

The Level 0 Pixel Trigger system for the ALICE experiment

G. Aglieri Rinella,^{a*} A. Kluge^a and M. Krivda^{ab}

On behalf of the ALICE Silicon Pixel Detector project

^a*CERN, European Organization for Nuclear Research,
CH-1211, Geneva 23, Switzerland*

^b*Institute of Experimental Physics,
Kosice, Slovakia*

E-mail: gianluca.aglieri.rinella@cern.ch

ABSTRACT: The ALICE Silicon Pixel Detector contains 1200 readout chips. Fast-OR signals indicate the presence of at least one hit in the 8192 pixel matrix of each chip. The 1200 bits are transmitted every 100 ns on 120 data readout optical links using the G-Link protocol. The Pixel Trigger System extracts and processes them to deliver an input signal to the Level 0 trigger processor targeting a latency of 800 ns. The system is compact, modular and based on FPGA devices. The architecture allows the user to define and implement various trigger algorithms. The system uses advanced 12-channel parallel optical fiber modules operating at 1310 nm as optical receivers and 12 deserializer chips closely packed in small area receiver boards. Alternative solutions with multi-channel G-Link deserializers implemented directly in programmable hardware devices were investigated. The design of the system and the progress of the ALICE Pixel Trigger project are described in this paper.

KEYWORDS: Particle tracking detectors; Electronic detector readout concepts (solid-state); Modular electronics; Trigger concepts and systems (hardware and software).

*Corresponding author.

Contents

1. Introduction	1
2. System architecture	2
3. Development, design and status	3
3.1 Optical receiver module	3
3.2 G-Link deserialization	4
3.3 Design and status	6
3.4 Control and configuration	9
4. Conclusions	9

1. Introduction

ALICE is an experiment designed to study the physics of strongly interacting matter and the properties of quark gluon plasma in the collisions between heavy ions nuclei at the Large Hadron Collider (LHC) at CERN [1]. The ALICE apparatus allows particle identification over a broad momentum range, powerful tracking with good resolution from 100 MeV/ c to 100 GeV/ c and excellent determination of secondary vertices. These features allow important contributions also to the physics of proton-proton interactions. The low material budget and the moderate magnetic field make the apparatus suited for studying low transverse momentum phenomena in proton-proton collisions.

The ALICE Silicon Pixel Detector (SPD) is the innermost detector of the Inner Tracking System of the ALICE apparatus [2]. It is shown in figure 1. The SPD is a double layer barrel pixel detector. It is constituted of 120 half staves, 40 in the inner layer and 80 in the outer one. One half staffe comprises two 200 μm thick silicon pixel sensors. Each sensor has a matrix of 160×256 pixels of $425 \times 50 \mu\text{m}^2$. A sensor is bump bonded to 5 mixed signal readout chips realized in 0.25 μm CMOS technology. The linear array of 10 pixel chips in each half staffe is read out via a Multi Chip Module (MCM) [3]. This includes the Analog Pilot chip for biasing, the Digital Pilot Chip for readout and communication, the GOL chip [4], optical PIN receivers and a laser diode transmitter. Timing and commands are received from the control room on two 40 Mb/s serial links. Data are transmitted on an optical link with a wavelength of 1310 nm using the G-Link protocol [5] at a signaling rate of 800 Mb/s. The readout of the data is initiated at the reception of a Level 2 trigger signal. Hit data are stored in the readout chips memory until a positive or negative Level 2 command is received. The readout and control electronics [6] includes 60 Link Receiver boards connected as mezzanine cards on 20 Router boards.

The 1200 readout chips of the SPD feature a Fast-OR output signal. This digital output is active whenever at least one of the 8192 channels of the chip records a hit. This information is




Figure 1. The Silicon Pixel Detector and a detailed view of one of its 120 half staves. The positions on the half stave of the sensors, of the pixel chips and of the readout Multi Chip Module (MCM) are indicated.

not stored in the on detector memory. The 10 Fast-OR bits of each half stave are transmitted by the MCM every 100 ns on the output optical link. They are transmitted using the user field of the G-Link control words. The Fast-OR signals allow the SPD to be operated as a low latency and low granularity pad detector, with an equivalent pad size of $\sim 13 \times 13 \text{ mm}^2$.

The 1200 Fast-OR signals will be used to generate an input signal for the Level 0 trigger decision in the ALICE Central Trigger Processor (CTP). Various trigger algorithms taking into account the Fast-OR data have been investigated [7], including topology based and occupancy based ones. These studies have shown that event selection in heavy ions runs and background rejection in proton-proton interactions can be significantly improved using the Fast-OR data. The different algorithms can be implemented as combinational logic functions of the 1200 Fast-OR signals. This naturally suggests an implementation of the algorithm on a programmable hardware device.

The Pixel Trigger System for the ALICE experiment is required to extract and process the 1200 Fast-OR signals in order to provide a signal input to the Level 0 trigger decision in the CTP. Various user selectable processing algorithms shall be supported by the hardware platform. The overall time latency from the interaction to the activation of the input to the CTP is required to be less than 800 ns by the specification of the ALICE CTP [8]. This requires the system to be located as close as possible to the detector and the CTP. A limited space of one standard crate could be allocated in the electronic racks next to the CTP.

2. System architecture

The architecture of the Pixel Trigger System is shown in figure 2. The 120 optical fibers outgoing from the detector are connected to the inputs of a commercial passive optical splitter located in the rack next to the CTP. The splitter output fibers forward the data to the readout electronics in the control room at about 110 m from the ALICE apparatus and to the electronic boards of the Pixel Trigger system, located in the same rack of the splitter. The data readout path is not affected by the Pixel Trigger System and remains fully independent from it.

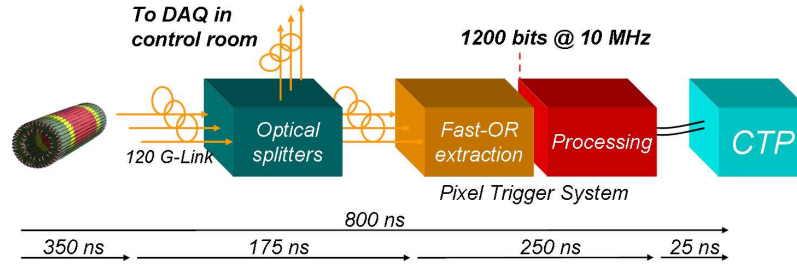


Figure 2. The architecture of the Pixel Trigger system.

The electronic system is subdivided in two subsystems. The first one deserializes the optical data and extracts the 1200 Fast-OR bits from the data flow. It is made up of a set of electronic boards with optical receivers and FPGAs. The second block implements the processing algorithm on the 1200 input bits and generates the output signal for the CTP. It is constituted of one electronic board based on a FPGA with a large number of pins and large logic space. The former architecture is naturally suggested by the need of providing all the 1200 Fast-OR bits as simultaneous inputs to the processing unit.

The overall latency budget of 800 ns must be subdivided among each of the processes along the data flow. The on-detector electronics takes 350 ns from the collision to transmit the Fast-OR bits. The optical fibers path length from the detector to the location of the system is ~ 35 m. This implies a signal propagation delay of ~ 175 ns. Therefore only about 250 ns can be allocated to the deserialization, extraction and processing phases. A detailed list of the latencies in the final system is given in section 3.3.

The algorithm to process the 1200 Fast-OR bits will be implemented in programmable hardware to allow fast execution, upgrading and reconfiguration by the user. The implementation on a large FPGA of some of the proposed trigger algorithms has been completed. The behavioral simulations showed that even the most complex function could be processed in less than 15 ns in a Xilinx Virtex 4 device. The critical delays in the system are therefore associated with the Fast-OR data deserialization, extraction and transfer between the peripheral FPGAs and the processing unit.

3. Development, design and status

Advanced optical receivers were chosen for the optical receiver cards after a qualification procedure. FPGA based solutions for the G-Link deserialization were also investigated. These development activities were motivated by the need of realizing the optical receiver cards as compact as possible. Some outcomes of these investigations are discussed in the following two subsections. The detailed description of the design is given after those.

3.1 Optical receiver module

Parallel optical fiber modules integrate an array of photodiodes and a receiver chip in a single device. Independent amplifying channels are realized on the chip. These devices allow for a significant reduction of the space required for the optical receivers section in a communication board. However devices operating at 1310 nm are not easily available on the market at the present




Figure 3. The Zarlink 12-channel Parallel Optical Fiber Module operating at 1310 nm. It features a male MPO/MTP fiber bundle connector on the optical input side and a 100 pin MEG-Array® electrical connector on the host circuit side.

time. Engineering samples of parallel optical fiber modules operating at 1310 nm were provided to our group by a private company.¹ They are customized versions of commercial devices operating at 850 nm. A photograph of the device is shown in figure 3. It features a 12 position MPO/MTP connector on the optical input side. A 100 pin MEG-Array² board-to-board connector allows the component to be plugged on the host board.

The performance of the two samples was experimentally evaluated with respect to sensitivity, bandwidth, overall jitter and transmission error rate. For these measurements an optical communication setup [9, 10] was used. The transmitter side included a pattern generator, the GOL radiation hard serializer and laser driver and a laser diode. The Zarlink optical module was used at the receiver side and its outputs could be connected to an oscilloscope or to a deserializer/receiver ASIC.³

All the channels of the Zarlink receivers were operational with input optical power as low as -18 dBm. The module specified bandwidth is 2.7 Gb/s. The contribution of the optical module to the overall jitter of the communication chain was measured by the eye diagrams. A value of ~ 25 ps was obtained at -15 dBm. The rate of frame (20 bit word) errors in the transmission was measured at various optical power levels. For these measurements the GOL was configured to transmit using the 8B/10B protocol. A Frame Error Rate below 10^{-14} with -18 dBm of input optical modulation amplitude was measured. This figure allowed an estimation of a bit error rate lower than 5×10^{-16} at -18 dBm. The engineering samples satisfied and even exceeded the specifications for the Pixel Trigger receiver cards. The Zarlink parallel optical fiber module was therefore chosen to be used for the Pixel Trigger system.

3.2 G-Link deserialization

The serial data stream has to be deserialized and frame aligned to allow tapping the Fast-OR bits out of the data flow. The on-detector electronics transmits data formatted according to the G-Link

¹Zarlink Semiconductor, <http://www.zarlink.com>.

²MEG-Array® is a registered trademark by FCI, <http://www.fciconnect.com>.

³Texas Instrument TLK1501.

Device	Latency [ns]		
	Data sheet	Simulated	Measured
Agilent HDMP-1034	88	n.a.	87±7
Altera Stratix GX ¹	n.a.	190	190±5
Altera Stratix GX ²	n.a.	130	n.a.
Virtex II Rocket IO X ³	475	475	n.a.
Virtex II Rocket IO X ⁴	275	n.a.	n.a.

¹Fully operational hardware implementation, 20 bit word length, 40 MHz clock.

²Behavioral simulation, 10 bit word length, 80 MHz clock.

³Simulation and data sheet, 20 bit word length, 40 MHz clock.

⁴Data sheet, 10 bit word length, 80 MHz clock.

Table 1. Latencies of alternative solutions investigated for the 800 Mb/s G-Link deserializers/receivers. The Agilent HDMP-1034 is a single channel dedicated ASIC. The remaining four solutions are multi-channel receivers implemented on FPGAs.

protocol. This protocol is fully supported by the GOL ASIC. However few products supporting the G-Link protocol can be found on the market. Agilent manufactures a single channel G-Link receiver ASIC (HDMP-1034).

Latest generation of FPGAs feature modules for serial communication up to few Gb/s. These modules include serializer/deserializers, comma alignment blocks, protocol decoders and are typically implemented in dedicated circuitry on the FPGA chip.⁴ These devices respond to the present needs for faster chip to chip and board to board data communication while reducing the number of interconnecting lines and board layout complexity. A board design entirely based on such devices, including the receiving and deserializing section of the system, would offer significant advantages in terms of board occupancy and system compactness. The implementation of a multi-channel G-Link deserializer/receiver on Xilinx and Altera FPGAs with fast serial circuitry was therefore investigated.

A parallel 12-channel G-Link receiver was designed and implemented by us in Hardware Description Language. Implementations of this circuit both on a Xilinx Virtex 2 with Rocket I/O X and on an Altera Stratix GX were tested by behavioral simulation. The circuit was also implemented and tested in real hardware using a larger purpose optical receiver board based on an Altera Stratix GX and including a parallel optical fiber module. This electronic board had been previously developed at CERN [11] by the group of the ECAL (electromagnetic calorimeter) detector of the CMS (Compact Muon Solenoid) experiment. To our knowledge this was the first working implementation of a parallel multi-channel G-Link receiver on programmable gate arrays with dedicated fast deserializing circuitry.

The latencies of the alternative G-Link receiving solutions were determined and are listed in table 1. The latencies of the FPGA based circuits were higher than the one offered by the HDMP-1034 ASIC. This was due to the lack of the possibility to bypass several functional blocks of the dedicated modules of the FPGAs, even if they were not actually used. Some examples are the comma alignment module and the 8B/10B decoder. Several clock periods were therefore wasted

⁴Both Altera and Xilinx implemented these featured in their latest Stratix and Virtex FPGAs.




Figure 4. The prototype OPTIN electronic board. The Zarlink Parallel Optical Fiber Module, six Agilent HDMP-1034 chips and the Xilinx Virtex 4 FPGA are shown. Six more HDMP-1034 chips and the Compact PCI connectors are on the other side of the board.

in the data path of the FPGAs high speed modules. The latency constraint on our application is not compatible with the performances obtainable with the FPGA based circuits. Therefore the HDMP-1034 was chosen as deserializer for the receiver cards of the Pixel Trigger system. The potential advantages offered by the solutions based only on FPGAs could not be exploited.

3.3 Design and status

The choice of the optical receiver modules and of the HDMP-1034 ASIC determined the architecture of the optical receiver electronic boards (OPTIN boards). The OPTIN board has been designed and the first prototype produced. It is shown in figure 4. It is realized on a 12 layer printed circuit board. It has four connectors compatible with the IEEE 1386 Common Mezzanine Card standard. However, the board is larger than the envelope specified in the standard. Its dimensions are $160 \times 84 \text{ mm}^2$. The assignment of the connector pins are compatible with the Compact PCI specification.

Each OPTIN board acts as a Fast-OR extractor for 12 optical channels incoming from the on detector modules via the splitter. One Zarlink optical module receiver converts and amplifies the signals received on the incoming fiber bundle. Twelve HDMP-1034 deserialize and realign the data streams. A Xilinx Virtex 4 LX60 FPGA receives the parallel outputs of the deserializers and extracts the 120 Fast-OR bits received every 100 ns. They are transmitted on a dedicated parallel output bus. Ten OPTIN boards are needed in total for the full system. They operate, as the rest of the system, at the LHC clock frequency of 40.0786 MHz. The power consumption of the OPTIN board is $\sim 13 \text{ W}$. A large fraction of this power, up to 8 W, is needed by the twelve HDMP-1034 chips.

The processing board (BRAIN) is a 9U size ($400 \times 360 \text{ mm}^2$) electronic board. It is shown in figure 5. Most of the board area is reserved to connect the 10 OPTIN boards as mezzanine cards, 5 per each side of the processing board. This interconnection solution is made possible by the

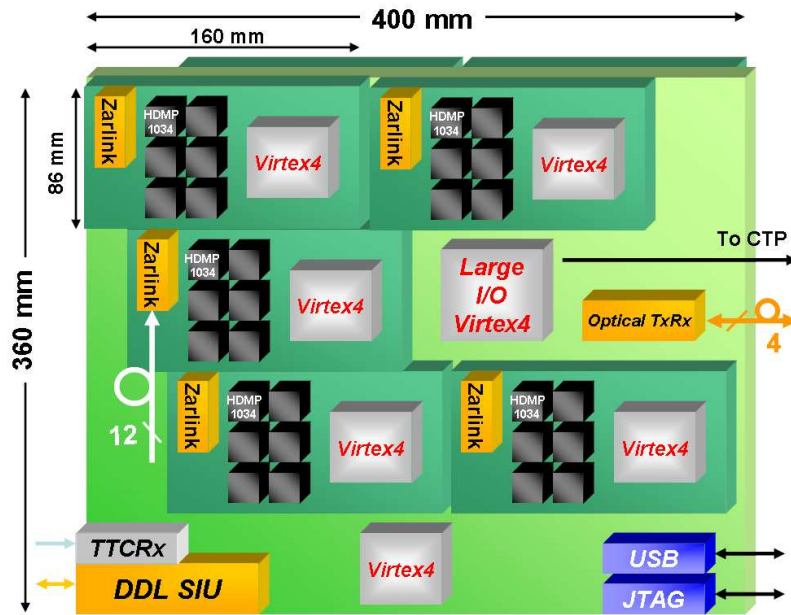


Figure 5. Schematic view of the BRAIN processing board with the OPTIN boards plugged as mezzanine cards. Five OPTIN cards are located on each side of the processing board. The processing and the control FPGAs are in evidence. The diagram also shows the interfaces of the system.

compact design of the OPTIN board. A similar approach has already been used for the SPD data readout electronics. The routing of the signals between the boards is thus simplified, avoiding the need of a backplane or wired interconnections.

A processing FPGA is the core of the BRAIN board. A large I/O space device, a Xilinx Virtex 4 LX100 with 1513 pins, has been chosen as the processing device because of the large number of parallel lines needed to receive the Fast-OR bits from the OPTIN boards with a minimum latency. The Fast-OR bits are transferred at Double Data Rate (80.16 MHz) on dedicated signal lines. There are 64 Fast-OR dedicated output lines on each OPTIN board. They pass through the OPTIN board connectors and reach dedicated input pins of the processing FPGA. In total 640 input pins of the processing FPGA are used for the Fast-OR parallel bus. This large degree of parallelism is needed to satisfy the constraint on latency. All these lines are routed as controlled impedance striplines. Source impedance matching is implemented using the Digitally Controlled Impedance DCI® feature of the Xilinx Virtex 4. Table 2 gives a detailed list of the latencies of the various processes along the data path in the system.

Several interfaces are available on the processing board. Ten main output lines (LVDS) allow the transmission of the result of the decision algorithm. The ALICE Detector Data Layer (DDL) interface [13] is the main slow control interface. USB and JTAG interfaces are available for *in situ* access and debugging. The TTCrx chip [14] is dedicated to receive the clock and timing information from the experiment Timing Trigger and Control distribution network. Auxiliary high speed serial lines and optical transceivers are available for high bandwidth communication and future upgrades. The design of the BRAIN board is finished. The layout and routing of the board is ongoing. The board shall need 12 routing layers.

Process	Duration [ns]	Integrated latency [ns]
Fast-OR readout (Digital Pilot output)	325	325
Serialization (GOL chip)	25	350
Transmission to the Pixel Trigger System (35 m long optical fibers)	175	525
Deserialization (HDMP-1034 chip)	100	625
Fast-OR data extraction	50	675
Transmission from OPTIN to BRAIN	75	750
Algorithm processing	25	775
Transmission to the CTP (~ 6 m)	25	800

Table 2. Latencies of the various processes on the Fast-OR data path for the final Pixel Trigger System architecture. Worst case values are considered for each process.

The input bandwidth of the system reaches 96 Gb/s (192 Gb/s) with the input links operating at 800 Mb/s (1600 Mb/s). The total power consumption of the BRAIN board and of the 10 OPTIN boards is estimated at about 140 W. The system layout implies that this power is dissipated in a single slot of the electronic rack. The dissipated power density is high and thermal verification of the system design was required. A finite element heat flow simulation of the system was implemented at CERN using dedicated software tools [12]. Detailed thermal models of all the electronic components and of the boards of the system were considered. The geometry of the simulated system was similar to the final one. The boards were located in the partially enclosed volume of an electronic rack, with no other boards in the adjacent slots as it is foreseen for the real case. The simulation considered forced air convection. The results of the calculation are shown in figure 6. The hot spots in the system are the voltage regulators on the OPTIN boards. They can reach a temperature of 110 °C. The regulators must operate with a junction temperature $T_j < 125$ °C. The result of the verification was judged acceptable because worst case conditions were assumed at each step in the definition of the thermal model.

The design was checked with respect to signal integrity because of the large number of lines, the length of the interconnections and the large number of loads on shared lines. A simulation study of the critical lines was done. It included IBIS⁵ models of the Virtex 4 output drivers as well as models of the signal propagation in the striplines. The signal analysis tools of the Cadence suite were used for this purpose. This study allowed for an optimization of the driving strength and switching speed of the FPGA output buffers. Signal buffering was shown to be required in some cases.

The location of the Pixel Trigger system is a radiation area with an expected neutron fluence [15] of 2.0×10^8 cm⁻² over 10 years. The iROC report [16] gives a measurement of the number of Single Event Upsets in a similar technology (90 nm) Xilinx Spartan 3 FPGA. That measurement can be used to estimate the number of upset events in the Pixel Trigger FPGAs. This calculation results in 18 Single Event Upsets (including configuration SRAM changes) and 5 Single Event Functional Interrupts (observable at the outputs) during 10 years of operation for each of the system FPGAs.

⁵I/O Buffer Information Specification, a standard to model the electrical properties of output line drivers.




Figure 6. Color map of the temperature of the system calculated by a finite element model. The highest temperature is around 110 °C. Worst case assumptions were made for the dissipated power and the thermal resistances.

3.4 Control and configuration

Full integration of the Pixel Trigger system in the ALICE Detector Control System is necessary. The main control is via the ALICE DDL interface. A second FPGA (Virtex 4 LX 40) on the processing board is dedicated to the slow control, to the system interfaces and to implement the reconfiguration of the main processing FPGA. Status monitoring and control is implemented via status and configuration registers in all the programmable devices of the system, accessed by the control FPGA via a local shared bus. Four JTAG chains are foreseen in the system. Two of them are used for two groups of 5 OPTIN boards, one is dedicated to the processing FPGA and the remaining to the control FPGA itself. JTAG players shall be implemented in the control FPGA to control the OPTIN boards and the processing FPGA chains. The processing FPGA will be reprogrammed remotely. The programming file for a given processing algorithm will be downloaded via the DDL link and the control FPGA to local SRAM memory. The control device shall then transfer the programming file to the Flash PROM connected to the processing FPGA by the JTAG players. Reconfiguration is finally launched via a JTAG command sent to the processing FPGA.

The control software shall be implemented in the ALICE PVSS framework [17]. An intermediate middleware layer (C++ library) will make the basic hardware functionalities available to the upper PVSS layer.

4. Conclusions

The Pixel Trigger System for the ALICE experiment has been designed and is being constructed. It allows for the 1200 Fast-OR signals from the detector to be used in the Level 0 trigger algorithm in the Central Trigger Processor. The system targets the stringent latency constraint of 800 ns. It is

fully independent from the existing data readout electronic chain. It allows for several algorithms to be implemented on the same programmable hardware platform. Remote configuration of the processing circuitry allows the user to change the processing algorithm at choice. The system has a modular design open to future upgrades and is based on several Field Programmable Gate Array devices. The system uses advanced parallel optical fiber modules in conjunction with 12 deserializer chips densely packed in the receiving boards. The ALICE experiment will be the only experiment at the LHC collider to include the silicon pixel detector in the Level 0 trigger decision from startup.

Acknowledgments

The authors thank A. Singovski, F. Vasey, P. Vichoudis for their support to the characterization of the parallel optical receiver modules and to the hardware implementation of the multi-channel G-Link receiver.

References

- [1] ALICE Collaboration, *ALICE Physics performance report*, CERN-LHCC-2003-049, *J. Phys. G* **30** (2004) 1517.
- [2] P. Riedler et al., *Overview and status of the ALICE silicon pixel detector*, Proceedings of the Pixel 2005 Conference, Bonn, Germany.
- [3] A. Kluge, *The ALICE silicon pixel detector front-end and read-out electronics*, *Nucl. Instrum. Meth. A* **560** (2006) 67.
- [4] P. Moreira et al., *GOL reference manual, gigabit optical link transmitter manual*, <http://proj-gol.web.cern.ch/proj-gol/>.
- [5] AGILENT TECHNOLOGIES, *Low cost gigabit rate transmit/receive chip set with TTL I/Os*, Technical Data, HDMP-1022/HDMP-1024 data sheet, December 2003.
- [6] M. Krivda et al., *Alice SPD readout electronics*, Proceedings of the 12th Workshop on electronics for LHC and future experiments, LECC 2006, Valencia.
- [7] J. Conrad et al., *Minimum bias triggers in proton-proton collisions with the VZERO and silicon pixel detectors*, Alice Internal Note, ALICE-INT-2005-025.
- [8] ALICE COLLABORATION, *Technical design report, trigger, data acquisition, high level trigger, control system*, CERN-LHCC-2003-062, ALICE TDR 010, 7 January 2004.
- [9] CMS COLLABORATION, J. Grahl and S. Corum, *Acceptance test report for NGK 12 channel opto-receivers for CMS ECAL data links*, CERN Note.
- [10] G. Aglieri Rinella, *Characterization of two Zarlink 1310 nm 12 channels optical receiver modules*, to be submitted as ALICE Internal note.
- [11] S. Reynaud and P. Vichoudis, *A multi-channel optical plug-in module for gigabit data reception*, Proceedings of the 12th Workshop on electronics for LHC and future experiments, LECC 2006, Valencia.
- [12] E. Dupont, *Electronics cooling with FLOTHERM: the Level 0 pixel trigger system for the ALICE silicon pixel detector*, CERN TS, CFD-2006-05, June 2006.

- [13] CERN ECP/ALD, RMKI RFFO, *ALICE detector data link user requirements document*, Alice Internal Note, ALICE-INT-1996-042.
- [14] J. Christiansen et al., *Receiver ASIC for timing, trigger and control distribution in LHC experiments*, *IEEE Trans. Nucl. Science*, **43** (1996) 1773.
- [15] A. Morsch and B. Pastircak, *Radiation levels in the ALICE detectors and electronic racks*, Alice Internal Note, ALICE-INT-2002-028.
- [16] IROC TECHNOLOGIES, *Radiation results of the SER test of Actel, Xilinx and Altera FPGA instances*, October 2004, <http://www.iroctech.com>.
- [17] PVSS SERVICE, <http://itcobe.web.cern.ch/itcobe/Services/Pvss/welcome.html>.