Design of the Front-End Driver card for CMS Silicon Microstrip Tracker Readout.

Baird S.A., Bell K.W., Coughlan J.A., Halsall R., Haynes W.J., Tomalin I.R.

CLRC Rutherford Appleton Laboratory, Oxfordshire, UK
j.coughlan@rl.ac.uk

Corrin E.

Imperial College, London, UK

Abstract

The CMS silicon microstrip tracker has approximately 10 million readout channels. The tracking readout system employs several hundred off-detector Front-End Driver (FED) cards to digitise, sparsify and buffer analogue data arriving via optical links from on-detector pipeline chips (APVs).

This paper describes the baseline design of the Front-End Driver card which is implemented with a 96 ADC channel (10 bits) 9U VME board. Under typical LHC operating conditions the total input data rate per FED after digitisation is 3 GBytes/s and must be substantially reduced. The required digital data processing is highly parallel and heavily pipelined and is carried out in several large FPGAs.

The employment of modern FPGA simulation tools in the design of a VHDL model of the FED is discussed.

I. INTRODUCTION

The CMS experiment [1] is due to begin operating at CERN's Large Hadron Collider facility (LHC) in 2005. A key component of the CMS detector is the tracker [2] [3] which is designed to provide robust particle tracking and detailed vertex reconstruction within a strong magnetic field in the high luminosity environment of the LHC. The tracker is implemented using silicon microstrips complemented by a pixel vertex detector with a separate readout system [4].

The silicon microstrip tracker readout system consists of approximately 10 million detector channels and, at expected track occupancies, will generate over 70% of the final data volume at CMS. The readout system is clocked at the LHC bunch crossing rate of 40 MHz and is designed to operate at sustained Level 1 trigger rates of up to 100 kHz. Silicon microstrips are read out using analogue electronics. A schematic of the system is shown in Figure 1.

Microstrip signals are amplified and stored in analogue pipeline memory chips (APV) located on the detector [5]. On receipt of a Level 1 trigger and after some elementary signal processing and multiplexing the

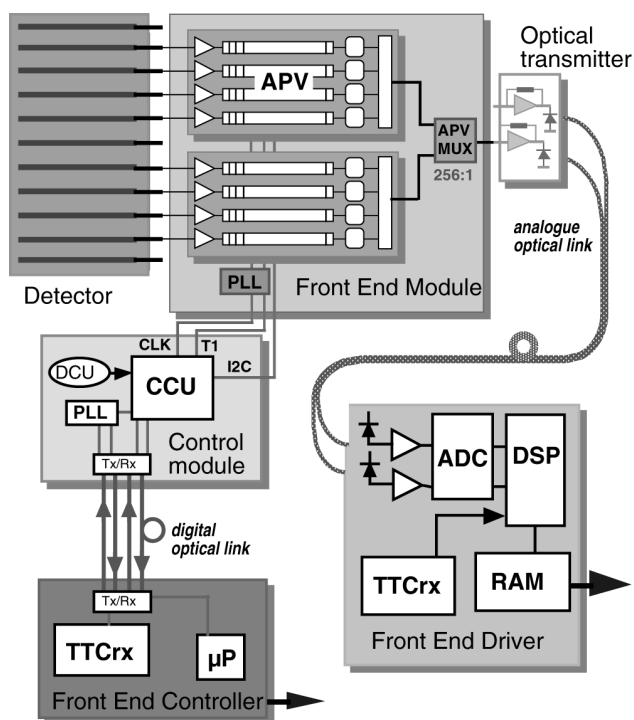signals are transferred to the counting room via analogue optical links [6].

Figure 1: The microstrip tracker readout and control system.

In the counting room the analogue optical data is converted back to electrical and digitised on Front End Driver (FED) cards. Each ADC channel processes data serially from a total of 256 microstrips (constituting an APV frame). The FEDs provide digital signal processing, including cluster finding, before storing the data in local memory buffers until required by the higher levels of the central data acquisition system. Each FED also receives information from the central Timing and Trigger Controls system (TTC) via a TTCrx ASIC [7].

The FED module has gone through two prototyping phases [8][9]. The principle functions of the FED are now clear and the project is now in the stage of finalising the design of the production version of the module.

II. FED ARCHITECTURE

A block diagram showing the proposed architecture of the FED is shown in Figure 2. The baseline FED design receives 96 analogue channels, each of which is digitised by a fast ADC and digitally processed before transmission to the next level of the CMS data acquisition system. The design comprises of several analogue and digital stages which are described in the following sections.
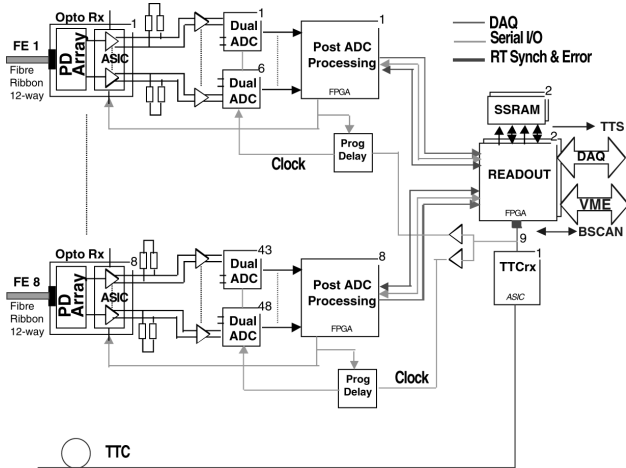

Figure 2: The FED architecture.

A. Analogue Stages

Each FED receives analogue data from the tracker from a cable containing 96 optical fibres which are grouped into 12-way ribbons. The optical signal is amplitude modulated and requires conversion to electrical levels before it can be digitised and processed by the FED. The opto-receiver packages consist of 12 channels of p-i-n diodes with the amplification stages provided by a custom ASIC. The parameters of the analogue optical link system are now fixed.

The ADC will be a 40 MHz, 10 bit commercial device. Experience of such components has been gained on prototype modules, but the choice of component is not yet decided. A number of dual ADC packages are under investigation [10].

The APV frame arriving at each ADC channel consists of a 40 MHz stream of alternate data from a pair of APVs (with data associated to the same trigger). The data output from each APV consists of a digital header, containing synchronisation and error status information, followed by the analogue data samples. It should be noted that if consecutive triggers arrive within a period of 7 μsecs there is no gap between the frames sent to the FED for this pair of triggers, thereby generating a so called back-to-back frame.

B. Digital Stages

The data from groups of 12 ADCs are fed to post ADC processing blocks (Figure 3). The data from each ADC are processed by their own independent pipelined processing logic clocked at 40 MHz.

The description of the processing stages, which follows, represents a baseline algorithm, which is intended to evolve into the final design. Each stage processes an APV frame as a continuous block after a fixed latency. The process can handle back-to-back frames with up to two frames in the pipeline simultaneously.
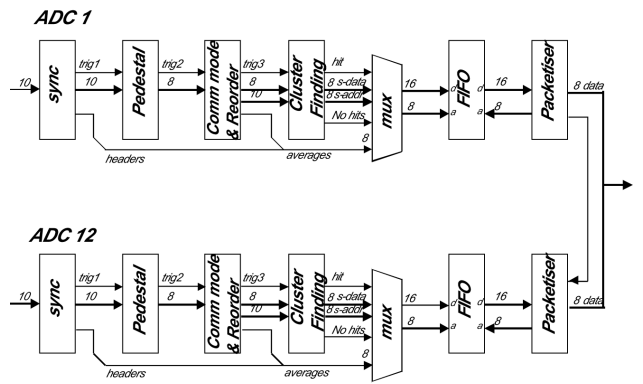

Figure 3: The post ADC processing block.

Synchronisation: The first stage recognises the start of an APV frame, producing a start frame signal in synchronisation with the first analogue sample, which triggers the subsequent processing stage. The circuit also extracts the pipeline addresses from the frame header and performs a nearest neighbour synchronisation check. The logic employs an auto-synchronisation method which locks directly to the APV data stream requiring no external trigger. This is necessary because there is no simple means of predicting the arrival of frames relative to the sending of the L1 trigger (except by state machine emulation) owing to the complex logic of the APV pipeline buffer.

Pedestal correction: Pedestal values are stored in a 256 location Look Up Table (LUT) and are subtracted from the data stream as they are clocked through. Values are sequenced from the LUT by an 8-bit counter which is triggered by the start signal from the previous processing stage. The stream of pedestal values from the LUT are fed into one half of a subtractor with the other input connected to the frame data stream. The output of the subtractor is a pedestal-corrected data frame which is passed onto the next stage.

Re-ordering: The APV outputs detector channels in a non-sequential, but fixed, order so that physically adjacent strips are not adjacent at the digitisation stage.

The cluster-finding logic requires the data to be in channel order and thus the APV frame must be re-ordered before the hit finding process can be performed.

The circuit uses a Dual Port Memory (DPM) to re-order the data. The input address is provided by a counter which counts with the APV output sequence. When an APV data frame is stored with this address sequence it is re-ordered in memory by channel order. The APV output sequence is emulated in real time by manipulating the output bits of a binary counter for addressing. After re-ordering, the data are immediately read from the other port using a binary counter for addressing. Both counters run at 40 MHz and thus, with the output sequence running immediately after the input sequence, the circuit appears externally to behave like a 256 cycle shift register. This delay is exploited by the common mode correction circuit.

Common mode correction: The common mode effect is assumed to be a common DC offset, which although expected to be the same for every channel within an APV, can be different for each APV and for every trigger. In the baseline algorithm the average of the data frame is taken as the common mode estimate.

The required level is calculated by summing each incoming APV stream in an independent registered accumulator. Since the APV streams are interleaved, each accumulator adds alternate samples which, after 256 cycles, produces two independent sums. These are divided by the number of non-faulty channels to produce the average value and stored.

Cluster-finding: The purpose of the cluster-finding circuit is to identify which channels are hit and to only write their associated data to the output buffer ready for readout by the FED-DAQ interface. The remainder of the data frame is discarded. Each hit value must be tagged with a detector channel address (8 bit). In the simplest case each hit is coded to two bytes which gives a data compression factor of around 60:1 for 1% occupancy. The sparsified data are clocked into local output buffers (FIFOs) ready for transfer to the Readout block.

The details of the algorithms to be implemented have not yet been finalised. Promising results have been obtained in early studies using simple algorithms based on a seed strip chosen on the basis of a suitable threshold and neighbouring strips satisfying a lower threshold. Alternative algorithms are currently under investigation in which the thresholds depend on the rms noise on each strip.

Each post ADC processing stage is designed so that it takes a fixed number of clock cycles. The total processing latency at 40 MHz is expected to be around 600 clock cycles and is dominated by the re-ordering and cluster-finding steps.

C. Readout block

The main function of the Readout block, shown in Figure 4, is to create the output data block for each event and transmit this to the DAQ system via the FED-DAQ link module. The data block is created by gathering the 96 individual data packets from the output buffers of the post ADC processing blocks. This data is then re-formatted and merged with the TTC and other header information.

The Readout block provides some dual-ported buffering with the aid of Synchronous Static Random Access Memory (SSRAM) units between the FED and the DAQ link to balance the flow of data between the two systems. However, the main buffering is located in Readout DPMs at the inputs to the event builder switch. Ancillary functions such as data integrity checking, exception monitoring and signalling to the Trigger Throttle and Synchronisation system [11] are also carried out in the Readout block.

There is a second (slower) data path via the VME backplane to the FED crate controller for local monitoring, and for control and downloading. This interface may also be programmed to perform as a logic analyser to spy on the internal functioning of the module.
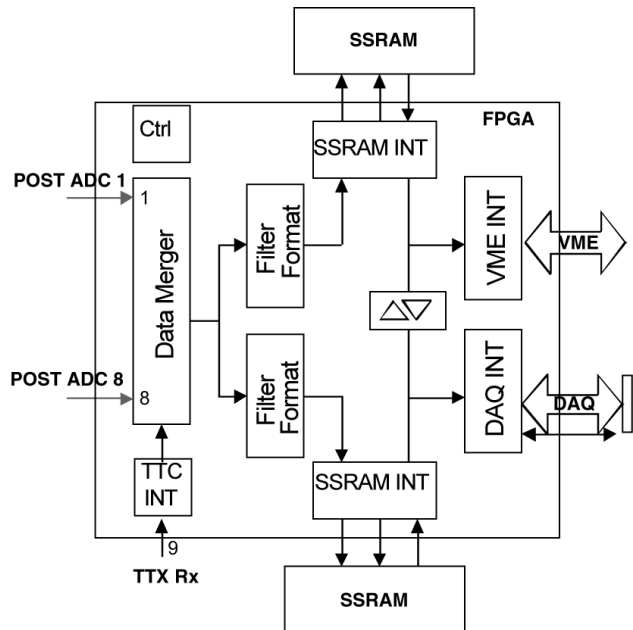

Figure 4: The Readout block.

III. FED IMPLEMENTATION

Figure 5 shows a proposed layout of the 96 ADC channel FED in a 9U by 400mm VME format. The opto-receiver packages are located at the front panel. Behind these are placed the dual ADC packages. The post ADC processing and Readout blocks are implemented in FPGAs. The employment of FPGAs (rather than custom ASICs) for the digital processing stages greatly increases the flexibility of the design. In particular, it may allow

the cluster finding algorithms to be further optimised once the tracker is in operation. The use of FPGAs has been made possible by the recent advances in performance to cost ratios of these devices. The DAQ link module is not shown, but may be implemented on a Common Mezzanine Card (CMC) plug-in.
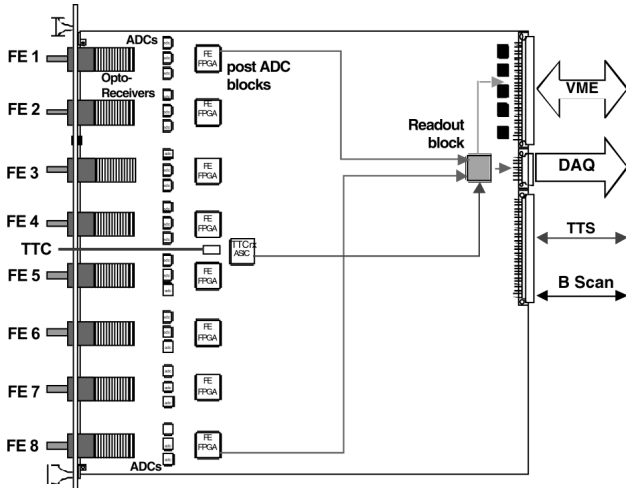

Figure 5: The FED 9U layout.

A summary of the tracker readout system parameters with this implementation is given in Table 1. The figures assume an optimal use of FED channels and a level 1 trigger rate of 100 kHz.

Table 1: Tracker readout system parameters

Total Number of channels	10,000,000
Total Number of APVs	78,000
Number of APVs per FED	192
Number of ADCs per FED	96
Number of FEDs	430
Number of crates	22
Single FED input rate	3.1 GBytes/sec
Single FED output rate	50 MBytes/sec/%occ + 10 MBytes/sec header

The modularity of the FED data inputs is matched to that of the 96-way fibre optic cables. The final modularity of the output FED-DAQ links will be determined by the average hit occupancy of the channels feeding a given group of FEDs. Recent Monte-Carlo studies show that the occupancy is expected to vary from 4% in the inner tracker layer to less than 1% in the outermost one [12]. In the outer regions of the tracker the merging of several FEDs will be necessary to make the most efficient use of the FED-DAQ links. As there is still some uncertainty in the expected occupancies and hence data rates the link system must be kept rather flexible. The FED-DAQ link must also be common to all CMS detectors and is now undergoing prototype testing as part of the CMS TRIDAS project [13].

IV. FED MODELLING

The digital logic implemented in the FPGA blocks will be programmed in VHDL. Modern design tools incorporate sophisticated simulation packages which allow the detailed behaviour of the model to be investigated.

A model of a prototype FED digital design (derived for an early version of APV) has recently been ported into one such design tools package [14]. This model has the basic elements of the post ADC processing required for cluster finding. Simulation studies are now under way which pass realistic Monte-Carlo generated tracker data as test vectors to the model. The required FIFO buffer depths for various cluster finding parameters and input data rates are under investigation.

The model is highly modular and will soon be extended to match the APV25 data frame characteristics. It will thereby evolve into the full post ADC processing block shown in Figure 3. In particular, it is intended to study the response of alternative clustering algorithms being developed in the offline environment. By synthesising these designs in candidate FPGAs an optimum clustering algorithm which balances hit finding efficiency and logic complexity will be arrived at.

It is also proposed to exploit software/hardware co-simulation tools in order to test and debug the low level software drivers required to control the system. By understanding the detailed behaviour of the design with simulations it is hoped to minimise the time spent on testing and debugging the final modules.

V. STATUS AND PLANS

The prototyping needs of the tracker community in laboratory and beam tests [15] have been shown to be well satisfied by the FED-PMC prototype. The evolution of the final FED can continue for longer than most parts of the system because it is in the external counting room and does not limit the design of the detector's internal elements.

The design of the final FED is expected to continue until mid 2001. The modelling and simulation studies of the internal digital processing blocks will continue in parallel with the finalising of the DAQ link interface. In particular, the decision on the optimum cluster finding algorithm is hoped for by the end of this year.

Market surveys of suitable ADC packages will begin soon. Following the selection of the ADC detailed work can commence on the final layout of the 9U board.

It is intended to have a prototype of the final FED for beam tests in 2002 when most other components of the tracker readout and controls system will be in or close to their final forms. Final FED production would then be expected to begin in 2003.

VI. REFERENCES

[1] The Compact Muon Solenoid, Technical Proposal, CERN/LHCC 94-38.

[2] The CMS Tracker Project, Technical Design Report, CERN/LHCC 98-6.

[3] Addendum to the CMS Tracker TDR, CERN/LHCC 2000-016.

[4] D. Kotlinski et. al., "The CMS Pixel Detector", these proceedings.

[5] G. Hall et. al., "The CMS Tracker APV25 0.25μm CMOS readout chip", these proceedings.

[6] F. Vasey et. al., "Project status of the CMS optical links", these proceedings.

[7] B. Taylor et. al., "LHC machine timing distribution for the experiments", these proceedings.

[8] R. Halsall et al., "Front End Readout Developments in the CMS Data Acquisition System" *Third Workshop on Electronics for LHC Experiments*, CERN/LHCC/97-60.

[9] J. Coughlan et. al., "A PMC based ADC card for CMS readout", *Fifth Workshop on Electronics for LHC Experiments*, CERN/LHCC/99-33.

[10] Signal Processing Technologies, Inc., *http://www.spt.com/*

[11] A. Racz et. al., "Trigger Throttling System for CMS DAQ", these proceedings.

[12] A. Caner, I.R. Tomalin, "On Balancing Data Flow from the Silicon Tracker", CMS Technical Note in preparation.

[13] CMS Trigger and Data Acquisition System, *http://cmsdoc.cern.ch/cms/TRIDAS/html/tridas.html*

[14] Mentor Graphics Inc., *http://www.mentor.com.*

[15] N. Marinelli et. al., "The CMS Tracker front-end and control electronics in an LHC like beam test", these proceedings.