TIM, J. Lane, M. Postranecky, M. Warren, J. Butterworth, and D. Hayes. (2003, July). [Online]. Available: <http://www.hep.ucl.ac.uk/atlas/sct/tim>