Available online at www.sciencedirect.com
SCIENCE DIRECT
Nuclear Instruments and Methods in Physics Research A 560 (2006) 494–502
NUCLEAR INSTRUMENTS & METHODS IN PHYSICS RESEARCH Section A
www.elsevier.com/locate/nima

The LHCb DAQ interface board TELL1 ☆

G. Haefeli[a,*], A. Bay[a], A. Gong[b], H. Gong[b], M. Muecke[c], N. Neufeld[c], O. Schneider[a]

[a] *LPHE, Ecole Polytechnique Fédérale de Lausanne (EPFL), Switzerland*
[b] *Tsinghua University, Beijing, China*
[c] *CERN, Switzerland*

Received 23 November 2005; accepted 16 December 2005
Available online 23 January 2006

Abstract

We have developed an electronic board (TELL1) to interface the DAQ system of the LHCb experiment at CERN. 289 TELL1 boards are needed to read out the different subdetectors. Each board can handle either 64 analog or 24 digital optical links. The TELL1 mother board provides common mode correction, zero suppression, data formatting, and a large network interface buffer. To satisfy the different requirements we have adopted a flexible FPGA design and made use of mezzanine cards. Mezzanines are used for data input from digital optical and analog copper links as well as for the Gigabit Ethernet interface to DAQ. The LHCb timing and trigger control signals are transported by a dedicated optical link, while the board slow-control is provided by an embedded PC running a Linux kernel.

Keywords: LHCb; CERN; Electronics; DAQ; FPGA; Gigabit Ethernet; Silicon strip detectors; Common mode; Zero suppression

1. Introduction

The LHCb [1,2] experiment at CERN's Large Hadron Collider (LHC) is dedicated to the studies of CP violation and rare decays in the "beauty" sector (i.e. the set of particles containing a b quark). It is currently under construction and will as soon as the first LHC collisions are delivered in 2007 become operational. The physics of interest in LHCb is carried by particles emitted mostly at small angles with respect to the beam. This is why the detector (Fig. 1) covers only a limited angular region (pseudorapidity in the range 1.9–4.9). The proton beams will collide with $\sqrt{s} = 14\,\text{TeV}$ in the region of the Vertex Locator (VELO), shown on the left-hand side of Fig. 1. The produced particles will cross two Ring Imaging CHerenkov counters (RICH1 placed just after the VELO, and the RICH2 placed before the calorimeters), the tracking system consisting of a 4 Tm dipole magnet, and four stations implemented as solid state technology (ST) and straw chambers in the outer part (OT), a calorimetric system consisting of the PreShower (PS), Scintillating Pad Detector (SPD), ECAL and HCAL, and five muon chambers (MUON). To protect against radiation a shielding wall confines the detector from the surroundings. The different sub-detectors will be used to identify particles and measure their energies and directions of flight. In particular, the VELO can locate the decay vertex of a B meson with a precision of a few tens of micrometers. It is foreseen that LHCb will not run at full LHC luminosity but that the LHC beams will be defocused to keep the maximal luminosity below $10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$. To reduce the event rate on tape to about 2 kHz an efficient and flexible trigger is foreseen [3], based on custom electronics and on a cluster of about 2000 PCs. Three levels are foreseen: Level-0 (L0) [4], Level-1 (L1) [5] and the High Level Trigger (HLT). Flexibility is mandatory to adapt to unforeseen signal/background conditions.

A short description of the Trigger system follows. The first level of trigger, L0, reduces the rate from the LHC bunch crossing rate of 40.078 MHz down to 1 MHz. L0 has access to calorimetric and muon chamber data and exploits the fact that b-hadrons produce relatively large p_T particles

☆ This project is funded in part by the Swiss National Science Foundation (SNSF), and by the National Nature Science Foundation of China (NSFC) under contract No. 10225522.
*Corresponding author. Tel.: +41 21 693 04 97; fax: +41 21 693 04 77.
E-mail address: Guido.Haefeli@epfl.ch (G. Haefeli).
URL: http://lphe.epfl.ch/~ghaefeli.

0168-9002/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.nima.2005.12.212

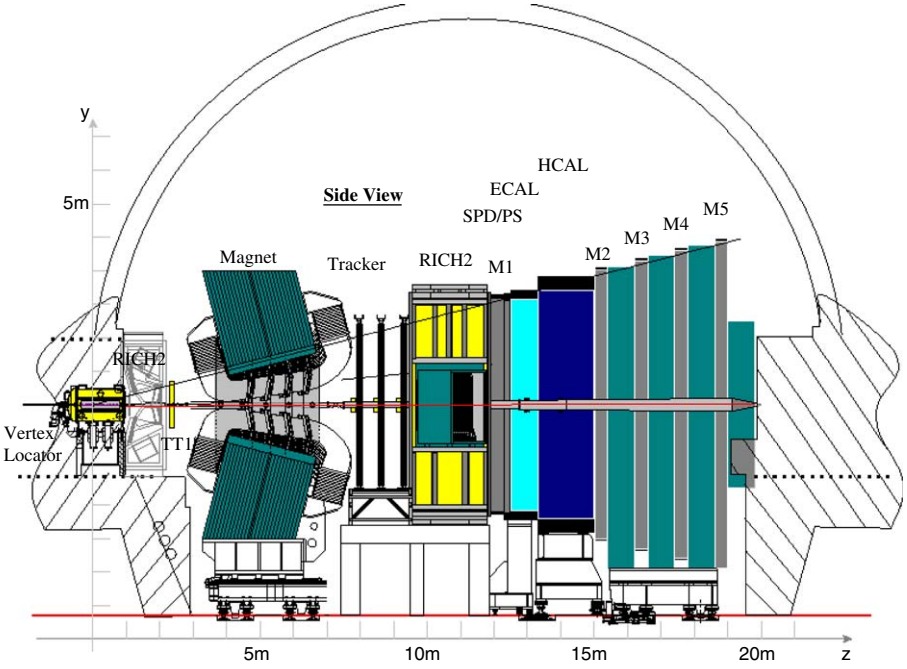

Fig. 1. The LHCb detector.

in their decays. The L0 system is fully synchronous with a fixed latency of 4 μs, and its implementation is based on custom designed electronics. The L1 algorithm running in the PC farm uses information from L0, VELO, and Tracker. A 3D track reconstruction is performed and events are selected based on large p_T tracks with displaced vertices (i.e. the signature of a b-hadron decay). In the end the HLT will have access to the full physics information and will be able to reconstruct a set of channels of interest reducing the event rate to about 2 kHz, corresponding already to $\sim 10^{15}$ byte of data storage per year.

In LHCb we call "L0 electronics" the front-end part of the DAQ which comes before the L0 decision. "L1 electronics" is the electronics handling L0 accepted data, before transmission to the DAQ. The implementation of each level is in general specific to each sub-detector. On the other hand, it has been recognized that all the sub-detectors require very similar L1 functionalities. Hence, it was decided to develop a common "L1 board" with enough flexibility to adjust to the different needs, with evident rationalization in the development effort, construction, and maintenance.

2. Overview of the data acquisition of LHCb

The general architecture of each level is the following: an input stage followed by a buffer defined by the trigger latency, an interface to receive the trigger decision, and an output stage with a buffer to "derandomize" the data transmission to the next trigger level. The derandomization is needed to adjust to the input bandwidth of the subsequent level.

Normally the L0 electronics is placed close to the detector and requires radiation-hard or radiation-tolerant components. To minimize the amount of electronics prone to the radiation effects, the L1 electronics is placed behind a shielding wall, quite far from the interaction point (more than 60 m). To preserve synchronization over large distances LHCb has a global clock and fast signal transmission network, the Timing and Fast Control (TFC) [6], implementing the Timing and Trigger Control (TTC) devices developed at CERN [7]. In particular, a radiation-hard chip (TTCrx [8]) is placed on all front-end electronics to receive the system clock, trigger decisions, resets and special commands from the TFC. In turn the TFC is driven by the Readout Supervisor (RS) [9], responsible for scheduling the distribution of the trigger decisions. Finally, the overall DAQ is slow-controlled by the Experiment Control System (ECS). The system overview is shown in Fig. 2. To cope with the radiation problem, the high density of detector channels, and the fast event rate L0 front-end ASICs have been developed. In particular, the Beetle front-end chip [10] will be used by VELO and Tracker. A fixed latency between the pp collision and the arrival of the L0 trigger decision of 4 μs has been chosen (i.e. 160 LHC bunch crossing, occurring each 25 ns). This latency accounts for the particle time of flight in the detector, transmission delays in the cables and in the front-end electronics, leaving about 2 μs for the actual trigger logic to derive the decision. The front-end is required to send out L0 selected events in 900 ns resulting in a maximal L0 accept trigger rate of 1.11 MHz. To avoid buffer overflows, the RS emulates the L0 derandomizer and "throttles" the L0 accept rate when its occupancy becomes close to the defined maximum value of 16.

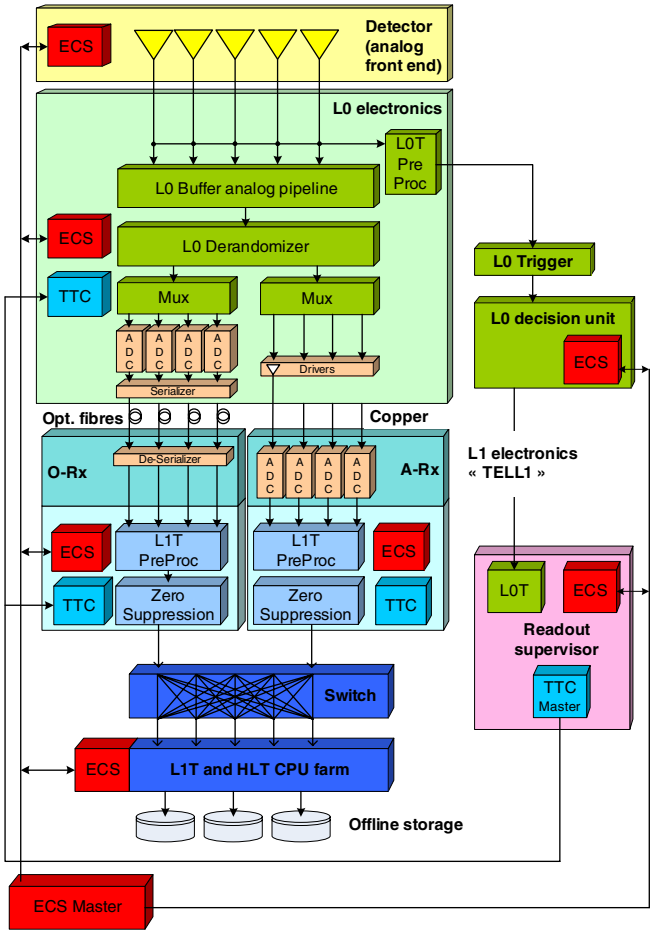

Fig. 2. LHCb trigger and data acquisition architecture overview.

After the L0 accept signal, the data are transferred by copper or optical lines to the counting houses. There the TELL1 boards perform deserialization or AD conversion, noise filtering and zero suppression, followed by transmission of the compressed data to the L1 and HLT trigger CPU farm. The TELL1 based L1 electronics architecture will be presented in Section 3. Section 4 will discuss how the TELL1 is synchronized with the rest of the LHCb electronic systems. Section 5 will present the data processing algorithms implemented. The description of the TELL1 hardware and its testing features will be given in Section 6.

3. TELL1 architecture

Fig. 3 shows a simplified block diagram of the TELL1 [11] data-flow. Physically, the TELL1 is composed of a motherboard hosting several mezzanine cards. Its main processing elements are five large FPGAs. A description of each element is given (see Table 1 for the FPGA resource usege).

A-Rx and O-Rx mezzanines: The TELL1 motherboard provides four connectors to accept the two kinds of receiver mezzanine boards developed so far: the analog

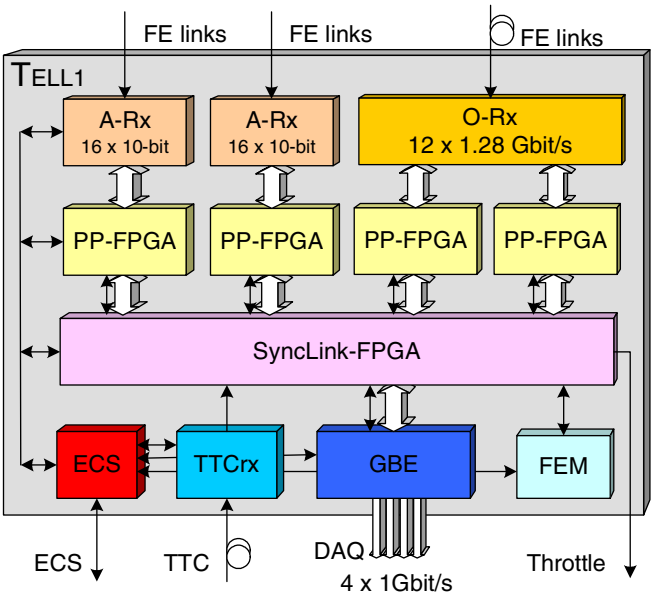

Fig. 3. Simplified block diagram of TELL1. The two options for optical and analog receiver cards are shown.

Table 1
The resource usage for the VELO implementation is summarized

Design part	Ligic elements (LE)	Memory bits
PP-FPGA		
Synchronization	5000	25 000
FIR	200	40 000
Pedestal	1000	16 000
CM suppression	2500	65 000
Derandomization	400	600 000
Zero suppression	2700	0
Linking	2000	68 000
ECS, test, divers	6600	442 000
Total used	20 400	1256 000
Total available	25 600	1 944 000
SyncLink-FPGA		
Event linker	2000	278 000
MEP assembly	500	67 000
MEP buffer (external)	500	0
Ethernet and IP framer	700	6000
ECS, test, divers	5100	632 000
Total used	8800	983 000
Total available	25 600	1 944 000

For the FIR filter lookup tables using only memory are used to implement the multipliers. In the CM suppression, a large number of embedded multipliers are used (not taken into account in this table). For buffering the MEP data at the SyncLink-FPGA an external SRAM (QDR) providing 2 Mbyte of memory is used.

version (A-Rx) for the reception and digitization of the VELO signals and an optical receiver (O-Rx) [12] for the other subsystems. Each one of the four A-Rx cards on the TELL1 provides 16 channels of 10-bit ADCs sampled at 40.078 MHz resulting in a total of 64 10-bit receivers for the completely equipped board. On the other hand, each O-Rx card uses two connectors and implements a 12-way

optical receiver resulting in a total of 24 channels per TELL1 board. The total user bandwidth per TELL1 is therefore $64 \times 10\,\text{bit} \times 40.078\,\text{MHz} = 25.6\,\text{Gbit/s}$ for the analog and $24 \times 1.28\,\text{Gbit/s} = 30.7\,\text{Gbit/s}$ for the optical version.

PP-FPGA: Each one of the four "PreProcessor" FPGAs receives from the Rx cards its share of L0 accepted data over a parallel link at a maximal event rate of 1.11 MHz. The processing is performed at full rate and consists of noise filtering, zero suppression and subsequent transfer to the SyncLink-FPGA.

SyncLink-FPGA: The "Synchronization and Link" FPGA is connected to the TTCrx and distributes synchronization signals as clock, trigger and event identification to the PP-FPGAs. Moreover, it is in charge of collecting and merging the data fragments from the PP-FPGAs, to assemble multiple events into so-called multi event packets (MEP), and to perform Ethernet and IP framing before transmission to the readout network.

Quad Gigabit Ethernet (GBE): The readout network is interfaced by a four-port Gigabit Ethernet card [13]. The network card is a custom design using the industry-standard SPI-3 interface delivering 4 Gbit/s of simultaneous Rx and Tx data rate. The IP address which has to be used to send the data is dynamically allocated and transmitted to the board by the TTC channel. In LHCb only the transmission part is used. Nevertheless, the receiver part is also implemented for connectivity tests. For this particular application, the received packets can be re-transmitted via packet mirroring implemented on the SyncLink-FPGA. A custom protocol on top of IP has been defined for this purpose [14].

TTC: The TELL1 receives the fast control signals via the on-board TTC receiver chip TTCrx. Fast control signals are the clock (locked at 40.079 MHz), the reset, the L0 trigger and the DAQ IP destination address.

FEM mezzanine: To minimize the number of required connections for the VELO readout, the control signal indicating ongoing data transmission (data valid signal) is not transmitted along the data. To reconstruct this information a Beetle chip is implemented on a small mezzanine card as a Front-end EMulator (FEM). The FEM's Beetle is connected with clock, reset and L0 trigger signals in a way to exactly reproduce the timing of the given data-valid signal.

Throttle system: After zero suppression, the length of each event is variable. Derandomizing buffers are employed to average the data rate and the data processing time. To prevent any overflows, each buffer can generate a throttle signal that is sent to the readout supervisor. The readout supervisor suspends the trigger signal until the buffers have recovered.

ECS: The on-board slow control is managed with a commercial credit card sized PC (CCPC) running a Linux kernel. A small adapter card, linked to the CCPC via PCI, provides all necessary interfaces to control and monitor all the other devices. The FPGAs and the quad Gigabit Medium Access Controller (GBE MAC) are on a shared microprocessor bus driven with a PCI bridge. I^2C and JTAG is also provided for convenient programming of the FPGA firmware EEPROM and monitoring the devices. The disk-less CCPC is attached over the network to a boot server.

External buffer: Until the final readout board was designed the L1 trigger included only the L0 information, VELO and part of the tracker data. Therefore, only a partial read out of the detector was performed at L0 rate while the complete detector data was stored in an on-board buffer waiting for the L1 decision. To store the data during L1 latency, a buffer based on DDR SDRAM of 384 Mbyte per board has been implemented and is still available for special DAQ modes.

4. Synchronization tasks

The overall fast control of the experiment makes use of the LHC-wide standard TTC system. Differences in particle flight time and in the timing distribution system will be compensated by an appropriate setting of the programmable delays in the TTCrx. However, the data transmitted synchronously from the detector cannot be expected to be synchronous anymore at the receiver side in the TELL1. For instance, in the case of the VELO the signal skew in the 60 m copper cable requires an adjustable sampling phase for each analog link on the A-Rx cards.

Event synchronization (event number tag) and source identification (data association to the source sensor) are solved in a specific way for each sub-detector. The simplest situation is given by Calorimeter data: this data have passed already some L0 trigger FPGAs based electronics and has been tagged with the complete event and data source identification. A similar situation exists for the Muon detector since they employ digital front end chips allowing to attach all necessary information to the data. For both Calorimeter and Muon simple cross-checks can be performed at the TELL1 reception: all the input channels must show the same event number and the source identification has to correspond to a specific value according to the cabling. If this is not the case an error flag is set. The VELO and the Tracker represent the most difficult cases. In order to reduce the number of lines to the strict minimum only a limited synchronization information is transmitted from the front-end electronics to the TELL1 boards. After L0 accept, the data are multiplexed by the Beetle chip and sent out over four analog links together with four pseudo-digital bits per link. The resulting 16-bit header carries an 8-bit error status plus an 8-bit Pipeline Column Number (PCN). The PCN defines the memory location at which the event was stored in the Beetle internal analog pipeline before read out. Neither event number nor source identification are attached to the data. With no redundant information the TELL1 relies on the correct order of event arrival and in the VELO case, on a reference

information from the FEM. The cross-check is performed by verifying that all the input channels have the same PCN.

To cope with possible Single Event Upsets (SEU) in the Beetle, the optical receivers must foresee automatic re-synchronization. The consequences of SEU are "bit-flips" in the data which can in some case be detected by parity bit check, checksum errors, or loss of synchronization. With some 4000 optical links running at 1.28 Gbit/s a total of 5.1×10^{12} bits are transmitted per second. With an assumed bit error rate of 10^{-15} bit/s an error rate of the order of 18 bit/h is expected. In order to reduce the effects of synchronization loss, the key idea is to use a RAM-based input circuit instead of a simple FIFO. Note that each event consists of a fixed number of words, for instance 35 in the case of the Tracker. The number of words written in the RAM is counted and upon reception of a complete event in one of the six parallel links per PP-FPGA, the state machines for the RAM address generation of all six links are triggered to proceed to the next event. This allows to continue with the next event despite incomplete reception on one of the links. Incomplete events or transmission errors are tagged by an error flag in the header.

5. On board data processing

The data processing required on the PP-FPGA is sub-detector dependent. The adaptation of the FPGA code for different users (sub-detectors) is simplified with a general processing framework that includes the structure for all processing steps shown in Fig. 4. The VHDL-framework is guiding each user to follow similar concepts in his algorithmics, which in turn helps to reuse a significant part of the existing VHDL code. The framework also supplies the users with a well-tested implementation of all interfaces and basic functionality to allow a basic usage of the board without any VHDL development.

The Outer Tracker detector and the Muon detectors are read out digitally. All analog processing such as pedestal correction or common mode correction are not needed and a simple algorithm for zero suppression can be applied.

For the Calorimeters, some data format conversion is foreseen which will replace the common mode suppression stage indicated in Fig. 4.

For the VELO and the ST, after the input synchronization the pedestal subtraction and Common Mode (CM) noise corrections is applied. For the VELO data processing an additional filter stage called Finite Impulse Response (FIR) filter is foreseen to compensate for signal distortion due to the 60 m analog transmission line. Several algorithms for the CM subtraction have been studied and tested by mixing test beam data (supplying a model for the noise) and Monte Carlo events (the signal). The Linear CM Suppression (LCMS) algorithm described hereafter has become the baseline algorithm. The subsequent steps in data processing are the zero suppression, data formatting and the assembly of multi-event packets prior to transmission to the network.

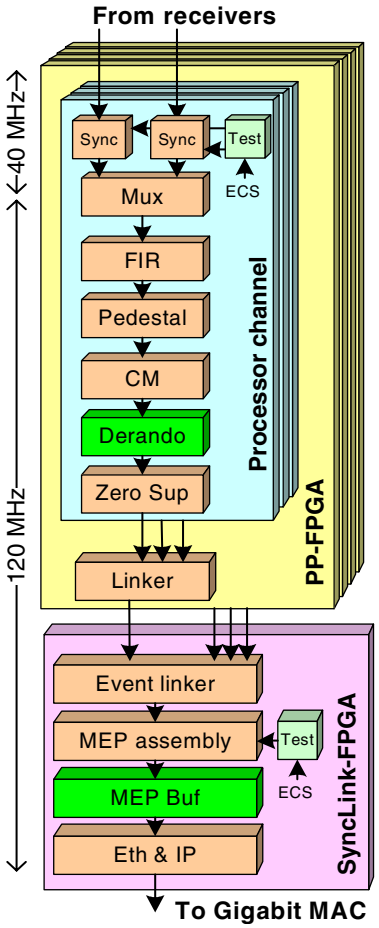

Fig. 4. Overview of the signal processing chain for the VELO and ST detectors.

5.1. Processing clock frequency

At the TELL1 input, the LHC frequency of 40.078 MHz is well below the maximum frequency achievable, the chosen FPGAs and memories could run at more than 120 MHz. The input stage is the natural place where to change the clock domain because a re-synchronization is needed anyway due to the different data arrival times from the different links. On the other hand, in order to fully load the four output ports of the GBE, a 120 MHz frequency is needed. To obtain a smooth data flow, the interconnection between PP-FPGAs and SyncLink-FPGA is also required to allow the same data transfer. Even if 80 MHz processing would be in principle sufficient for the PP-FPGA, in order to avoid multiple clock domains, we have chosen to maintain a 120 MHz clock. The increased frequency is used to process in time multiplexed mode, of the kind two-to-one.

5.2. Processing of the VELO analog signals

The signals from the Beetle chips are too weak to be sent over 60 m of copper wires. Amplification and frequency compensation are performed on "driver boards" placed

close to the detector. The lines are copper cables with shielded twisted pairs similar to category 6 Ethernet cable. On the TELL1 the A-Rx cards implement line receivers before the digitization with a 10-bit ADC. The choice of the sampling time is done individually for each ADC by phase adjustable clocks, which can be remotely programmed. In particular, this is needed to compensate for the time skew between different lines in the same cable which can be as large as 5 ns. Finally, we have foreseen 16 DAC channels on the mezzanine controlling the ADC gain for small gain corrections (about 5%). Fig. 5 shows the signal from a Beetle chip, measured at its destination at the ADC input. The internal Beetle pulse generator was switched on for two channels, with an amplitude corresponding to about 1.5 MIPS. It can be observed that the sampling point of the ADC can be chosen in quite a large region of 15 ns. Thanks to the frequency compensation, the channel cross-talk (inter symbol interference due to the limited bandwidth of the transmission line) is in general less than 5% of the pulse amplitude, while the effect after 50 ns is completely negligible. For an even more precise cross-talk correction a FIR filter of third order has been foreseen. The FIR acting on the 10-bit data is efficiently implemented with lookup tables and embedded or distributed multipliers. The FIR filter can also be used to adjust the gain for each analog link.

Each VELO physical channel (corresponding to a silicon strip) can have a different offset value due to the influence of electric fields, preamplifier tolerances, temperature variations, gain variations and so on. The pedestal is the mean value over some period of time of the amplitude registered by one strip with no physical signal. The pedestal can vary with time, therefore, we decided to implement an automatic procedure to continuously calculate the pedestal during data-taking. Assuming low detector occupancy it can be approximated by averaging over a large number of events. In the current implementation, a 20-bit pedestal sum per strip is implemented to calculate the sum over 1024 events, in such a way that the average is obtained by a simple 10 bit shift. For each event, the difference between the actual value and the pedestal value is computed for subsequent calculations. A very efficient way to follow the pedestal variation is described in Ref. [15].

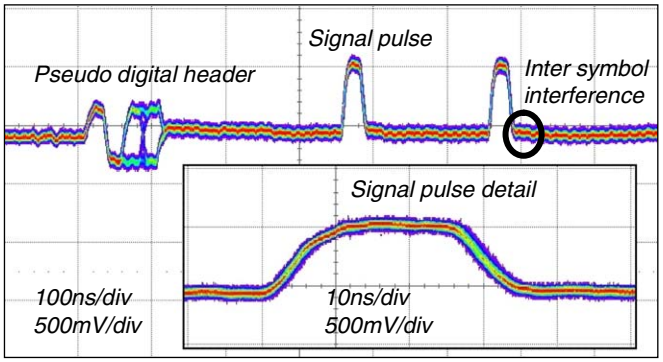

Fig. 5. Average of several analog signals generated by a Beetle chip, measured at input of the ADC, after 60 m of copper cable. The pseudo digital header can be seen on the left. The last two bits of the header correspond to the pipeline column number bits and are different for each event. Two channels were pulsed, a blow up of one of the pulses is shown on the bottom.

5.3. Linear Common Mode Suppression (LCMS) algorithm

The CM correction procedure (see Ref. [16] for details) is based on the assumption that the CM is varying in a time domain much larger than 25 ns. Thus, a linear approximation can be used over the data stream from a Beetle output channel (32 channels, transmitted in $32 \times 25\,\text{ns} = 800\,\text{ns}$). The LCMS consists of two identical iterations. During the first iteration, a least-square linear fit is performed over the 32 channels amplitudes and the result subtracted. Candidate "hit" channels are identified by comparison to a threshold and stored in a memory. The second iteration is applied to the results of the first with the candidate hits set to zero. The combination of the two iterations is subtracted from the original data which is then ready for the zero-suppression stage.

5.4. Zero-suppression and clusterization

Data compression is applied to the CM corrected data using detector specific hit finders and clusterization procedures. The concept for VELO and ST is very similar. In a first step, the seeding strips are found using a strip-individual Signal/Noise (S/N) thresholds and neighboring hits are combined into a cluster. The chosen data format allows to encode up to four strips in a cluster. Larger clusters are split. In a last step, the cluster position is calculated by a weighted average with 3-bit inter-strip precision.

5.5. Interface to DAQ

For this purpose, a quad port Gigabit Ethernet network card was developed suitable for the TELL1 specific needs [13]. Its interface to the SyncLink-FPGA is the FIFO-like industry standard SPI-3 [17] providing separate receiver and transmitter ports. The transmission performance is very close to the nominal bandwidth given by the 32-bit wide 120 MHz fast interface (3.84 Gbit/s) and is sufficient to load all four links simultaneously at 96%. Only one clock cycle per frame is used on the SPI bus for control purpose.

To obtain a high performance over the Ethernet-based network the packet size needs to be sufficiently large to reduce the protocol overhead given by the header of each packet. To obtain large packets from the small event fragments of about 150 byte each, multiple consecutive events are assembled into Multi Event Packets (MEPs). The number of events packed into a MEP is called the packing factor. For efficient network utilization, it is desirable to set the packing factor as high as reasonably possible. However, in order to avoid truncation due to the

limitation by the IP protocol (64 kbyte) it is necessary to limit it to a small number of the order of 10. The packing factor is a parameter of the readout system that can be changed according to the needs. Based on our studies the optimal packing factor is of about 10, and therefore the output packet rate reduced from 1 MHz to 100 kHz. It should also be noted that the gain of packing is practically lost once the resulting IP packet is very frequently larger than the maximum frame size allowed by the networking hardware which is 1500 byte for standard and 9000 byte for Jumbo frames.

Each node of the CPU farm represents a destination in the readout network. To cope with differing processing time per event due to varying event size or CPU speed, the destination for each MEP is assigned centrally over the TTC system by the readout supervisor. The RS uses the feedback from each CPU node to achieve dynamic load balancing.

5.6. Readout system

A "raw IP over Ethernet" protocol is used for the connection of the TELL1 boards to the switch. The protocol can be generated very efficiently inside the transmitting SyncLink-FPGA. It must be noticed that contrary to TCP, the protocol does not foresee any form of error recovery. To avoid packet loss in the system, sufficient connectivity is foreseen to keep the average link load below 80% and buffering is provided on the TELL1 as well as on the network switch. Data concentrator switches between TELL1 and the main switch are avoided since they have been identified as a source of packet loss, caused by insufficient data buffering capability. A large switch providing 1260 Gbit Ethernet ports will be employed for the complete readout. The number of links per TELL1 board going to the switch depends on the required bandwidth (see Table 2).

5.7. Test and monitoring

To verify the correct operation, test data can be injected at several stages of the signal processing. In the early development phase, the lack of input sources has shown the need for some form of data generator at the very first stage of the signal chain. Simulated event data can be stored in a data generator memory and injected by the L0 accept trigger command. Another data injector is inserted at the MEP assembly stage in the SyncLink-FPGA to support test of the DAQ network infrastructure.

Monitoring the processing via ECS is implemented at several levels. On the counting level all processing steps contain word, event and error counters giving in particular the possibility to verify the correct synchronization. At the data level, ECS access to event data buffers is provided allowing data integrity monitoring.

6. TELL1 hardware

The hardware implementation development of the first TELL1 version with the final number of channels, functionality, and choice of technology began in 2003. The first estimates of the required logic and memory were based on an Altera Apex [18]. At the same time, Altera announced and soon after released the first generation of Stratix devices with chip size, memory structure, and the implementation of embedded DSP blocks well fitting the TELL1 requirements. The later version Virtex II Pro from Xilinx [19] was discarded because of its cost, due to the integrated high-speed serializers not used in our application.

The possibility for migration to higher or lower density devices was implemented in the design. The PP-FPGA device chosen for the prototype was a Stratix with 20 k Logic Elements (LE) called 1S20, a 780-pin Fine pitch ball grid array (FBGA). This device can be replaced with all devices in the set from 10 to 40 k LE (1S10 to 1S40). For the SyncLink-FPGA a larger package was required since more IO is used. The 1S25 in 1020-pin FBGA package was chosen, which can also be migrated between 1S25 up to 1S80. In Table 1, a summary of the resource usage for the two FPGAs is given for each functionality.

For the final version, the 1S25 for the PP-FPGA was adopted (Fig. 6). The TELL1 board fits in a 9U VME crate with custom power backplane. Its power consumption of 50 W for the analog and of 40 W for the optical fully loaded board and its single slot dimensions allow to fit 21 boards in one crate. The data input connectors are located on the front side. On the back, there are the connectors for the four Gigabit Ethernet (RJ45), one 10/100 Ethernet (RJ45) for ECS, 1 optical fiber for TTC, and one connection to the throttle (RJ9), see Fig. 7.

For the LHCb experiment, a total of about 350 boards will be produced. Table 2 shows the number of TELL1 to be installed for each subsystem. It also gives the event size for each sub-detector and the number of Gigabit Ethernet

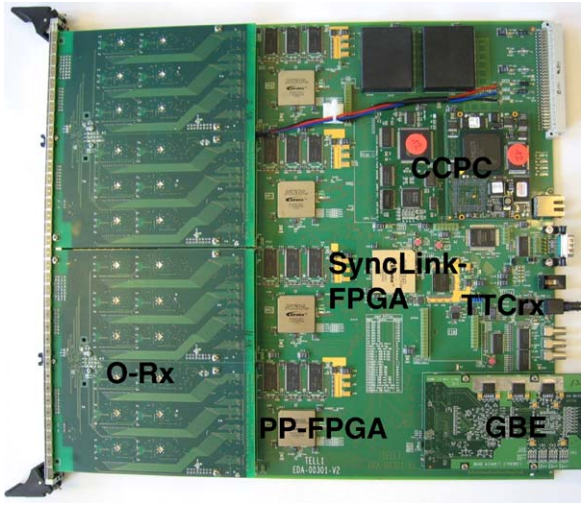

Fig. 6. Final version of the TELL1 readout board.

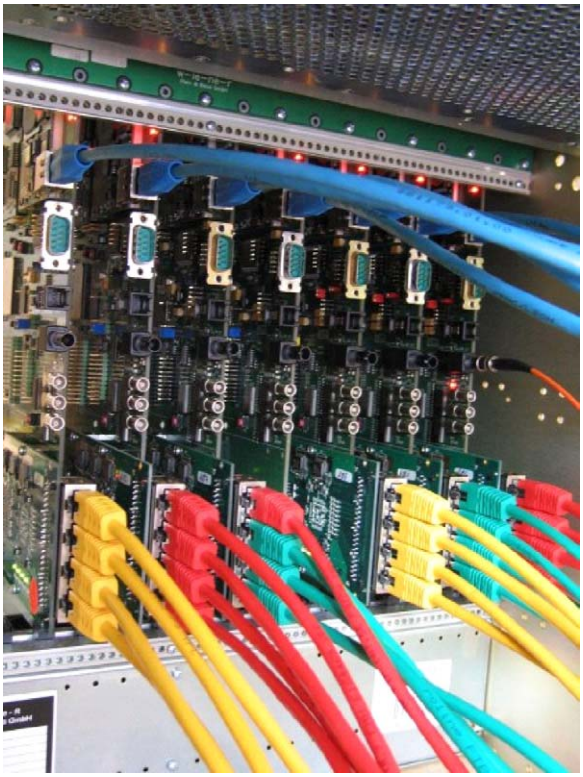

Fig. 7. TELL1 crate seen from the back. At the lower part of the picture one can see each network interface card connected to four cables; at the top of the visible part of the board is the ECS LAN and in the middle only one TTC fiber is plugged; the throttle connection is not cabled.

Table 2
The table gives the expected event size, number of TELL1 boards used, number of network links and maximal link load for each sub-detector

Subdetector	Event Size (byte)	Number of Tell1	Other L1 boards	Number of GBE links	Max. link load (%)
VELO	11 712	84		168	61
Trigger tracker	4650	48		68	53
Inner tracker	4173	42		54	61
Outer tracker	14 876	48		192	62
PS/SPD	1511	8		16	76
Ecal	7176	26		78	74
Hcal	1946	8		24	65
Muon	2239	12		30	65
Pileup system	561	4		8	56
RICH	5552		20	67	68
L0 pileup system	303	1		4	61
L0 cal	569	2		6	76
L0 decision unit	142	1		2	57
L0 muon	699	5		9	68
Readout supervisor	73		1	1	59
Totals	56 182	289	21	727	

The information for RICH and Readout Supervisor is also given to show a complete picture of the system.

connections that are foreseen to be used. As previously explained, the number of inputs and its sources are distributed such that the link load is below 80%.

6.1. Hardware testing

The board is equipped with JTAG boundary scan that covers a high percentage of the interconnection on the board. To verify the connectivity to the mezzanine cards, the JTAG scan is extended to the Gigabit Ethernet and also to a special input receiver mezzanine. This mezzanine is inserted in place of the O-Rx/A-Rx during testing. The boundary scan has been successfully used during pre-production testing and has proven to play a key role in the hardware verification.

In addition to the electrical test with the boundary scan, all buses and memories are systematically tested with varying bit patterns. This test can be repeated at any time and will be implemented as a part of the system self-test.

All optical links (4000 in the experiment) connected to the TELL1 can be verified with an integrated test mode. This special mode allows to measure the bit error rate at full link speed, for all links simultaneously. Moreover, an Optical Pattern Generator (OPG) module has been developed to provide a programmable data source to verify the correct functioning of the TELL1 including the receiver stage.

As described before, to test the network interface connectivity the SyncLink-FPGA can be placed in packet-mirror (loop-back) mode. From the trigger CPU farm test packets can be addressed to each link in the system and verified after reception.

7. Conclusion and outlook

We have developed the off-detector electronics (Level 1) for the LHCb data acquisition. The LHCb detector consists of several sub-detectors using digital optical links (or copper links for the vertex detector, VELO) to bridge the large distance between detector and the radiation safe counting house. The TELL1 board performs de-serialization for the optical links and digitization for the VELO, event synchronization, cable compensation for the analog transmission, pedestal following and subtraction, common mode suppression, zero suppression, data linking and interfacing to the Gigabit Ethernet readout network. The data processing is performed in four processor FPGAs accepting a total of 24 optical links running at 1.6 Gbit/s or 64 channels of 10-bit at 40 MHz for the analog version. Each processor FPGA provides a large amount of resources like 25 000 Logic Elements, 1.9 Mbit on chip memory and embedded multipliers. The data path to the readout network is provided by four Gigabit Ethernet links providing 4 Gbit/s output bandwidth.

The current TELL1 fulfills all LHCb requirements and will be produced as described in this note for the entire experiment. The board has been designed to be the most economical readout solution for the LHCb purpose. Two hundred eighty nine boards are needed (about 350 will be built, including spares).

The concept of readout with Gigabit Ethernet for the data path and 10/100 Ethernet for slow control has shown multiple advantages. No special crate bus is needed and the board can be efficiently remote controlled over the network. A second generation TELL1 with increased interconnection and output bandwidth using 10-Gbit Ethernet is under study. A complete upgrade of the board would also include a migration to the latest generation FPGA as, for example, the Altera Stratix II family.

Acknowledgments

We thank Jeremie Borel, Jorgen Christiansen, Dr Yuri Ermoline, Angel Guirao Elias, Raymond Frei, Jean-Phillippe Hertig, Guy Masson, Hans Mueller, Federica Legger, Laurent Locatelli, Erika Luethy and Prof. Beibei Shao for their help and advice.

References

[1] S. Amato, et al. (LHCb collab.), LHCb Technical Proposal: a large hadron collider beauty experiment for precision measurements of CP violation and rare decays, CERN/LHCC/1998-4, LHCC/P4, February 1998.
[2] LHCb reoptimized Detector Design and Perfomance Technical Desing Report, LHCb, CERN LHCC 2003-030.
[3] LHCb Collaboration, R. Antunes Nobrega, et al., Trigger Technical Design Report, CERN–LHCC/2003-31.
[4] J. Christiansen, Requirements to the L0 front-end electronics, LHCb Note 2001-014.
[5] J. Christiansen, Requirements to the L1 front-end electronics, LHCb Note 2003-078.
[6] R. Jacobsson, B. Jost, Timing and Fast Control, LHCb Note 2001-16.
[7] M. Ashton, et al., Status report on the RD12 project: timing, trigger and control systems for LHC detectors, CERN/LHCC 2000-002, LEB status report/RD12, 2000.
[8] J. Christiansen, A. Marchioro, P. Moreira, T. Toifl, TTCrx Reference Manual, CERN—EP/MIC, Geneva Switzerland ⟨http://ttc.web.cern.ch/TTC/TTCrx_manual3.10.pdf⟩.
[9] R. Jacobsson, B. Jost, Z. Guzik, Readout Supervisor Design Specifications, LHCb Note 2001-012.
[10] See URL: ⟨http://wwwasic.kip.uni-heidelberg.de/lhcb/⟩. Beetle 1.3 Reference Manual: ⟨http://wwwasic.kip.uni-heidelberg.de/lhcb/Publications/BeetleRefMan_v1_3.pdf⟩.
[11] G. Haefeli, et al., TELL1 Specification for a common read out board for LHCb, LHCb Note 2003-007.
[12] D. Wiedner, G. Haefeli, U. Uwer, A. Vollhardt, Prototype for an Optical 12 input Receiver Card for the LHCb TELL1 Board, LHCb Note 2003-137.
[13] H. Muller, A. Guirao, F. Bal, X. Tao, J. Christiansen, M. Muecke, N. Neufeld, G. Haefeli, Quad Gigabit Ethernet plug-in card, Cern Electronic Document Server ⟨https://edms.cern.ch/file/520885/2/gbe_manual_ver2.pdf⟩ EDMS 520885.
[14] B. Gaidios, M. Muecke, Ethernet test transport format for L1 boards, Cern Electronic Document Server ⟨https://edms.cern.ch/document/529434⟩ EDMS 529434.
[15] G. Haefeli, Contribution to the development of the acquisition electronics for the LHCb experiment, CERN-THESIS-2004-036.
[16] A. Bay, G. Haefeli, P. Koppenburg, LHCb VeLo Off Detector Electronics Preprocessor and Interface to the Level 1 Trigger, LHCb Note 2001-043.
[17] POS-PHY, SATURN Compatible Packet Over SONET Interface Specification for Physical Layer Devices (Level 3), PMC-980495, Issue 3, November 1998.
[18] ⟨www.altera.com⟩.
[19] ⟨www.xilinx.com⟩.