

The Serial Link Processor for the Fast TrackKer (FTK) processor at ATLAS

P. Luciano^{*bn}, A. Annovi^e, A. Andreani^{fg}, R. Beccherle^b, M. Beretta^e, N. Biesuz^a, W. Billereau^h, J.M. Combe^h, M. Citterio^f, S. Citraro^a, F. Crescioliⁱ, D. Dimas^l, S. Donati^a, C. Gentsos^m, P. Giannetti^b, K. Kordas^m, A. Lanza^c, V. Liberali^{fg}, D. Magalotti^{do}, P. Neroutsos^m, S. Nikolaidis^m, M. Piendibene^a, E. Rossi^a, A. Sakellariou^l, S. Shojaii^{fg}, C.-L. Sotiropoulou^m, A. Stabile^f, P. Vulliez^h

^aUniversity of Pisa, Largo B. Pontecorvo 3, 56127 Pisa, Italy

^bSezione di Pisa INFN, Largo Bruno Pontecorvo 3, 56127 Pisa, Italy

^cSezione di Pavia INFN, Via Agostino Bassi 6, 27100 Pavia, Italy

^dUniversity of Modena and Reggio Emilia, Via Università' 4, 41121 Modena, Italy

^eLaboratori Nazionali di Frascati INFN, Via Enrico Fermi 40, 00040 Frascati, Italy

^fSezione di Milano INFN, Via G. Celoria 16, 20133 Milano, Italy

^gUniversità degli Studi di Milano, Department of Physics, Via G. Celoria 16, 20133 Milano, Italy

^hCERN, CH-1211 Geneva 23, Switzerland

ⁱLaboratoire de Physique Nucléaire et de Hautes Energies, Couloir 12-22 étage 4 Place Jussieu, 75005 Paris, Francia

^lPrisma Electronics SA, El Venizelou 128 Nea Smyrni, 17123 Athens, Greece

^mAristotle University of Thessaloniki, 54124 Thessaloniki, Greece

ⁿUniversity of Cassino and Southern Lazio, Via Guglielmo Marconi 10, 03043 Cassino, Italy

^oSezione di Perugia INFN, Via Alessandro Pascoli, 06100 Perugia, Italy

E-mail: alberto.annovi@lnf.infn.it, Alessandro.Andreani@mi.infn.it,
Roberto.Beccherle@pi.infn.it, matteo.beretta@lnf.infn.it,
nicolo.vladi.biesuz@cern.ch, william.billereau@cern.ch,
jean-marc.combe@cern.ch, Mauro.citterio@mi.infn.it,
saverio.citraro@for.unipi.it, francesco.crescioli@pi.infn.it,
ddimas@prisma.gr, simone.donati@pi.infn.it,
cgentsos@physics.auth.gr, cgentsos@physics.auth.gr,
Kostas.Kordas@cern.ch, Agostino.Lanza@pv.infn.it,
valentino.liberali@unimi.it, pierluigi.luciano@pi.infn.it,
daniel.magalotti@pg.infn.it, wnerouts@physics.auth.gr,
snikolaid@physics.auth.gr, marco.piendibene@pi.infn.it,
enrico.rossi@pi.infn.it, asakel@prisma.gr,
seyedruhollah.shojaii@cern.ch, lsoti@physics.auth.gr,
alberto.stabile@unimi.it, pascal.vulliez@cern.ch



The Associative Memory (AM) system of the Fast TracKer (FTK) processor has been designed to perform pattern matching using the hit information of the ATLAS silicon tracker. The AM is the heart of FTK and it finds track candidates at low resolution that are seeds for a full resolution track fitting. To solve the very challenging data traffic problem inside FTK, multiple designs and tests have been performed. The currently proposed solution is named the "Serial Link Processor" and is based on an extremely powerful network of 2 Gb/s serial links. This paper reports on the design of the Serial Link Processor consisting of the Associative Memory (AM) chip, an ASIC designed and optimized to perform pattern matching, and two types of boards, the Local Associative Memory Board (LAMB), a mezzanine where the AM chips are mounted, and the Associative Memory Board (AMB), a 9U VME board which holds four LAMBs. We report also on the performance of a first prototype based on the use of a min@sic AM chip, a small but complete version of the final AM chip, built to test the new and fully serialized I/O. Also a dedicated LAMB prototype, named miniLAMB, with reduced functionality, has been produced to test the mini@sic. The serialization of the AM chip I/O significantly simplified the LAMB design. We report on the tests and performance of the integrated system of the miniAsic, miniLAMB and AMB.

*International Conference on Technology and Instrumentation in Particle Physics
2-6 June 2014
Amsterdam, The Netherlands*

*Speaker.

1. Introduction

The trigger system of a detector installed at a hadron collider must have high efficiency for the interesting physics processes and it must suppress the enormous QCD backgrounds. A multilevel trigger [1] is an effective solution for this task. The ATLAS trigger system [2, 3, 4] is organized in three levels. The hardware Level-1 Trigger quickly locates the regions of interest in the calorimeter and the muon system, operating with output event rates up to 100 KHz. The subsequent trigger levels, Level-2 and the Event Filter (EF) are collectively known as the high-level trigger (HLT). They consist of software algorithms running on a farm of commercial CPUs.

In the current ATLAS trigger, the information that comes from the silicon detectors are not used at Level-1 and only late and locally at Level-2. Early reconstructed tracks increase the performance of the system. The Fast TracKer processor (FTK) [5] is a hardware-based system designed for track reconstruction in the silicon detectors with offline quality and in time for Level-2 selections.

The FTK processor is highly parallel. The detector is segmented into $\eta - \phi$ towers, each one processed by its own tracking processor. Each processor covers one sixteenth of the detector in ϕ (22.5° , plus 10° overlap to maintain high efficiency) while the η range of each region is divided into four overlapping intervals, for a total of 64 $\eta - \phi$ towers. Consequently, a tower receives only a fraction of the silicon hits, and a track reconstruction processor has substantially fewer candidates to process. Within each tower, we distribute the silicon data on 12 parallel buses at the full 100 KHz Level-1 output rate.

The track finding inside each detector tower is executed in two steps working in pipeline: the pattern matching at low resolution and the track fitting at full resolution.

The time consuming pattern recognition stage, generally referred to as the “combinatorial challenge”, can be solved by the Associative Memory (AM) technology [6] exploiting parallelism to the maximum level: it compares the clusters found in the event (“hits”) to all pre-calculated “expectations” or “patterns” (pattern matching) at once, searching for candidate tracks called “roads”. This approach reduces the typical exponential complexity of the CPU based algorithms into a linear problem.

Figure 1 shows the architecture of the FTK processor of the ATLAS detector. Tracking is implemented in the Processing Unit (PU), that is made of an AM board (pattern matching, PM) and an AUXiliary card (AUX) for Track Fitting (TF). The pixel and microstrip detector data are transmitted from the front end ReadOut Drivers (RODs) to the Data Formatters (DFs) which perform cluster finding. The barrel layers and the forward disks are grouped into logical layers so that there are 12 layers over the full rapidity range. The DFs organize the detector data for the FTK tower structure of the core crates, taking the needed overlap into account. The cluster centroids in each logical layer are sent to the Data Organizers (DOs). The DO is a smart database, where hits are stored at full resolution, converted to a coarser resolution (super-strips, SS) and sent to the AM for pattern matching.

The AM boards contain a very large number of low resolution candidate tracks, called patterns. The AM compares each hit with all patterns nearly simultaneously. Patterns matching the event (roads) are sent back to the DOs which immediately fetch and send to the TF the associated full resolution hits. The TF is executed only for hit combinations inside the found roads. Because

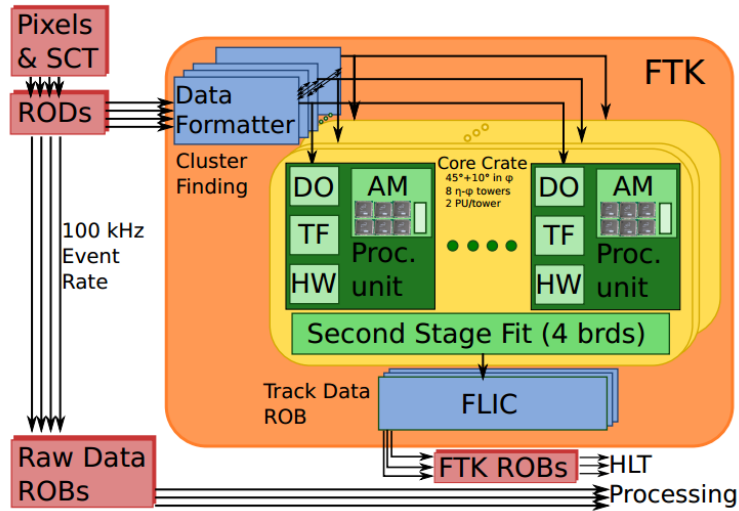


Figure 1: Architecture of the Fast TracKer processor

each road is quite narrow, the TF can provide high resolution helix parameters using the average parameters across the relevant tracking modules and applying corrections that are linear in the actual hit position in each layer. Fitting a track is thus extremely fast since it consists of a series of multiply-and-accumulate steps. In a modern Field Programmable Gate Array (FPGA), approximately 10^9 track candidates can be fit per second. In FTK, the first track fitting stage (TF) uses 8 silicon layers. All 12 layers are used in the Second Stage Fit. Duplicate track removal (the Hit Warrior function) is carried out among those tracks that pass the χ^2 cut. All these functions are executed in a pipeline inside FTK. The FLIC boards provide the connection to the HLT farms.

In summary: FTK has a very large number of devices organized in pipelines connected by thousands of serial links; there are 8200 dedicated custom chips (AM chips) that perform pattern matching and 2000 FPGAs for all other functions.

2. Associative Memory System

The AM system consists of the Associative Memory chip (AMchip), an ASIC designed and optimized for this particular application, and two types of boards, a 9U Versa Module Eurocard (VME) board (AMBSLP, where SLP means Serial Link Processor) on which are mounted four local associative memory boards (LAMBSLP), mezzanines that host up to 16 AMchips each.

Figure 2 shows the PU made of an AUX board and an AMBSLP assembled with a first LAMB prototype (in yellow). A network of high speed serial links characterizes the bus distribution on the AMBSLP: 12 input serial links (in red) that carry the silicon hits from the high frequency P3 connector (in green) to the LAMBS, and 16 output serial links (each blue arrow represents 4 links) that carry the road addresses from the LAMBS to P3. The data rate is up to 2 Gb/s on each serial link. Thus the AMBSLP has to handle a challenging data rate: a huge number of silicon hits must be distributed at high rate (24 Gb/s) with very large fan-out to all AM chips (8 million patterns will be located on 64 AM chips on a single AMBSLP) and a similarly large number of roads must be collected and sent back to the AUX (32 Gb/s). There are also diagnostic and test functions; through




Figure 2: The AMBSLP and AUX boards. Data paths are shown to and from the AUX board that sits in the same VME slot. The large yellow square is the mezzanine that holds the AM chips, the miniLAMBSLP.

the VME interface we can spy on the dataflow or simulate the arrival of silicon hits in the input. All these functions are configured in 4 Xilinx FPGAs. They are 2 Xilinx-Artix7 FPGAs which have Gigabit Transceivers (GTP) that provide ultra-fast data transmission (2Gb/s) while 2 other Xilinx-Spartan6 implement the data control logic. The incoming hits are received by the GTPs in the input FPGA (red lines in Figure 2) and saved in large derandomizing FIFOs that are 4k words deep per link. Outgoing road IDs from the LAMBS are sent to output FPGA (blue lines).

The LAMBSLP and the AMBSLP communicate through a high frequency and high pin-count connector placed in the centre of the LAMBSLP. The prototype shown in Figure 2, denoted as miniLAMBSLP, is a simplified version with only four small mini@sic AM chips using the final serialized I/O. The small mini@sics, placed in Quad Flat No-Leads Package (QFN64) packages since they are too small to use the final Ball Grid Array (BGA) package, are visible at the bottom left of the mezzanine. A low-jitter oscillator and a fan-out chip that distributes the 100 MHz clock are at the centre of the mini@sic group. Above the group are the fan-out chips for distribution of the hit serial buses. Two FPGAs are on board, one to program the mini@sics, the other for test purposes.

The new mezzanine has been specifically designed to test the network of serial connections and the compatibility between the FPGAs, the new fan-out chips and the new AM chip I/O. An 8b/10b encoding is used in the serial data stream in order to provide effective error detection, i.e. a 32-bit word is transmitted as 40 bits.

3. Associative Memory chip

3.1 Overview

The AM chips are hardware devices capable of solving the very demanding pattern matching task in real time by comparing all possible hits combinations with a pre-computed set of possible tracks. The AM chip is an evolution of the concept of Content Addressable Memory (CAM). A CAM is a device that implements the inverse function of a Random Access Memory (RAM).




Figure 3: Schematic internal structure of an Associative Memory

A RAM stores data at given addresses and retrieves the corresponding data if queried with the address, a CAM stores data at given addresses and retrieves a list of addresses matching the queried data. CAMs are commonly used in networking devices to implement routing tables or in CPUs to implement translation lookaside buffers. The FTK AM implements the CAM function, but the data is segmented into independent buses and the matching of an address is decided by a majority logic unit over the partial matches of each segment that are stored in flip-flops. This is done since the hits on various detector layers arrive at different times. With this feature the FTK AM is able to find matches between its own content and any combination of the input data, solving a complex combinatorial problem in real time.

In Figure 3 the internal structure of the Associative Memory is shown. Each row represents a pre-calculated trajectory stored in the memory (pattern) in the form of one hit per silicon detector layer. The hits coming from different layers of the tracking detector (up to 8 buses in the current ASIC implementations) are processed in parallel by the AM for pattern recognition. A comparison is made by a CAM cell for each pattern and for each layer. A local memory (“ff” in the Figure 3) latches the result of the comparison. Once a ff is set, it is not reset until the end of event. The pattern is matched if the number of layers matched is above a programmable threshold (majority).

3.2 The AMchip design from CDF to ATLAS

The first device implementing the FTK AM function described in the previous section was developed in the 90s: a Very Large-Scale Integration Application Specific Integrated Circuit (VLSI ASIC) for the Silicon Vertex Tracker (SVT) processor at the CDF experiment at Fermilab [7]. It had a capacity of 128 patterns and was designed in full custom 0.7 μm technology. In 2006 an improved ASIC called AMchip03 was developed for the SVT upgrade. AMchip03 was designed using UMC 180 nm Standard Cells technology; it had a capacity of 5k patterns, 6 input buses and a power consumption of 1.8 W at an operating frequency of 40 MHz.

The requirements for the ATLAS FTK application are more demanding than those for SVT: a bigger silicon detector with higher granularity requires more patterns and more input buses. Higher trigger frequency requires higher operating frequency while the total power consumption must be contained. The AMchip04, the 4th generation of AM devices and the first prototype aimed at FTK application, introduced a mixed architecture: full custom blocks for the CAM cells, standard cell logic for everything else (JTAG, input and output logic, majority logic, priority encoder). The technology chosen for AMchip04 is TSMC 65 nm. The use of full custom CAM cells enabled a higher pattern density with respect to AMchip03 and also the use of advanced techniques to reduce power consumption, more than what was expected from simple node scaling from 180 nm to 65 nm. Another important feature was introduced with the AMchip04 and continued in the newer generation: ternary logic bits. Some bits in the CAM cell can store ternary values (1, 0, don't care) and can be used to achieve a variable resolution pattern. The idea of variable resolution pattern is essential in FTK to have a high efficiency pattern bank without increasing the size of the AM system over the foreseen capacity of one billion patterns [8]. The AMchip04 had the following key characteristics: a) 8192 pattern storage capacity, b) 8 input buses, 15 bits wide inputs (parallel), c) 3 bits with ternary logic (configurable up to 6 bits with 12 bit wide inputs), d) 100 MHz operating frequency.

The AMchip05 will utilize the same TSMC 65 nm technology, but with further improvements in power consumption and new serialized I/O.

3.3 The Serialized I/O

The AMchip03 and AMchip04 used parallel buses for I/O. This led to extreme complexity in the design of the mezzanine boards that host the AM chips (16 AM chips per board). Furthermore for AMchip05 it is foreseen to use different power domains (0.8 V or 1.0 V for the AM core, 1.0 V or 1.2 V for the standard cells, 1.2 V and 2.5 V for I/O) increasing again the routing complexity of the board. In order to solve this board routing issue and be able to produce a reliable and relatively simple mezzanine board, we decided to switch from parallel buses to high speed serial buses. The package of the AM chip also changed from Thin Quad Flat Package (TQFP208) to 23x23 mm²BGA in order to have many pins for the many power domains and a small number of pins for the serial I/O. The main features required for the AM chip serial links (SERDES) are: a) data rate at least 2 Gbps to match 32-bit words @ 50 MHz, b) separate serializer and deserializer macros (the AM chip has many input buses but one output bus for patterns), c) 32-bit input/output bus, d) driver and receiver circuits compatible with the LVDS standard, e) comma detection and word alignment, f) built-in self-test capabilities for fast debugging, g) low power.

We have bought a SERDES IP by Silicon Creations that satisfy all of our requirements [9].

4. Tests

The mini@sic and the miniLAMBSLP have been produced to test the new IP CORE [9] providing a serialized I/O to our AMchip. This system has been specifically designed at low cost, to test the network of serial connections and the compatibility between the FPGAs, the new fan-out chips and the new AM chip I/O, before producing the final board and AM chip version (much more



Figure 4: Test stands available in our lab. The left one is a VME crate while the right one is a stand-alone test for the LAMB board.




Figure 5: The results of our measures are shown: a) Eye Diagram (top); b) Bathtub Curve (middle); c) Jitter Histogram (bottom);

expensive) to reach the final pattern density (128000 patterns) and final package (a flip-chip 23mm x 23 mm BGA).

Figure 4 shows the two test stands available in our lab. The one shown on the right is a standalone test stand where the miniLAMBSP is tested by an evaluation board [10] that has Ethernet connection. We use the IP bus system [11] to access the board and execute the tests. These tests will be used at production time to validate the LAMBs before integrating them into the more complete tests executed in the VME crate, shown on the left of Figure 4.

The VME crate shown in the figure contains one AMBSLP board assembled with one miniL-

AMBSLP board. We execute VME procedures to configure and test the system, with full compatibility with the ATLAS TDAQ standards. The serial links of the all the FPGAs and mini@sics were tested in both test stands. Signals have been correctly transmitted by the FPGAs to the mini@sics, correctly received by the FPGA on the evaluation board and the AMBSLP through the fan-out buffers that are used to replicate and refresh the signals. We used a PRBS (pseudorandom binary sequence) generator activated inside the mini@sics and the FPGAs; for the latter we used a Xilinx macro named Ibert. We also transmitted control and data words, which were correctly received. In our serial link performance analysis we measured the Eye diagram and the BER (bit error rate) parameter. Finally we performed the jitter analysis . Results, shown in Figure 5, were very good.

5. Conclusion

The miniLAMBSLP board and the mini@asic AM chip were built in order to test the performance of the gigabit transceiver and serial I/O links on the AM system. We successfully tested the new AM system, integrating the new boards (AMBSLP and miniLAMBSLP) and the mini@sic chip.

Acknowledgments

The Fast Tracker project receives support from Istituto Nazionale di Fisica Nucleare; the US National Science Foundation and Department of Energy; Grant-in-Aid for Scientific Research from the Japan Society for the Promotion of Science and MEXT, Japan; the Bundesministerium für Bildung und Forschung, FRG; the Swiss National Science Foundation; and the European community FP7 People grant FTK 324318 FP7-PEOPLE-2012-IAPP.

References

- [1] W. Smith, “Triggering at LHC Experiments”, 2002 Nucl. Instr. and Meth. A, vol. 478, pp. 62–67.
- [2] ATLAS Collaboration, “The ATLAS Experiment at the CERN Large Hadron Collider”, JINST, vol.3, pp.437, 2008.
- [3] ATLAS Collaboration, “Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics”, arXiv:0901.0512 [hep-ex], pp.549, 2008.
- [4] ATLAS Collaboration, Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System, CERN-LHCC-2013-018, ATLAS-TDR-023. <https://cds.cern.ch/record/1602235/>
- [5] ATLAS Collaboration, Fast TracKer (FTK) Technical Design Report, CERN-LHCC-2013-007, ATLAS-TDR-021. <http://cds.cern.ch/record/1552953?ln=it>
- [6] A. Andreani et al., “The AMchip04 and the processing unit prototype for the FastTracker”, JINST, vol. 7, C08007, 2012.
- [7] S. R. Amendolia et al., “The AMchip: a full-custom CMOS VLSI associative memory for pattern recognition,” Nuclear Science, IEEE Transactions on, vol. 39, no. 4, pp. 795–797, 1992.
- [8] A. Annovi et al., “A new variable-resolution associative memory for high energy physics,” in Advancements in Nuclear Instrumentation Measurement Methods and their Applications (ANIMMA), 2011 2nd International Conference on, pp. 1–6.

- [9] Silicon Creation, “Serializer - Deserializer IP.” [Online]. Available:
<http://www.siliconcr.com/products/serdes-interfaces>
- [10] http://hitechglobal.com/Boards/Virtex6_PCIExpress_Board.htm
- [11] <https://svnweb.cern.ch/trac/cactus>