

# L1 track trigger for the CMS HL-LHC upgrade using AM chips and FPGAs

Giacomo Fedi<sup>1,a</sup>

<sup>1</sup> *Università di Pisa and INFN Pisa, Italy*

## Abstract.

The increase of luminosity at the HL-LHC will require the introduction of tracker information in CMS's Level-1 trigger system to maintain an acceptable trigger rate when selecting interesting events, despite the order of magnitude increase in minimum bias interactions. To meet the latency requirements, dedicated hardware has to be used. This paper presents the results of tests of a prototype system (pattern recognition mezzanine) as core of pattern recognition and track fitting for the CMS experiment, combining the power of both associative memory custom ASICs and modern Field Programmable Gate Array (FPGA) devices. The mezzanine uses the latest available associative memory devices (AM06) and the most modern Xilinx Ultrascale FPGAs. The results of the test for a complete tower comprising about 0.5 million patterns is presented, using as simulated input events traversing the upgraded CMS detector. The paper shows the performance of the pattern matching, track finding and track fitting, along with the latency and processing time needed. The  $p_T$  resolution over  $p_T$  of the muons measured using the reconstruction algorithm is at the order of 1% in the range 3-100 GeV/c.

## 1 Introduction

For the High Luminosity (HL) upgrade of the Large Hadron Collider (LHC), scheduled for 2026, the number of proton-proton interactions per bunch crossing (pile-up) is predicted to reach 140 at the design instantaneous luminosity of  $5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ . With an increase of  $\approx 5$  times of the pile-up count with respect to the current LHC operations, the CMS trigger system needs to be redesigned to improve its performance.

CMS's trigger system is divided in two stages. The first stage, called level-1 or L1 trigger, decides whether or not to save the information coming from the detector within few microseconds after the bunch crossing (event). This system is mainly implemented using hardware and firmware logic. The second step is called the High Level Trigger (HLT), and makes a further decision on the level-1 accepted events. While the HLT system can use all the available event information, the L1 trigger makes its decision based on partial and imprecise information.

Currently, the L1 trigger does not use the tracker information in its decisions. A way to mitigate the increase of pileup could come from the use of high-resolution spatial information from the silicon tracking detectors, at the cost of adding a few microseconds of latency and with a large increase in

---

<sup>a</sup>e-mail: giacomo.fedi@cern.ch

the amount of data that needs to be processed compared to the current L1-trigger. The trajectories of charged particles in CMS are helices since they are bent by the magnetic field generated by a solenoid. When a charged particle passes through the silicon tracker it deposits a small amount of energy in the silicon sensors and the passage of the particle is recorded (“hit”). The CMS tracker for the HL-LHC upgrade will be composed of six (five) layers of silicon modules in the central (forward) region. Each proton-proton interaction and their interaction with the detectors can generate hundreds of charged particles so that the silicon tracker generates thousands of hits in each event. The large number of hits makes the reconstruction of the trajectory of charged particles in few microseconds very challenging.

For the HL phase of the LHC, the outer tracker of CMS will be redesigned, adopting a novel sensor technology to reduce the event information that is needed to be transferred from the detector to the data acquisition system. Each outer tracker layer will be composed of modules called “ $p_T$  modules”, and each module is made of pairs of silicon sensors. The  $p_T$  module concept relies on the fact that the strips of both sensors are parallel to the beam axis in the barrel and nearly radial in the endcap, and uses the correlation of signals in closely-spaced sensors (a few millimeters) to make an angular measurement of the pair of hits. These pairs of hits are called “stubs”. Using the stub angular information and assuming the tracks are coming from the interaction volume, the  $p_T$  of the track that is generating the stub can be measured. By imposing a  $p_T > 2$  GeV/c requirement, a reduction by  $O(10)$  in the amount of data that needs to be transmitted to the trigger can be achieved.

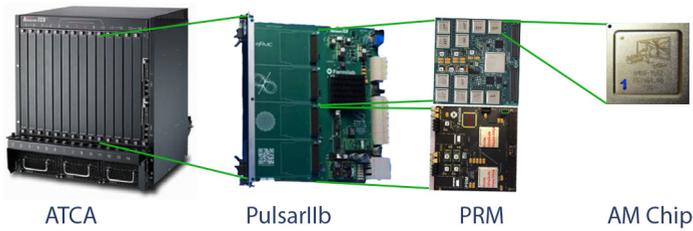
CMS is exploring three different approaches to tackle the track reconstruction at L1 trigger. Two are based on pattern matching and track reconstruction done with FPGAs [1, 2]. The third, to be discussed in this report, is based on the use of an associative memory (AM) ASIC for the pattern matching and FPGAs [3] for the track reconstruction. The purpose of our R&D activity is to prove that CMS’s tracking requirements can be met by combining pattern recognition based on AM custom chips and computing capabilities available in state-of-the-art FPGA devices. Track finding based on Associative Memories [4] has been successfully used in the CDF experiment [5] and, more recently, is being exploited in the Fast-Track processor [6, 7] for the ATLAS Level-2 trigger system.

In the AM plus FPGA approach, the CMS outer tracker will be divided in  $6 \times 8$  regions in  $\eta \times \phi$  (pseudo-rapidity and azimuthal angle) called “trigger towers”. In each trigger tower, the track finding will be performed using data from  $p_T$  modules belonging to that region. Each tower will receive an average of  $\approx 100$  hits per layer for each bunch crossing. We built a prototype of the system with currently available technology to demonstrate that we will be able to build the full system for the HL phase of the LHC. Our demonstration requires only a small extrapolation of the ASIC technology and cost of commercial devices, and depends on the assumption of a strong R&D program on the AM chip.

## 2 Demonstrator System Overview

A small prototype (demonstrator) of the system has been built to prove that the AM plus FPGA approach can satisfy CMS’ HL-LHC requirements. The demonstrator is composed of two ATCA shelves; one ATCA shelf (data-source ATCA) is emulating the data output of the CMS outer tracker, while the second ATCA shelf receives the data and is equipped with the ATCA boards and mezzanine boards that reconstruct the tracks. The two ATCA shelves are physically connected by 480 10 Gbps fiber-optic links, leading to a maximum data rate of 4.8 Tbps.

Both ATCA shelves are equipped with several instances of ATCA custom boards, called Pulsar IIB [8]. These boards were designed at FNAL for data delivery in high energy physics experiments. Two types of PRM have been developed: one with a full AM chip configuration and one with two Xilinx Ultrascale FPGAs and a socket which can accommodate VIPRAM [9] chips. Each Pulsar IIB



**Figure 1.** Picture showing a part of the system. An ATCA shelf accommodates ten Pulsar IIB boards. Each board can accommodate two PRMs.

board can accommodate two Pattern Recognition Mezzanine (PRM) boards; these boards perform the pattern recognition and track fitting. Figure 1 shows a schematic view of the track reconstruction system based on ATCA technology.

Each PRM board contains one or two Xilinx FPGAs, implementing data flow management and track fitting, and one or multiple AM devices. Each PRM board receives full resolution data and temporarily stores them in a data-buffer. Full resolution data are the coordinates of the lower sensor cluster and the angle information. Lower resolution data, called SuperStrip ID (SSID), are evaluated and transmitted to custom chips where the AMs are implemented (AM chip). The AM chip acts as a Content Addressable Memory (CAM) and quickly matches the SSIDs with the pre-loaded patterns. Pre-loaded patterns are based on simulation of the charged particles originating from the proton-proton collision and traversing the tracker. The tracks mostly originate in a well-defined region of the detector and the curvature of the helices is restricted to a certain region of interest.  $p_T$  modules are able to discriminate stubs associated with track  $p_T$  above 2 GeV/c, but in the current implementation of the prototype,  $p_T > 3$  GeV/c threshold has been explored instead, in order to simplify the prototype.

All the patterns present in the AM chip are simultaneously compared with the input data for each layer. When a match is found on a layer, a match bit is set and it remains set until the end of the event processing, when an internal reset signal is propagated. If in a given pattern a match is detected for all the layers, or for a majority of them, the pattern address is transmitted to the FPGA. In the current implementation only one missing layer is allowed. The full resolution data associated with these SSIDs are retrieved from the FPGA data-buffer, filtered and propagated to the Track Fitting module which fits the set of stubs with the same majority set in the AM chips. The processing time for each event in a PRM board is of the order of few  $\mu\text{s}$ , while each PRM board receives events with a period of 500 ns. Therefore, it is necessary to time-multiplex the data to 10 Pulsar IIB boards using a round-robin mechanism, exploiting the full mesh backplane of the ATCA shelf. Tracks that are found in the PRM will then be transferred to the Global Trigger System [10]. The final goal of the prototype board is to evaluate the performance of the real-time system described above using state-of-the-art AM chips. The prototype is also used to extrapolates the future needs, in terms of system requirements, to match the processing time constraint of 4-5  $\mu\text{s}$ .

For this purpose, two prototype processors have been developed: the Pattern Recognition Mezzanine PRM06 at INFN and the FNAL PRM. The processors are designed to cope with the pattern recognition task [11] using the new version of the AM device. The track parameters are then measured for each pattern using a precision track fit [7] performed with the latest generation of FPGA devices.



**Figure 2.** Left: INFN pattern recognition mezzanine board. Right: FNAL pattern recognition mezzanine board.

### 3 PRM architecture

The PRM board is shown in Figure 2. It is a custom  $14.9 \times 14.9$  cm<sup>2</sup> printed circuit board housing the following components:

- FMC connectors (2x);
- Xilinx Kintex Ultrascale FPGA (1x INFN PRM, 2x FNAL PRM);
- Associative Memory devices (12x INFN PRM, 1x FNAL PRM);
  - FNAL AM: FPGA emulated chip with 4k patterns;
  - AM06: 65 nm custom chips with 128k patterns each;
- Power regulators; and
- DDR Pattern Memory (2x INFN PRM, 1x FNAL PRM).

The two FMC connectors are used to interface the PRM to the Pulsar IIb host board. They carry signals and power supply voltages. The host board provides 3.3 V and 12 V, for a maximum available power of 150 W for each PRM. The power regulator group is used to generate the 1.0 V, 1.2 V, 1.8 V and 2.5 V voltages required by the FPGA, the AM devices and the pattern memory. Six bidirectional high-speed serial links (up to 10 Gbps, limited by the FMC connectors), with three links through each FMC connector, are used to send and retrieve data from the PRM with a total input and output bandwidth of 60 Gbps. Moreover, 68 additional LVDS pair pins are used for monitoring and to provide additional bandwidth for data transfer.

The Xilinx Kintex Ultrascale FPGA manages the data flow inside the mezzanine. It receives and temporarily stores the full resolution hit data from the Pulsar IIb board and evaluates and distributes the lower-resolution hit data used in the pattern recognition to the AM chips. Finally, the external pattern memory provides up to 576 Mbit for INFN PRM and up to 36 Mbit for FNAL PRM of memory resources to store a copy of the pattern banks of the AM chips.

The FNAL PRM is equipped with two Xilinx Ultrascale 040/060 FPGAs. Fast serial links connect the two FPGA to each other with a maximum data rate of 16.3 Gbps. One of the FPGA is also fully connected with six high-speed serial links to the FMC connectors. The PRM also has additional external communication links consisting of QSFP+ transceivers.



**Figure 3.** Picture showing an ATCA shelf containing a Pulsar IIb board (leftmost slot) equipped with a PRM06, and an ATCA switch board (center slot).

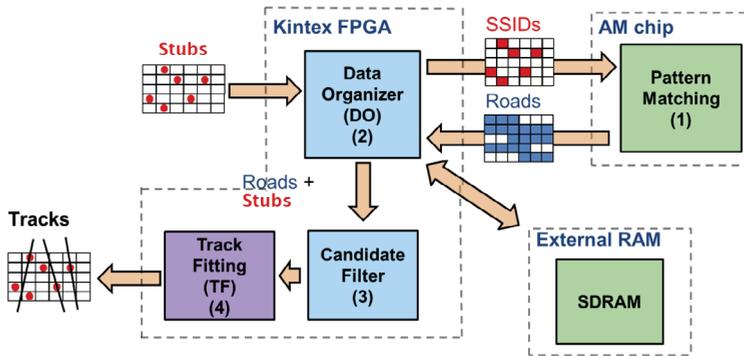
The INFN PRM board houses 12 AM custom chips version 06 (AM06). Each chip can store up to 128k patterns, corresponding to a total capacity of 1.5M patterns per PRM. Each pattern is made of 8 independent 16-bit words, one for each layer. The input data bus distribution from the FPGA to the AM chips uses stages of fan-out buffers to distribute data in parallel to all the AM chips. SSIDs can be fed to the AM chips using two independent input buses. Each input bus is connected to a set of six AM chips. The input serial links can sustain a data rate of 2 Gbps, while the output data of each AM chip is connected directly to the FPGA using high speed serial lines running at 2.4 Gbps. The direct connection between the FPGA and the AM chips has been chosen to eliminate daisy chains and to reduce as much as possible the latency of the Level-1 trigger decision and to integrate all the functionalities in one single board.

The FNAL PRM board has a socket which can accommodate the soon-to-arrive VIPRAM [9] chip. One of the FPGAs is used to emulate the AM chip using synthesized HDL code, while the other FPGA is used to store the hits and to fit the matched patterns. The emulated AM chip can store up to 4k patterns. This mezzanine cannot manage the full number of patterns required for a trigger tower, but is designed to minimize the latency of the pattern-finding process. To this end, the FNAL PRM is equipped with a powerful FPGA with high speed serial links with data rates up to 16.3 Gbps, and the synthesized FPGA code has been designed to minimize latency. The aim of this PRM is to show what we can achieve with the future AM custom chips.

## 4 PRM hardware validation

Both INFN and FNAL PRM boards were validated regarding the electrical circuitry and both showed the expected behavior. The INFN PRM board was tested with a test stand consisting of a PRM06 board, an adapter card, and an evaluation board. An HTG Virtex 6 evaluation board was used for the main tests reported in this text, but electrical and logic connections were also tested using a Pulsar IIb board.

The HTG evaluation board has limited high-speed lanes to the FMC connectors. Only four high speed serial links are connected and the data rate is limited to 6.2 Gbps. To test the fast serial links of the INFN PRM, an Ultrascale Evaluation board was used. The evaluation board is connected to a host PC via Ethernet. It provides power, LVDS parallel connections and high-speed serial connections to



**Figure 4.** Schematic view of the modules of the track finding algorithm.

the PRM board. The IP-bus protocol is used to access devices in the Virtex 6 and in the PRM FPGAs (e.g.: control and status registers, source FIFOs containing test hits, monitoring FIFOs containing PRM outputs).

All the serial links between the FPGAs and to/from the AM chips on the PRM board were successfully tested. JTAG connections to the AM chips (2 JTAG chains with 4 AM chips each, and 2 JTAG chains with 2 AM chips each) have been tested to verify the ability to configure and program the AM chips. The eight input links to each AM chip were characterized. Pseudo Random Bit Sequences (PRBS) were used to test the links between the FPGA and the AM chips. The FMC connector links have undergone tests by using a loopback card and the IBERT tool, provided by Xilinx, that allows measuring the Bit Error Ratio (BER) and producing the corresponding eye diagrams.

Signal integrity up to 12.5 Gbps was tested by measuring the BER and the corresponding eye-scans in the receiver links of the PRM and in the receiver links of the Ultrascale board. The BER has been measured using a PRBS-7 sequence. The measured BER was less than  $1 \times 10^{-15}$  in both directions. The power consumption of the board has been measured by powering and configuring the four groups of the AM chips, resulting in a static power consumption of about 40 W. The mezzanine board has been successfully integrated with the Pulsar IIb board; see Figure 3. During these tests, the mechanical and electrical compatibility have been verified and both the LVDS and high-speed differential pairs were fully tested up to 10 Gbps. The two external DDR memories (RLD3RAM) were successfully tested using a dedicated Xilinx IP core.

The FNAL PRM board has been tested in a standalone configuration and with the Pulsar IIb board. The high speed serial links between the two FPGAs and towards the FMC connectors have been tested; the links between the FPGAs were validated at a rate of 16.3 Gbps, while towards the FMC connectors they were validated at a rate of 10 Gbps. The LVDS links between the FPGAs and the FMC connectors were successfully validated. The QSFP+ transceivers were also tested and no errors were found in the transmitted pseudo-data.

## 5 PRM functionalities

On both INFN and FNAL PRM boards, a synthesized HDL code has been developed to implement functionalities needed for the track finding: hit buffering, hit retrieving, and hit fitting. The tracking algorithm implemented in both PRM boards consists of the following sequence of operations (see Figure 4):

1. Full resolution stub data received by the host board are decoded and stored in a smart data buffer, called Data Organizer (DO), and coarse resolution stub data, SSIDs, are generated and transmitted to the AM chips;
2. AM chips perform pattern recognition on SSID input data and identify sets of SSID matching patterns, called “Road”, that are transmitted back to the FPGA;
3. Roads are used to retrieve the associated SSID from an external pattern memory;
4. Stub data stored in the Data Organizer and belonging to the SSID retrieved from the pattern memory are propagated to the filtering module or to the combination builders;
5. The Track Fitter performs the fitting algorithm using the full resolution data from the Data Organizer to evaluate the track helix parameters and the goodness of the fit ( $\chi^2$ ).

Two different track finding algorithms were developed. Each one was designed to fulfill a different goal. In one case the algorithm focused on the reduction of the track finding latency; it is hosted by the FNAL PRM board. In the other case the algorithm focused on the handling of large pattern banks and multiple AM devices; it is hosted by the PRM06. The full chains of track finding have been developed and tested using HDL simulations and the actual FPGAs for simulated events.

The PRM06 reconstruction algorithm was intended to demonstrate the validity of the PRM design with multiple AM chips working in parallel. The PRM06 boards were equipped with the AM06 chip, which was not developed for stringent latency requirements. The latency of the pattern matching due to the AM06, about 1  $\mu$ s, exceeds the 500 ns time period which needs to be met to manage the incoming data with the time multiplexing period of 20 of the proposed system. Although the AM06 does not meet the latency requirements, the number of patterns per chip are already enough to fulfill the CMS track trigger requirements. The PRM06 reconstruction algorithm was developed to show that a single mezzanine can deal with the number of patterns needed in a trigger tower. Indeed, the number of patterns needed to cover a trigger tower with a satisfying track reconstruction efficiency were evaluated to be in a range between 0.5M and 1M depending on the angular ( $\eta$ ) region of the trigger tower.

Another specific feature of the PRM06 reconstruction algorithm is the way it filters the stubs that are retrieved from the DO after the pattern matching phase. The filtering phase, after the stub retrieval from the DO, is a necessary phase to clean out spurious stubs from the obtained stub set. This is due to the fact that coarse resolution stub data (SSIDs) are used to match the AM chip patterns and multiple stubs can end up creating the same SSID. Once the stubs related to the SSIDs of a pattern are retrieved from the DO, more than one stub per layer might be found. Since the track fitter algorithm accepts only track candidates with a single stub per layer, a filter stage is implemented with a pattern finding algorithm, which uses combinations of pairs of stubs coming from the first three layers of the outer tracker as seeds to find the other track-compatible stubs. The filtering module outputs only the best combination of stubs, fulfilling the requirements of a single stub per layer and only one possible missing stub in one layer.

The track reconstruction algorithm of the PRM06 board was checked by comparing the output of the test stand with the HDL simulation and with the C++ emulation of the system. Various simulated event samples were used to test the board and matching AM06 pattern recognition output was obtained when comparing the emulation and HDL simulation, demonstrating the correct performance of the PRM06 board.

The FNAL PRM reconstruction algorithm was intended to measure the latency of the track reconstruction algorithm, emulating the operational behavior of the future AM chip and testing the

algorithms with the current available FPGAs. To this end, all track reconstruction steps shown in Figure 4 were optimized to reduce the latency to a minimum and to foster pipelines in the system.

Apart from the different goals of the FNAL PRM and PRM06, the main differences in the track reconstruction algorithm are the different ways of dealing with the spurious stubs after the stub retrieval phase from the DO and the track fitting method precision. The FNAL PRM algorithm does not filter the spurious stubs out, instead implementing an algorithm that generates all possible permutations of one stub per layer and all possible permutations of one stub per layer with one stub missing in one layer. This method has low latency when there are only a few stubs per layer, but it is suboptimal in case of large number of stubs per layer: in this case, the number of permutations can be very large, increasing the latency of the system. To prevent the number of permutations from affecting the overall performance of the system, a truncation on the number of permutations considered is set to 480. The truncation limit has never been met with any simulated event samples considered (even in the worst case:  $t\bar{t}$  events with pileup of 250). In both FNAL PRM and PRM06, a linearized  $\chi^2$  track fitter (PCA) was implemented with HDL code in the FPGA. The FNAL PRM track fitter was designed to give more precise parameters by adding higher-order corrections to the linearized fitting matrices.

## 6 Demonstrator performance

Using simulated event samples, we tested the prototype by splitting the total time of processing and latencies into three main logical blocks. The first logical block is the data delivery latency. It is the time that is needed to bring the stubs of an event from the outer tracker front-end electronics to the mezzanine interface of the Pulsar IIb. It was measured in hardware using the two-ATCA shelf prototype and it also accounts for the latency needed for the data formatting and data delivery (serialization and deserialization latency is included). The second logical block is the pattern recognition latency; it includes the latency due to the AM chip pattern recognition, the latency due to the retrieving of pattern SSIDs from the memory, and the retrieving of the stubs per each SSID from the DO. The third logical block is the track fitting latency; it includes the latency of the permutation builder and the latency of the PCA track fitter.

For the latency and total time of reconstruction measurements we used the FNAL mezzanine that was developed for this purpose. All the HDL algorithms implemented in the FPGA were running at the frequency of 240 MHz. The FPGA emulating the AM chips can contain only about 4k patterns, but they were enough to test single events, since the number of matching patterns is one order of magnitude less. In the test configuration, the FNAL mezzanine could serve one fourth of the trigger tower, since there was only one emulated AM chip, while in the future mezzanine, at least four chips are expected. To optimize the pipelining of the algorithms three instances of permutation builder and PCA track fitter were implemented.

The three logical blocks in which the latency is split are fully pipelined. Thus, when the first data exits one block it is immediately computed by the next block. In Table 1, we report the time measurements performed with the demonstrator. The table reports the starting and ending time for each logical block. The total time for the track reconstruction is estimated to be  $2.53 \mu\text{s}$ .

The track parameter performance of the system is generally good. Muons are reconstructed with an efficiency greater than 95% for  $p_T$  greater than 3 GeV/c. Good performance is obtained also for pions. The track parameter resolutions are excellent; it is worth mentioning that the  $p_T$  resolution over  $p_T$  of the muons is at the order of 1%.

**Table 1.** Measured time for each logical block of the demonstrator.

Logical block	Starting time ( $\mu\text{s}$ )	Ending time ( $\mu\text{s}$ )
Data delivery to the PRM	1.2	1.7
Pattern recognition	1.85	2.3
Track fitting	2.03	2.53

## 7 Conclusions

A low latency track reconstruction demonstrator for the L1 trigger of CMS for the HL phase of the LHC has been built. In this approach, we made use of custom AM chips and commercial FPGAs. The demonstrator was composed of two ATCA shelves in order to demonstrate that large amounts of data can be delivered to the pattern reconstruction mezzanines. Two PRMs were tested. One shows that with the available custom AM chips, we can build a working mezzanine that can deal with the number of patterns needed for a trigger tower. The other mezzanine was used to emulate the future AM chip features in order to measure the latency of the pattern recognition and track fitting stages. It was demonstrated that extrapolating the current technology to a reasonably more performant technology, the system can reconstruct the tracks in less than 4  $\mu\text{s}$  with very good track parameter resolution and efficiency.

## Acknowledgments

This work is partially funded by MIUR Italy, under contract 2012Z23ERZ of PRIN 2012 H-TEAM and the People Program (Marie Curie Actions) of the European Union's Seventh Framework Program FP7/2007-2013/ under REA grant agreement no. 317446 INFIERI "INtelligent Fast Interconnected and Efficient Devices for Frontier Exploitation in Research and Industry".

## References

- [1] G. Hall [CMS TMTT], Nucl. Instrum. Meth. A **824** (2016) 292. doi:10.1016/j.nima.2015.09.075
- [2] L. Skinnari [CMS Collaboration], EPJ Web Conf. **127** (2016) 00017. doi:10.1051/epjconf/201612700017
- [3] F. Palla *et al.* [CMS Collaboration], JINST **11** (2016) no.03, C03011. doi:10.1088/1748-0221/11/03/C03011
- [4] M. Dell'Orso and L. Ristori, Nucl. Instrum. Meth. A **278** (1989) 436. doi:10.1016/0168-9002(89)90862-0
- [5] J. Adelman *et al.* [CDF Collaboration], "The Silicon Vertex Trigger upgrade at CDF," Nucl. Instrum. Meth. A **572** (2007) 361.
- [6] F. Alberti *et al.*, JINST **8** (2013) C01040. doi:10.1088/1748-0221/8/01/C01040
- [7] C. Gentsos, F. Crescioli, P. Giannetti, D. Magalotti and S. Nikolaidis, "Future evolution of the Fast TracKer (FTK) processing unit," PoS TIPP **2014** (2014) 209.
- [8] J. Olsen, T. Liu and Y. Okumura, JINST **9** (2014), C01041. doi:10.1088/1748-0221/9/01/C01041
- [9] T. Liu, G. Deptuch, J. Hoff, S. Jindariani, S. Joshi, J. Olsen, N. Tran and M. Trimpl, JINST **10** (2015), C02029. doi:10.1088/1748-0221/10/02/C02029
- [10] M. Jeitler *et al.*, JINST **2** (2007) P01006. doi:10.1088/1748-0221/2/01/P01006

- [11] A. Andreani, A. Annovi, R. Beccherle, M. Beretta, N. Biesuz, W. Billereau, R. Cipriani, S. Citraro *et al.*, “The Associative Memory Serial Link Processor for the Fast TracKer (FTK) at ATLAS,” *JINST* **9** (2014) 11, C11006.