

Evaluation of an Optical Data Transfer System for the LHCb RICH Detectors.

N.Smale, M.Adinolfi, J.Bibby, G.Damerell, C.Newby, L.Somerville, N.Harnew, S.Topp-Jorgensen;
University of Oxford, UK
V.Gibson, S.Katvars, S.Wotton; University of Cambridge, UK
K.Wyllie; CERN; Switzerland

Abstract

This paper details further development of the front-end readout system for the LHCb Ring Imaging Cherenkov (RICH) Hybrid Photon Detector (HPD) which was reported in ‘Development of an Optical Front-end Readout System for the LHCb RICH detectors’ [1]. The HPD is a silicon pixel detector with an encapsulated binary readout ASIC*. The proof-of-principle single HPD readout chain is presented, with particular attention given to the data packing, transmission error detection and TTCrx synchronisation from the radiation-tolerant Level_0 electronics to the off-detector Level_1 electronics. The Level_0 transmits 17x36x2 bits per HPD at a sustained Level_0 accept trigger rate of 1MHz over multimode fibre within 900 ns; a non-zero-suppressed data block, plus address ID, error codes and parity checks. . FPGA interface chips, GOL ASICs, QDR memories, multimode fibre and VCSEL devices are used in the transmission of data with a 17 bit wide Glink protocol.

I. Introduction

The CERN/DEP-developed Hybrid Photon pixel detector (HPD)[2] has active elements which comprise a photo-cathode, electrostatic imaging system, encapsulated 32x32 pixellated silicon detector and binary readout ASIC [3][†]. LHCb has two RICH detectors, with approximately 450k 2.5x2.5mm² HPD pixels[‡] in total. These are read out at the LHC beam crossing frequency, 40 MHz. The ASIC, referred to as the Pixel chip, stores one bit per pixel in an 160 deep digital pipeline until a Level_0 trigger is received. The data are either discarded on a Level_0 reject decision or accepted, then multiplexed at 1024:32, and driven off chip in 800ns at an average 1MHz Level_0 accept rate. Figure I-1 shows the RICH modular conception, which incorporates the HPDs, which is optimised with mechanical, electronic, and module replacement constraints in mind. For the two RICHes a total of 220 on-detector and 54 off-detector modules will be used. The following is a brief overview of the component parts; for further detail the reader is referred to reference [1].

* Application Specific Integrated Circuit.

[†] The HPD was specifically designed for use within LHC.

[‡] The pixel size being determined by the spatial measuring resolution needed.

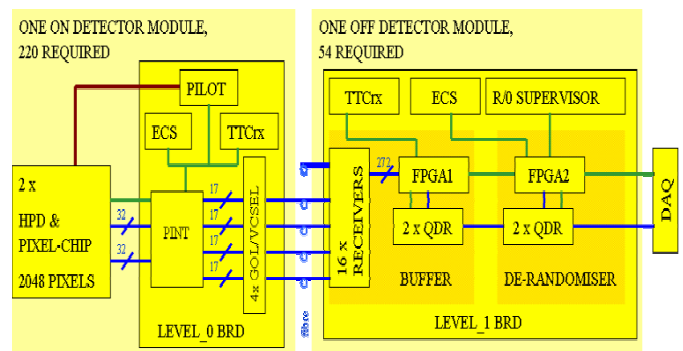


Figure I-1 Block diagram of modular system for HPD readout.

Referring to Figure I-1, the PInt logic splits the incoming data block from an HPD into two 16-bit wide blocks so as to be compatible with the Gigabit Optical Link (GOL)[4] when operating in an 800Mbs mode. Each 16-bit wide data block has headers and trailers added, comprising Bunch Count ID (BCID), error information, column parity and Hamming code; and a 17th bit for row parity. This brings the total transmission from PInt to GOLs to 17x36x2 bits/HPD/Level_0 trigger accept in 900ns. This block is referred to as the event block. The antifuse ACTEL AX1000 FPGA device will be used to implement the PInt logic and will be the centre of synchronisation and communication for the CERN developed TTCrx, Pilot and GOL chips and will communicate to the Experiment Control System (ECS)[5] using the JTAG protocol. The Timing Trigger Control ASIC (TTCrx)[6] provides clock, trigger, reset and fast control distribution. The Pilot ASIC is used for the biasing of the Pixel chip. The GOL serialises data from the PInt and drives Vertical Cavity Surface Emitting Lasers (VCSEL)[7] for optical transmission over 100m of multimode fibre to the Level_1 board. All components of the on detector module should be tolerant to the LHC radiation environment, see section II. The Level_1 module, designed using COTS[§] components, has two regions, namely the buffer and de-randomiser. The buffer region is able to receive 16 fibres (four Level_0 boards). Under the control of FPGA1, the event blocks are checked for correct beam count ID, parity and Hamming code. Any subsequent errors found in the event block or in the FPGA1 built-in-logic-observer (BILBO) are appended to the side of the event as an extra column; the ECS is also notified of an error occurrence via FPGA2. The event block, now 18x36 bits, is

[§] Commercial Off The Shelf.

stored into a Quad Data Rate SRAM block for the duration of the Level_1 trigger latency**. The QDR SRAM is a four-burst device and requires only one write address to be generated to store four consecutive 18-bit words. Data can be read in and read out on the same clock edge at a rate of 333Mbits/sec. The ECS interface and TTCrx are used in the same way as for the Level_0 board but further demands are made of the TTCrx with the full use of the channel-B to receive Level_1 triggers and user commands. On a Level_1 trigger accept, FPGA2 (de-randomiser area) receives data from the QDR buffers in a double data rate mode at 80MHz, the event blocks are put into the de-randomising QDR SRAMs for temporary storage so as to meet the required readout speed of the DAQ. FPGA2 can perform different functions on the event block such as zero-suppression, further error checking and some local error correction if requested. All communication between ECS and the readout supervisor is achieved with FPGA2.

II. Environment and Radiation Tolerance

The Level_0 electronics will be situated in the ~30 Gauss magnetic fringe field of the spectrometer magnet and will experience radiation doses of 3krad/year [8]. The HPD's are sensitive to magnetic fields of a few Gauss. A shell of ferromagnetic material has therefore been proposed to shield the Level_0 electronics, reducing the fields to less than 10 Gauss [9]. Except for the Pint, all Level_0 electronics will be fabricated in the 0.25µm process; the PInt is an antifuse FPGA from ACTEL. The analogue part of the Level_0 electronics suffers from charge trapping in thick oxides when irradiated which causes the threshold voltage (VT) of FET




Figure II-1 a) Standard FET layout, b) ELT FET

devices to shift, and for parasitic paths to be created that circumvent the normal conduction channel. Figure II-1a shows a standard FET layout; trapped charges can cause either a conduction path between drain and source in the surrounding thick oxide or effectively turn on/off the device by charge build up under the gate region. When ASICs are fabricated in a 0.25µm process the oxide under the gate region

** How many QDRs are used depends on the Level_1 latency and the QDR memory depth available.

is so thin, 6.2 nm, that charge is able to tunnel out, therefore inherently removing the problem of shifting in VT. To remove the parasitic paths between drain and source††, enclosed layout techniques (ELT) are used, see Figure II-2b, ensuring that the only available path is under the gate region. However a decrease in gate oxide thickness requires a decrease in supply voltage (and a decrease in VT) to prevent oxide punch-through due to the electric fields; the scaling factor used for 0.25µm technologies results in a 2.5V supply. The reduction in gate area and VT causes sensitive nodes for digital logic. A change in bit state or transient can occur by an ionising particle depositing enough energy in the gate region of a sensitive logic node to cause a single event upset (SEU). To minimise the effects of SEU, two forms of self-correcting digital control logic cells have been used.

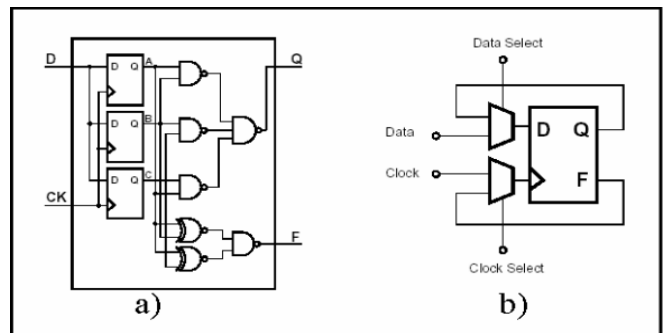


Figure II-2 a) Triple redundant flip-flop, b) Self correcting cell. [10]

Figure II-2a depicts a control logic flip-flop cell that would be used in, for example, a state machine. In this case the input is fed to three standard flip-flops and a majority-voting scheme is used to determine the output state. There is also a flag to show that a SEU has occurred. Figure II-2b shows a cell that would be used as part of configuration registers i.e. a register that is loaded from the JTAG interface. Internally it uses a triple redundant cell. In this case an analogue switch is used to select between feedback or configuration data. When in normal running mode (feedback) a SEU causes the flag to set, which is used to clock the majority vote back into the register therefore making the correction, normally achievable in less than 1ns, depending on technology. The Level_1 electronics are situated in the counting room about 100m away from the Level_0 area in a non-radiation and non-magnetic field region. The counting room can be considered as an electronic friendly environment and therefore standard COTS components can be used. This has the advantage of availability, maintenance, and cost effectiveness, with a broad range of products, and allows the use of SCRAM FPGA devices. Error checking, error correction and self test algorithms will be built into the Level_1 electronics to ensure that synchronisation is not lost and corrupt data are not being transmitted either to or from the Level_1 region.

†† The surrounding oxide is still relatively thick.

III. Transmission Format and Errors

Transmission errors can be a missing clock, an inversion of a data bit or a loss of event block synchronisation. These errors can be detected and a correction mechanism used for repair. The detect and repair method used will depend on: the number of errors that need to be corrected per second, the time taken to detect and repair, and the overhead in logic density to implement such a module. One method is to perform a reset of all the electronics, but this should be kept to a minimum of occurrences, as several hundreds of thousands of events will be lost during the reset period. Another option, which is used in this case, is to make full use of the allowable readout time from Level_0 to Level_1 of 900ns/event [11]. The data block from the Pixel chip is 32 words deep and is read out in 800ns. This allows four extra words that contain error-checking information to be added to the data block within the specified maximum Level_0 to Level_1 data transmission time of 900ns[11]. These words can be used to implement error detection and correction codes that can then be used downstream; it is more prudent to correct downstream as the number of event blocks that need correcting has been reduced by a factor of one thousand due to the Level_0 and Level_1 triggers, and it greatly simplifies the Level_0 logic algorithms. A further extension can be made to the event block when using GLINK in the 17th flag bit mode as the flag-bit is given up to the user between the first and last word.

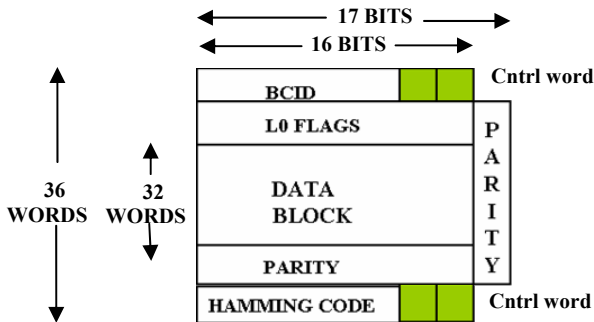


Figure III-1 Data format of an event block.

Figure III-1 depicts the event block for one fibre after being formatted by the PInt chip. The PInt splits the data block from the Pixel chip in to two blocks of 16x32 bits/HPD. Four extra words and a parity column are appended to both data blocks making them independent of each other with regards to transmission. The various parts of the event block are discussed in the following subsections.

A. G-Link Protocol and Control Words

The GOL transmitting and HDMP-1034 receiving devices are used in the 17-bit G-LINK protocol mode, the 17th bit being a flag. In hardware this gives: various configuration pins, 16 user data pins, a flag pin, and two pins for marking words as either control, data or idle words. Idle words are used in

periods when no event blocks are sent so that the transmitting and receiving devices can remain locked to each other; the clock for the receiver is recovered from transmission data. The first and last words of an event block are restricted to being 14 bits wide and marked as control words. The flag bit is also not available at this time. Ideally the control words should have a data-word that is predictable and recognisable by the Level_1 logic for synchronicity checks. The BCID from the TTCrx^{††} is common to both Level_0 and Level_1 regions and therefore makes a good option. By checking that the BCID is a control word and is correct, then a misalignment of event block data, a loss of clock or a wrong event block sent, can all be detected and action taken^{§§}. On receiving the BCID the Level_1 logic can increment a counter to check that no event words are missing between the start and end control words. The end control word is utilised as a Hamming code see subsection C. The remaining part of the event block is marked as data words and are 16-bits wide plus a flag-bit. The flag bits have been used for data block row parity see subsection B.

B. Parity Checks

Taking the 'XOR' of a row or column generates parity flags. By comparing row against parity flags an error grid can be generated allowing the correction of single errors. Double errors are identifiable but not necessarily correctable; this depends on the error rate within a block. Triple errors are not recognisable or correctable. The parity generation for columns requires storage of a running parity over 32 clock cycles, which can also suffer from SEUs. An error in the parity flags results in false error detection, which is impossible to distinguish or correct. Studies have shown that a Hamming code is more robust in terms of reducing false errors while not so good at being able to correct real errors, while the opposite is true for parity checking. The use of both Hamming and Parity brings the predicted bit error rate to less than 10^{-12} with 3.2×10^{-5} being correctable; the false errors are less than 5.5×10^{-18} [12]. The PInt and FPGA1 of Level_1 both use exactly the same parity algorithms to ensure compatibility.

C. Hamming Code

The Hamming code uses a given polynomial of the N^{th} power, $N+1$ bits, by which the data stream is modulo-2 divided. There is no carry operation between places, i.e. each place is computed separately, which makes the division simpler. The quotient is disregarded and the remainder of N bits is appended to the data stream. When the data with the appended remainder is divided by the same polynomial, the remainder should be zero if no errors occurred [12]. The Hamming code can be applied to the full event block, as it is the last word to be appended. The most suitable polynomial for the Level_0 to Level_1 transmission is $x^{11} + x^{10} + x^4 + x^3 + x + 1$, (11000011011), which gives an 11-bit remainder[12]. This will detect single errors, double errors or

^{††} The BCID is given on a Level_0 accept trigger.

^{§§} These errors would result in a flag to the ECS for a reset.

triple errors. Downstream, the single errors can be corrected, the double errors are recognisable and the triple errors look like single errors. Hardware for Hamming coding is simple and compact. For full checking of an event block an 11-bit shift register and five 'OR' gates are required. The PInt and FPGA1 of Level_1 both use exactly the same Hamming code algorithms to ensure compatibility.

IV. Proof of Principle

The proof of principle requires one HPD to be operating at a 40MHz rate and to be read out to a DAQ system over fibre optics. The hardware*** should emulate, and if possible use components that will be found in the final implementation. This proof of principle differs in five aspects from the proposed final solution: The Level_0 board uses a non radiation tolerant Xilinx FPGA as it offers multiple re-programming capabilities, whereas the ACTEL anti-fuse is expensive and is programmable only once. The physical hardware is considerably bigger for testing purposes. The Level_1 receives only two fibres and the FPGA multiplexes data into the QDR at a double data rate of 40MHz instead of 80MHz; this allows easy use of a four-burst QDR device when connected to only one HPD. The Pilot DAC ASIC chip

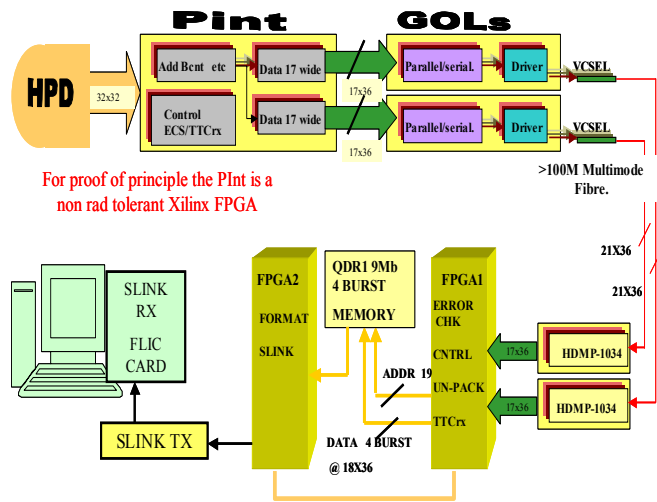


Figure IV-1 Proof of principle readout scheme.

used for external biasing of the Pixel chip is not yet available, so COT DACs have been used. The Level_1 to DAQ-PC uses a readily available and simple CERN developed optical Gbit Link (ODIN) which is based on an S-Link specification, also a CERN development[13]. The ODIN comprises an optical link capable of a maximum data rate of ~1.2Gbit/s and transmission and receiver board. Figure IV-1 shows the proof of principle readout scheme.

A. Hardware

Figure IV-2 shows the proof of principle Level_0 board. The board size is 28x14cm² and along the front edge is the pin

connector for the HPD. The TTCrx plugs into the back of the board, not shown here. Group a) is the power supply, b) the GOLs and VSCELS, c) the DACs for Pixel biasing and d) is the Xilinx FPGA.




Figure IV-2 The proof of principle Level_0 board.

Figure IV-3a shows the proof of principle Level_1 board. The board size is 23x17cm². The board can connect into a VME crate for power or use bench supplies. The FPGAs are close to the QDR memories to optimise speed performance. The interface boards for: Level_0 fibre optic receiver, four fibres in total, TTCrx and S-Link card, all plug onto the face of the board see Figure IV-3b.

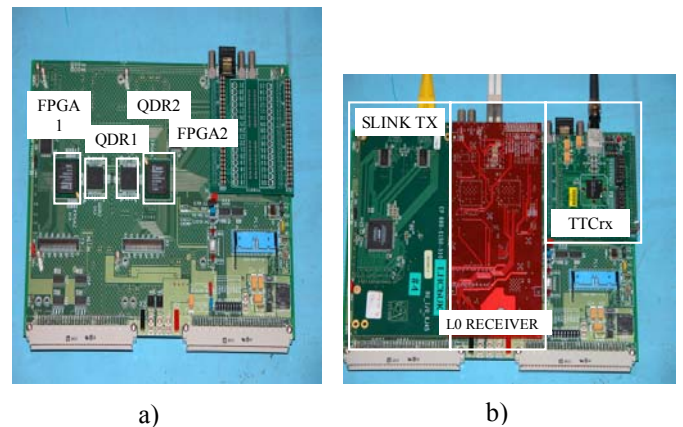


Figure IV-3 Level_1 board a) not loaded with interface cards, b) loaded.

B. FPGA and QDR memory

Xilinx Spartan II XC2S200 FG456 FPGAs are used as controller devices. The device offers 200,000 system gates, >5000 logic cells. 284 I/Os and 16 selectable I/O standards with access times of 200MHz, internal clock speeds of 333MHz, and is a low cost item, around £40 sterling. Internal Delay Lock Loops (DLL) are used for clock multiplication. The Spartan-II is programmable directly by JTAG or

*** And chip algorithms in the case of FPGAs.

EEPROM. The algorithms have been written in VHDL^{†††}. For the present testing of the readout chain FPGA2 is used only to interface to the S-Link as follows: accept event blocks from the QDR memory @ 1.44 Gbits/S; this is two event blocks at double data rate. Strip the 17th bit (parity) and 18th bit (error code) and format them into end words. Concatenates the two 16 bit data words into a 32-bit word. Generate all the necessary control and clocks for the S-Link. Clock out the data to the S-Link at 20 MHz.

The QDR SRAMs are samples from Micron. The MT54V512H18EF is internally organised as a memory bank of 9Mbits (18x512) and can store up to about 7000 events from an HPD pending the Level_1 trigger decision. Data can be read in and read out on the same clock edge at a rate of 333Mbits/sec. The QDR memory is a four-burst device and requires only one write address to be generated to store four 18-bit words. The package is ball grid array with a 13x15mm² 1mm pitch.

C. Delay Lock Loop (DLL)

The Level_1 FPGA1 takes advantage of the Xilinx internal DLLs to multiply the TTCrx 40.08MHz clock by a factor of 2 and 4 and adjusts the phase of the 80MHz QDR clocks by 90 degrees. To ensure TTCrx/DLL compatibility, a study was undertaken to investigate whether the lock signal can be used to indicate missing input clock pulses, how a 2x output responds to a missing input clock pulse, and whether the DLL can tolerate the jitter from a TTCrx. TTCrx version 3 was found to have a long-term jitter, or drift, of 400ps and a cycle-to-cycle jitter of +/-300ps. There are TTCrx chips now available with improved jitter. The test was achieved using a Lecroy oscilloscope to histogram changes in the lock signal^{***} and the jitter between a reference and x2 DLL clock. The reference clock was also used to insert missing clock pulses. The results showed that there was no degradation found in the DLL x2 clock when using the TTCrx as a system clock. It takes 100µs to lose the lock signal after the TTCrx is turned off, i.e. no system clock. An irregular 14 out of 16 input clock pulses can be missing before the out of lock signal is given. The Lock signal can be used for Locked-on but not loss of lock.

V. Conclusion and Future Plans

A complete prototype HPD to DAQ readout system for one HPD that closely resembles the final solution of a single chain has been developed. Error rates from SEU and transmission have been studied and a detection method designed and incorporated. Bit error rates have been predicted to be less than 10^{-12} for the full readout chain and measured to be 3.87×10^{-15} for the Level_1 receiver board when in stand-alone operation. The TTCrx has been found to be compatible with

^{†††} Very high speed integrated circuit **H**ardware **D**escription **L**anguage.

^{***} The lock signal is given from the DLLs when synchronisation is achieved.

the DLLs used on the Level_1 board, and is fully operational on both Levels. Algorithms for the FPGAs have been written using VHDL, synthesised and using JTAG, downloaded. These algorithms have been tested at a modular level and work fully in functionality. The Level_1 board can store incoming events at 40MHz, perform the necessary formatting for the S-Link and send data out to a PC on a Level_1 trigger from a TTCrx. The optical link from Level_0 to Level_1 is in the process of being verified. The immediate task is to verify the optical transmission and to carry out extensive robustness tests in terms of transmission of known data over the entire chain. Known data are patterns that are generated in the Pint; errors can be injected to test the error detection methods. Work is also ongoing to port the Pint algorithms into an antifuse ACTEL AX1000 for radiation qualification. The general specifications are the same as the Spartan although the AX1000 has additional I/O.

VI. References

- [1] 'Development of an Optical Front-end Readout System for the LHCb RICH detectors', 7th Workshop on Electronics for LHC Experiments, CERN 2001-005
- [2] LHCb TP CERN/LHCC 98-4 LHCC/P4, 20 February 1998
- [3] LEB 1999 CERN 99-09 CERN/LHCC/99-33
- [4] GOL REFERENCE MANUAL, Preliminary version March 2001 CERN-EP/MIC, Geneva Switzerland
- [5] http://lhcb-elec.web.cern.ch/lhcelec/html/ecs_interface.htm
- [6] <http://micdigital.web.cern.ch/micdigital/ttcx.htm>
- [7] <http://www.lasermate.com/transceivers.htm>
- [8] CERN/LHCC/2000-0037 LHCb TDR3, 7 September 2000
- [9] CERN/LHCC 98-4 LHCC/P4, 20 February 1998
- [10] Daniel Baumeister, Development and Characterisation of a Radiation Hard Readout Chip for LHCb-Experiment, Thesis, Universitat Heidelberg, 2003.
- [11] Jorgen Christiansen. Requirements to the L0 front-end electronics, LHCb Technical Note, second version, revision 1.0, LHCb 2001-014, created July 1999, last modified, July 3, 2001.
- [12] Gillian Damerell. Error Checking in the LHCb_RICH Readout Chain, 2002 Paper in preparation. Contact Particle and Nuclear Physics dept, Oxford University.
- [13] Eric Brandin, "Development of a Prototype Read-out Link for the Atlas Experiment", Master Thesis, June 2000.

