Emulating the GLink Chip-Set with FPGA Serial Transceivers in the ATLAS Level-1 Muon Trigger

Alberto Aloisio, Francesco Cevenini, Raffaele Giordano and Vincenzo Izzo

Abstract—**Many High Energy Physics experiments based their serial links on the Agilent HDMP-1032/34A serializer/deserializer chip-set (or GLink). This success was mainly due to the fact that this pair of chips was able to transfer data at $\sim$ 1 Gb/s with a deterministic latency, fixed after each power up or reset of the link. Despite this unique timing feature, Agilent discontinued the production and no compatible commercial off-the-shelf chip-sets are available. The ATLAS Level-1 Muon trigger includes some serial links based on GLink in order to transfer data from the detector to the counting room. The transmission side of the links will not be upgraded, however a replacement for the receivers in the counting room in case of failures is needed.**

In this paper, we present a solution to replace GLink transmitters and receivers. Our design is based on the gigabit serial IO (GTP) embedded in a Xilinx Virtex 5 Field Programmable Gate Array (FPGA). We present the architecture and we discuss parameters of the implementation such as latency and resource occupation. We compare the GLink chip-set and the GTP-based emulator in terms of latency, eye diagram and power dissipation.

Index Terms—**Serial links, fixed latency, FPGAs.**

I. INTRODUCTION

TRIGGER systems of High Energy Physics (HEP) experiments need data transfers to be executed with fixed latency, in order to preserve the timing information. This requirement is not necessarily satisfied by Serializer-Deserializer (SerDes) chip-sets. The Gigabit link, or GLink, chip-set [1], produced by Agilent, was able to transfer data at data-rates up to 1 Gb/s with a fixed latency even after a power-cycle or a loss of lock. Serial links of data acquisition systems of HEP experiments have been often based the GLink chip-set. For instance it has been deployed in the Alice [2], ATLAS [3], Babar [4], CDF [5], CMS [6], D0 [7] and Nemo [8] experiments (just to cite some of them). The chip-set became so popular, that CERN produced a radiation hard serializer compatible with it [9]. Unfortunately, a few years ago Agilent discontinued the production of the chip-set and users needing replacements are looking for alternative solutions. Latest FPGAs include embedded multi-Gigabit SerDes, which offer a wide variety of configurable features. The benefit from the integration of such a device in FPGA is obvious in terms of power consumption, size, board layout complexity, cost and re-programmability. The Level-1 Barrel Muon Trigger of the ATLAS experiment

includes GLink serial links in order to transfer data from the detector to the counting room. The transmission side of the links is on-detector and will unlikely be upgraded, however a replacement for the receivers in the counting room in case of failures is needed. We developed a replacement solution for GLink transmitters and receivers, based on the gigabit serial IO (GTP) embedded in Xilinx Virtex 5 Field Programmable Gate Array (FPGA). Our solution preserves the fixed-latency feature of the original chip-set. In the coming sections we will introduce the present L1 Barrel Muon Trigger and the GLink chip-set, then we will describe the architecture and the implementation of our design. Eventually we will present some test results about our emulator, comparing them also with the GLink chip-set.

II. ATLAS BARREL MUON TRIGGER AND DAQ

The ATLAS detector [10] is installed in one of the four beam-crossing sites at the Large Hadron Collider (LHC) of CERN. The detector has a cylindrical symmetry and it is centered on the interaction point. ATLAS consists of several subsystems, among them there is a muon spectrometer, which in the barrel region is built in the loops of a air-core toroidal magnet and includes Resistive Plate Chambers (RPCs). RPCs are arranged in towers used for the Level-1 (L1) muon trigger (Fig. 1). The spectrometer is divided in two halves along the axis and each half is in turn divided in 16 sectors. A physical sector is segmented in two trigger sectors, including 6 or 7 RPC towers each.

The whole trigger system is implemented as a synchronous pipeline, with a total latency of 2.5 μs, clocked by the Timing, Trigger and Control (TTC) system [11] of the LHC. The TTC distributes timing information such as the bunch crossing clock (at about 40 MHz) and the L1 trigger. It also provides synchronization signals for the bunch crossing identificator (BCID) and the event identificator (EVID) counting the L1 accepted triggers. The trigger system handles more than 350k detector channels and provides an average L1-accept rate of 100 kHz. The algorithm is based on geometrical coincidences within detectors of a tower or adjacent towers. It is able to identify tracks and classify muons with respect to their transverse momentum. It also tags the tracks with the pertaining BCID and EVID.

The read-out and trigger electronics of the barrel muon spectrometer includes an on-detector part and an off-detector one. For each RPC tower there are two PAD boards: the "Low-p_TPAD box" mounted the middle detector and the "High-p_TPAD box" on the outer detector. Each PAD hosts four Coincidence Matrix ASICs (CMAs) [12], which receive data
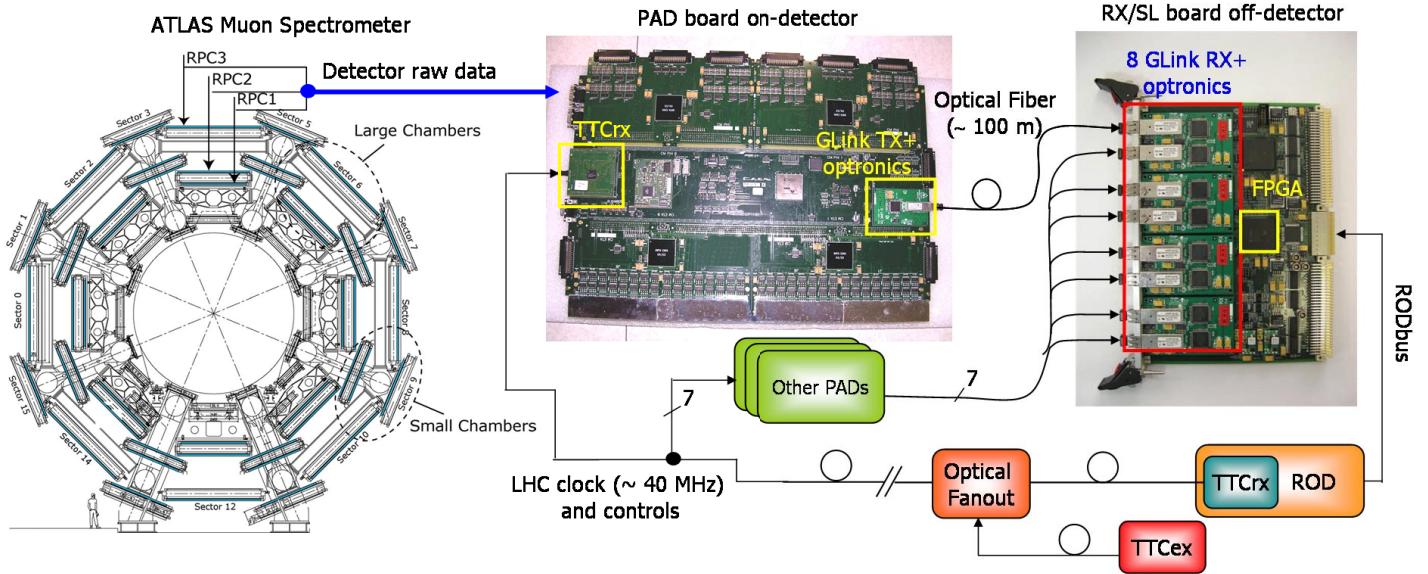
Figure 1. Left: Cross section of the ATLAS muon spectrometer. Right: Level-1 Trigger and DAQ for the spectrometer.

from the front-end electronics, execute geometrical trigger algorithms and provide read-out functionalities. The High-p_TPAD box can send triggers and data from the detectors (and from the Low-p_T PAD box) to a off-detector VME board, the Sector Logic/RX (SL/RX) [13], via an 800-Mbps serial link based on the GLink chip-set. Each SL/RX board includes 8 GLink receivers and two FPGAs handling the received data and the communication with other off-detector boards: the Read Out Driver (ROD) [14] and the Muon Central Trigger Processor Interface (μCTPI). The boards lie in a VME crate, but due to the large amount of data to be transferred with a fixed latency, they are also connected by a dedicated bus (RODbus).

There is one SL/RX board for each trigger sector, thus each board receives data and trigger information from 6 or 7 High-p_TPADs. The RX/SL executes a sector trigger algorithm and sends results to a μCTPI, which in turn transfers trigger information to the Muon Central Trigger Processor (μCTP). The μCTP collects trigger frames from all the interfaces and decides whether or not to accept the event with a fixed-latency of 2.5 μs. During the decision, data are stored by the on-detector electronics. If the event is validated, a L1 accept signal is broad-casted to the PADs, which tag data with the EVID and BCID and transfer it to the RX/SL. The RX/SL boards send data to the adjacent ROD, which checks the data frames, performs the event building within the sector and formats the data according to a specification common to all the ATLAS RODs. Data are eventually transferred to the Read Out System for further elaborations or storage.

III. THE GLINK CHIP-SET

The GLink chip-set consists of a serializer (HDMP-1032A) and a deserializer (HDMP-1034A). The chips work with data-rates up to 1 Gb/s and encode data according to Conditional Inversion Master Transition (CIMT) protocol. In order to read serial data, the receiver extracts a clock from the CIMT stream

and locks its phase to the master transition. The recovered clock synchronizes all the internal operations of the receiver and it is available as an output. Received data is transferred out of the device synchronously with the recovered clock and the chip-set architecture is such that the overall link latency is deterministic. Moreover, by means of the dedicated Parallel Automatic Synchronization System (PASS), it is also possible to output data synchronously with a local receiver clock, provided that it has a constant phase relationship with the transmission clock (like it happens in the ATLAS L1 barrel muon trigger, which is clocked by the LHC machine clock).

We now briefly introduce the CIMT encoding protocol. A CIMT stream is a sequence of 20-bit words, each containing 16 data bits (D-Field) and 4 control bits (C-Field). The C-Field flags each word as a data word, a control word or an idle word. Idle words are used in order to synchronize the link at start-up and to keep it phase-locked when no data or control words are transmitted. The protocol guarantees a transition in the middle of the C-Field and the receiver checks for this transition in received data in order to perform word alignment and to detect errors. Two encoding modes are available: one compatible with older chip-sets and an enhanced one, which is more robust against incorrect word alignment. However, previous studies indicated that at start-up the receiver can achieve fake lock conditions if the word sent is not an idle word [15]. The DC-balance of the link is ensured by sending inverted or unaltered words in such a way to minimize the bit disparity, defined as the difference between the total number of transmitted 1s and 0s. By reading the C-Field content, the receiver is able to determine if a word is inverted or not and restore its original form.

IV. GLINK EMULATION

We built our GLink emulator around the Xilinx GTP transceiver [16], embedded in Virtex 5 [17] FPGAs. Other FPGA vendor offer embedded SerDes, for instance Altera
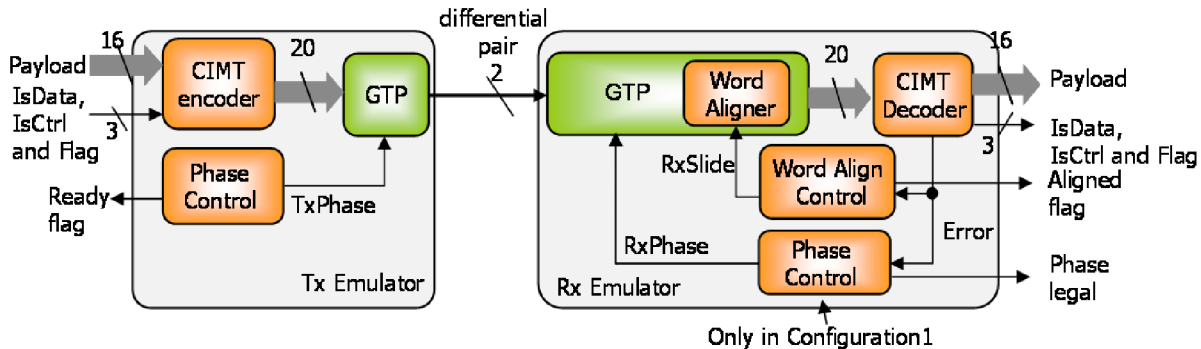
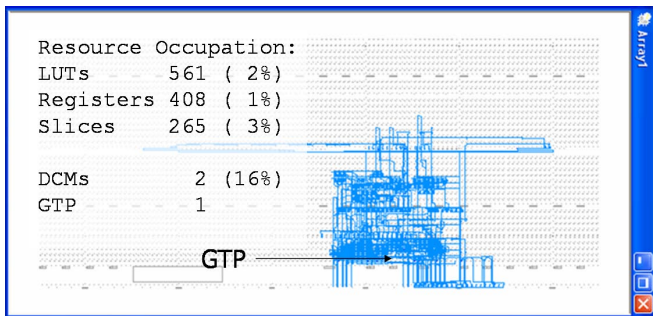Figure 2. Simplified block diagram of the emulator.

Figure 3. Resources used by an implementation of a GLink transmitter/receiver in a Xilinx Virtex 5 LX50T.

with the GX and Lattice with the flexiPCS. However, the fixed-latency characteristic of our emulator is deeply-based on some hardware features of the GTP. For a discussion about the possibility to implement a fixed-latency link with FPGA-embedded SerDes see [18].

A. Architecture

The GTP can serialize/de-serialize words 8, 16, 10 and 20 bit wide. We configured it to work with 20-bit CIMT-encoded words at 40 MHz, to achieve a 800 Mb/s link. The receiver clock has an unknown, but fixed, phase offset with respect to the transmitter clock. In order to transfer data with minimum latency the GTP allows to skip internal elastic buffers, one being in the data-path of the transmitter and the other one in the data-path of the receiver. When skipping buffers, all phase differences must be resolved between the external parallel clock domain and a clock domain internal to the device. We set up the transmitter to work without the elastic buffer, while we left two options for the receiver: the first one without the buffer and with an improved latency (Configuration1), but with some constraints on the relative phase between transmission and reception clocks and the second one without any phase constraint, but with a higher latency (Configuration2).

On the transmitter, a phase control logic instructs the GTP to align the phase of the internal clock to the transmission clock and asserts the Ready signal when done. A logic encodes input 16-bit words into 20-bit CIMT words and transfers them to the GTP (Fig. 2). The encoder is able to send data, control or idle words and supports an input flag bit exactly like the original chip-set.

On the receiver side, when working in Configuration1, the phase align and control logic checks whether or not it is possible to retrieve data from the link with the assigned parallel clock phase. If it is not possible the phase must be changed either in the FPGA or outside. In Configuration2 every phase offset is legal, therefore no checks are performed. In order to align received data to the correct word boundary, we added to the GTP: a CIMT decoder and a word align control logic. The decoder checks the C-Field of incoming CIMT words and, if it is not valid, flags an error to the word align control logic. When that happens, the logic drives the RxSlide signal causing the GTP shift parallel data by one more bit, then it ignores further errors, for a number of clock cycles equal to the overall latency of the shifting operation and the CIMT decoder. This process is repeated for each error received. If, for 256 consecutive clock cycles, no error is found, the align control logic assumes the parallel data is correctly aligned to the word boundary and asserts the ALIGNED signal. Of course, the decoder determines if the received word is an idle, a control or a data word and the status of the flag and activates the corresponding outputs.

For the sake of completeness, we inform the reader that our emulator supports all the CIMT encoding modes of the HDMP-1032/34A chip-set, but not the 20/21-bit modes of the older HDMP-1022/24.

B. Implementation

A full-duplex emulator (transmitter and receiver) requires around 500 Look Up Tables (LUTs) and 400 Flip Flops (FFs), which are 3% of the logic resources available in a Xilinx Virtex 5 LX50T FPGA (Fig. 3). Such a tiny resource requirement, will allow to integrate all the eight GLink receivers of the RX/SL board in the FPGA and the impact of this integration will be just a 6% of the fabric resources.

The latencies of the transmitter and the receiver are respectively 6.75 and 5.25 clock cycles (6.75 in Configuration2). Details about the contribution of internal blocks are given in Tab. I. For each component we report the latencies in terms of clock cycles and the absolute value. For comparison with the latencies of our solution we recall that latencies of the GLink transmitter and receiver are respectively 1.4 and 3.0 parallel

Table I
LATENCY OF THE BUILDING BLOCKS OF THE LINK (RECEIVER IN CONFIGURATION1).

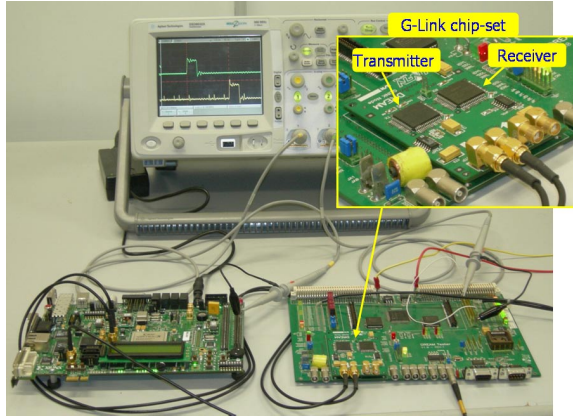|	# of clock cycles	Block latency (ns)
Transmitter		
Total Encoding Latency (fabric)	4.5	112.5
Total GTP Latency	2.25	56.25
Total Transmitter Latency	6.75	168.75
Receiver		
Total GTP Latency	4.75	118.75
Total Decoding Latency (fabric)	1	25
Total Receiver Latency	5.25	143.75
Total Link Latency	12	312.5

Figure 4.   Experimental setup for compatibility tests with the GLink chip-set.

clock cycles. Hence, our emulator has a higher latency with respect to the original chip-set, however this is not an issue for our application.
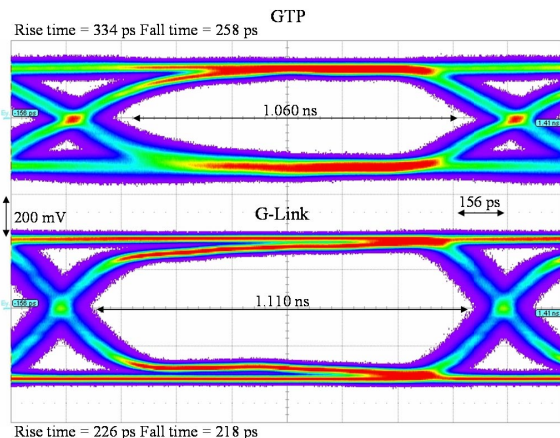

Figure 5.   Eye diagram comparison between GLink and the GTP.

V. TEST RESULTS

In order to test our link, we deployed two off-the-shelf boards [19] built around a Virtex 5 LX50T FPGA. The boards route the serial I/O pins of one of the GTPs on the FPGA to SMA connectors. We connected the transmitter and the receiver GTPs with a pair of 5 ns, 50 Ω impedance coaxial cables. Transmitted and received payloads were available on single ended test-points as well as on LVDS SMA connectors and were monitored by an oscilloscope to observe latency variations. We used a dual channel clock generator providing two 40-MHz clock outputs with fixed phase offset. This way, we emulated the TTC system of the ATLAS experiment, which is used to clock data in and out from the link.

We checked that our emulator is able to correctly transmit (receive) data toward (from) an Agilent GLink receiver (transmitter) chip in all the encoding modes supported by the HDMP-1032/34A chip-set. In order to perform this test, we deployed a ML-505 board and a custom board hosting a GLink transmitter and a receiver (Fig. 4). The test showed that the emulator correctly exchanges data with a GLink chip in both the CIMT encoding modes.

We present an eye diagram comparison between the Agilent GLink transmitter and the GTP (Fig. 5). We fed the transmitters with the same payload, a 16-bit pseudo random word sequence. We probed the signal on the positive line of the differential pair, at the far end of a 5 ns 50 Ω coaxial cable. Between the transmitter and the cable, there was a 10 nF decoupling capacitor. We terminated the negative line on its characteristic impedance to keep the differential driver balanced. We notice that the GLink eye width is 50 ps wider than GTP's. Despite the GTP smaller voltage swing (400 mV) with respect to GLink (600 mV), the latter has rise and fall times respectively around 30% and 15% lower. The timing jitter on GTP's edges is $\sim$ 210 ps, while for Agilent transmitter is $\sim$ 180 ps. This difference could be due to the fact that the generation of high-speed serial clock, from the 40-MHz oscillator, requires only the internal PLL for GLink, while in our clocking scheme for the GTP we deployed an FPGA DLL to multiply the 40-MHz clock to obtain the 80-MHz clock. Therefore, the total jitter on the transmitted serial stream includes the contribution of the jitters of both the PLL and the DLL. Moreover, we used a single ended oscillator to source the PLL of the GTP, while the User Guide recommends to use a differential oscillator.

We performed Bit Error Ratio (BER) measurements on the link implemented with our emulator. We deployed a custom Bit Error Ratio Tester (BERT) [20], checking the received payload against a local copy and flagging an error when a difference occurred. More than 10^{13} bits have been transferred and no errors have been observed, corresponding to a 10^{-12} BER, estimated with a 99% confidence level [21].

VI. CONCLUSIONS

SerDes embedded in FPGAs have a lower power dissipation with respect to external SerDes chip-sets and their data-rates and transmission protocols can be changed by simply re-programming the FPGA. By suitably configuring a GTP

transceiver and adding few logic resources from the FPGA fabric (~ 3% of the total), we have been able to achieve a complete replacement for the GLink chip-set. Our emulator transfers data with a fixed latency, which was a crucial feature of the original chip-set. We experimentally verified the compatibility of our emulator with GLink both in transmission and reception. The emulator has a tiny footprint in terms of logic resources and, in a future upgrade of the RX/SL, it will allow us to integrate all the GLink receivers on the board in a single FPGA, still leaving most of the device resources free to for trigger and readout tasks. Hence, the layout of the upgraded board would be simplified with respect to the present. A GLink receiver dissipates ~ 800 mW (typical @1 Gb/s) while each GTP pair (transmitter and receiver) dissipates ~ 300 mW (typical @3 Gb/s), hence also the overall power dissipation of the board will be lowered in the upgrade.

ACKNOWLEDGMENT

REFERENCES

[1] (2001). Agilent HDMP 1032-1034 Transmitter-Receiver Chip-set Datasheet, Agilent [On-line]. Available: http://www.physics.ohio-state.edu/~cms/cfeb/datasheets/hdmp1032.pdf

[2] Kluge, A. for the ALICE SPD collaboration, "The Read-Out system of the ALICE pixel detector", Proceedings of the International Workshop on Semiconductor Pixel Detectors for Particles and X-Rays (PIXEL2002)

[3] J. Garvey et al., "The ATLAS Level-1 Calorimeter Trigger Architecture", IEEE Transactions on Nucl. Sci., Vol. 51, no. 3, June 2004

[4] P. Sanders, "The BaBar trigger, readout and event gathering system", IEEE Trans. on Nucl. Sci., Vol. 45, Issue 4, Part 1, August 1998 pp. 1894-1897

[5] W. Ashmanskas et al., "Silicon vertex tracker: a fast precise tracking trigger for CDF", Nuclear Instruments and Methods in Physics Research Section A, Vol. 447, Issues 1-2, June 2000, pp. 218-222

[6] B. Lofstedt On behalf of the CMS ECAL Collaboration, "The digital readout system for the CMS electromagnetic calorimeter", Volume 453, Issues 1-2, 11 October 2000, pp. 433-439

[7] James T. Linnemann, "The D0 Level 2 Trigger" International journal of modern physics. C, Physics and computers, vol. 16, issue 1, 2001, p. 1166

[8] F. Ameli, "The Data Acquisition and Transport Design for NEMO Phase 1", IEEE Trans. on Nucl. Sci., Vol. 55, Issue 1, Part 1, Feb. 2008 pp. 233-240

[9] P. Moreira, T. Toifl, A. Kluge, G. Cervelli, F. Faccio, A. Marchioro, J. Christiansen, "GLink and gigabit Ethernet compliant serializer for LHC data transmission", Nuclear Science Symposium Conference Record, 15-20 Oct. 2000, Vol. 2 , pp. 9/6 - 9/9

[10] ATLAS Collaboration, ATLAS Detector and Physics Performance - Technical Design Report - Volume I, May 1999 [On-line]. Available: http://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/TDR/physics_tdr/printout/Volume_I.pdf

[11] B.G. Taylor for the RD12 Project Collaboration, TTC Distribution for LHC Detectors, IEEE Trans. on Nucl. Sci., Vol. 45, No. 3, June 1998, pp. 821-828.

[12] F. Pastore, E. Petrolo, R. Vari, S. Veneziano, "Performances of the Coincidence Matrix ASIC of the ATLAS Barrel Level-1 Muon Trigger", Proceedings of the 11th Workshop on Electronics for LHC Experiments, Heidelberg, Germany, 12-16 Sept 2005.

[13] G.Chiodi, E.Gennari, E.Petrolo, F.Pastore, A.Salamon, R.Varia, S.Veneziano, "The ATLAS barrel level-1 Muon Trigger Sector-Logic/RX off-detector trigger and acquisition board", Topical Workshop on Electronics for Particle Physics, Prague, Czech Republic, 03 - 07 Sep 2007, pp.232-237

[14] A. Aloisio et al., "The Read-Out Driver for the RPC of the ATLAS Muon Spectrometer", IEEE Trans. Nuclear Science, Vol. 55, no.1, part 1, 2008, pp. 265-271

[15] A. Aloisio, F. Cevenini, V. Izzo, "Do's and Don'ts With the Agilent's GLink Chipset ", IEEE Trans. on Nucl. Sci., vol. 53, no. 3, June 2006

[16] (2008). Virtex-5 FPGA RocketIO GTP Transceiver User Guide - UG196 (v1.7), Xilinx [On-line]. Available: http://www.xilinx.com/support/documentation/user_guides/ug196.pdf

[17] (2008). Virtex-5 FPGA User Guide - UG190 (v4.3), Xilinx [On-line]. Available: http://www.xilinx.com/support/documentation/user_guides/ug190.pdf

[18] A. Aloisio, F. Cevenini, R. Giordano, V. Izzo, "High-speed, fixed-latency serial links with FPGAs", Nuclear Science Symposium Conference Record, 2008, Dresden, Germany, pp. 2137-2142

[19] (2008). ML505/ML506/ML507 Evaluation Platform User Guide - UG347 (v3.0.1), Xilinx [On-line]. Available: http://www.xilinx.com/support/documentation/boards_and_kits/ug347.pdf

[20] A. Aloisio, F. Cevenini, R. Cicalese, R. Giordano, V. Izzo, "Beyond 320 Mbyte/s With 2eSST and Bus Invert Coding on VME64x", IEEE Transactions on Nuclear Science, Volume 55, Issue 1, Feb. 2008, pp. 203-208

[21] (2007). Statistical Confidence Levels for Estimating Error Probability, Maxim [On-line]. Available: http://pdfserv.maxim-ic.com/en/an/AN1095.pdf