

PAPER • OPEN ACCESS

Intel Stratix 10 FPGA design for track reconstruction for the ATLAS experiment at the HL-LHC

To cite this article: A. Camplani *et al* 2023 *JINST* **18** P06029

View the [article online](#) for updates and enhancements.

You may also like

- [Stability of Corrosion-Resistant Magnéli-Phase \$\text{Ti}_n\text{O}_{2n-1}\$ -Supported PEMFC Catalysts at High Potentials](#)
Tsutomu Ioroi, Hiroshi Senoh, Shin-ichi Yamazaki *et al.*
- [Analysis of structural and electronic properties of \$\text{Pr}_2\text{NiO}_4\$ through first-principles calculations](#)
S M Aspera, M Sakaue, T D K Wungu *et al.*
- [Tailorable Negative Permittivity of Carbon Materials Derived from Microcrystalline Cellulose at Different Carbonizing Temperature](#)
Yingjie Wang, Huan Ren, Chuanbing Cheng *et al.*



PRIME
PACIFIC RIM MEETING
ON ELECTROCHEMICAL
AND SOLID STATE SCIENCE

HONOLULU, HI
Oct 6–11, 2024

Abstract submission deadline:
April 12, 2024

Learn more and submit!



Joint Meeting of

The Electrochemical Society
•
The Electrochemical Society of Japan
•
Korea Electrochemical Society



Intel Stratix 10 FPGA design for track reconstruction for the ATLAS experiment at the HL-LHC

A. Camplani,^{a,*} S. Dittmeier,^{b,*} A. Annovi,^c K. Axiotis,^d R. Beccherle,^c N. Biesuz,^c R. Brenner,^e S. Débieux,^d M. Ellert,^e P. Francavilla,^c P. Giannetti,^c K. Kordas,^f M. Mårtensson,^e P. Mastrandrea,^c C. Noulas,^f J. Oechsle,^a M. Piendibene,^c R. Poggi,^d A. Schöning,^b A. Sfyra,^d C.L. Sotiropoulou,^{c,1} J. Steentoft,^e T. Tsiakiris,^f S. Xella^a and J. Zinßer^b

^aNiels Bohr Institute (NBI), University of Copenhagen,
Blegdamsvej 17, 2100 Copenhagen, Denmark

^bPhysikalisches Institut, Ruprecht-Karls-Universität Heidelberg,
Im Neuenheimer Feld 226, 69120 Heidelberg, Germany

^cINFN Sezione di Pisa and Dipartimento di Fisica, Università di Pisa,
Largo Bruno Pontecorvo, 3, 56127 Pisa, Italy

^dDépartement de Physique Nucléaire et Corpusculaire, Université de Genève,
Quai Ernest-Ansermet 24, 1205 Genève, Switzerland

^eDepartment of Physics and Astronomy, Uppsala University,
Lägerhyddsvägen 1, 75237 Uppsala, Sweden

^fDepartment of Nuclear and Elementary Particle Physics,
Aristotle University of Thessaloniki,
Faculty of Sciences bldg, A.U.Th. Campus, 54124, Thessaloniki, Greece

E-mail: alessandra.camplani@cern.ch, sebastian.dittmeier@cern.ch

ABSTRACT: The fast reconstruction of charged particle tracks with high efficiency and track quality is an essential part of the online data selection for the ATLAS experiment at the High Luminosity LHC. Dedicated custom designed hardware boards and software simulations have been developed to assess the feasibility of a Hardware Tracking Trigger (HTT) system. The Pattern Recognition Mezzanine (PRM), as part of the HTT system, has been designed to recognize track candidates in silicon detectors with Associative Memory ASICs and to select and reconstruct tracks using linearized algorithms implemented in an Intel Stratix 10 MX FPGA. The highly parallelized FPGA design makes extensive use of the integrated High-Bandwidth-Memory.

In this paper, the FPGA design for the PRM board is presented. Its functionalities have been verified in both simulations and hardware tests on an Intel Stratix 10 MX development kit.

KEYWORDS: Online farms and online filtering; Pattern recognition, cluster finding, calibration and fitting methods; Trigger algorithms; Trigger concepts and systems (hardware and software)

ARXIV EPRINT: [2302.13609](https://arxiv.org/abs/2302.13609)

¹Now with Computer Aided Software Technologies (CAST), Inc.

*Corresponding author.



Contents

1	HL-LHC and the ATLAS experiment	1
1.1	The TDAQ Phase-II upgrade	2
1.2	The Hardware Tracking for the Trigger system	4
2	The Pattern Recognition Mezzanine	5
2.1	Functional description	5
2.2	Requirements	8
3	Block description of firmware for standalone testing	9
3.1	Data Generator	9
3.2	ASIC Emulator	10
3.3	Data Organiser	11
3.4	High-Bandwidth Memory Interfaces	12
3.5	Track Fitter	13
3.6	Monitoring	14
4	Implementation and results	14
4.1	Resource usage	15
4.2	Clock and power studies	15
4.3	Performance studies	16
4.4	Considerations for a final design	18
5	Conclusion	20

1 HL-LHC and the ATLAS experiment

The High Luminosity Large Hadron Collider (HL-LHC) [1] is currently planned to begin operation in 2029. Protons will collide at $\sqrt{s} = 14$ TeV, delivering an instantaneous luminosity of up to $\mathcal{L} = 7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. This implies that the mean number of proton-proton collisions in the same bunch crossing (pile-up) is expected to reach up to $\langle \mu \rangle = 200$, an increase of more than a factor three compared to current conditions. In order to achieve optimal physics performance under these harsh conditions, all LHC experiments will undergo a series of upgrades.

For the HL-LHC era, the ATLAS experiment [2, 3] foresees, among others, two major upgrades: the installation of a completely new tracking detector, the Inner Tracker (ITk), composed exclusively of silicon pixels [4] and strips [5] as depicted in figure 1; and the redesign of the Trigger and Data Acquisition (TDAQ) system [6].

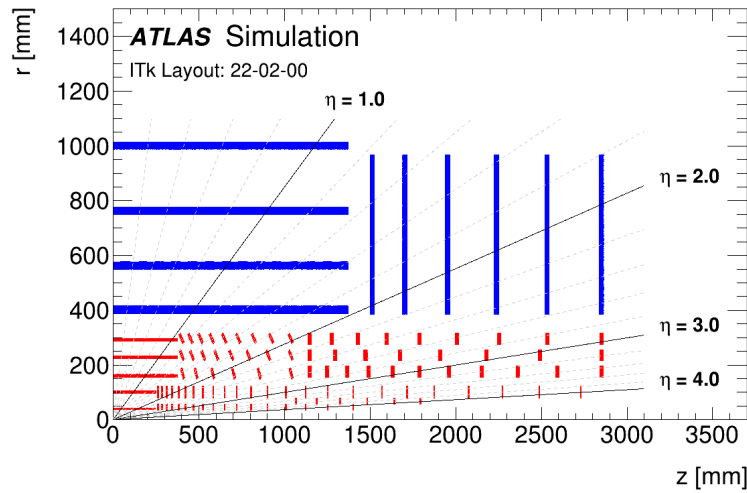


Figure 1. A schematic view of the layout of the ATLAS ITk detector used for this study, with a five-layer Pixel detector (red) surrounded by the Strip detector (blue). Only the positions of the active sensors are shown. The final ITk layout can be found in [7]. Adapted from [7]. CC BY 4.0.

1.1 The TDAQ Phase-II upgrade

The ATLAS detector currently produces more than 60 TB/s of data during operation, but only a small fraction of these collision events contain interesting characteristics that might lead to new physics discoveries. As a consequence, a specialized multi-level system, the TDAQ system, selects the events that are potentially interesting for physics analyses in real-time. The TDAQ role is to reduce the data rates to manageable levels in order to be able to transport the data to the storage system for offline analysis.

At the HL-LHC, the ATLAS TDAQ system will face various challenges, such as a much more complicated event selection as a result of higher pile-up; readout of a tracking detector with more than 5 billion channels (compared to the ~ 100 million channels of the current detector); and a full event size of up to ~ 5.2 MB (compared to the current ~ 2 MB).

In addition, it is paramount for the trigger selection to remain inclusive enough to allow for electroweak scale physics as well as a broad range of physics searches. Therefore, the TDAQ system will undergo a major upgrade, with several operating parameters updated with respect to the LHC era. Among these the most relevant are:

- the hardware trigger latency will be increased from $2.5 \mu\text{s}$ to $10 \mu\text{s}$
- the detector readout rate will be raised from 100 kHz to 1 MHz
- the rate of data sent to permanent storage will be increased from 1 kHz to 10 kHz

Track reconstruction is a crucial element in the trigger selection, as it can provide extremely precise measurements in both position and momentum. This is of particular importance in order to distinguish hard-scattering processes from the abundant low-momentum background particles from concurrent pile-up events.

The trigger system is composed of two levels: the hardware based Level-0 Trigger system, which relies only on data from the calorimeters and muon detectors and is operated with a fixed

latency, and the software based Event Filter system, potentially deploying accelerator hardware, that provides high-level trigger functionality. Both trigger systems are connected to a complex Data Acquisition system. A schematic view can be found in figure 2.

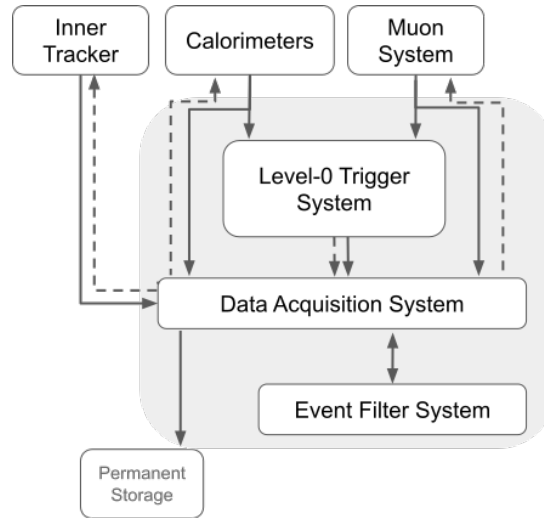


Figure 2. TDAQ system structure for Phase II. The dashed lines represent the Level-0 accept signal, while the solid lines represent the readout data.

The Event Filter design in the Technical Design Report (TDR) [6] includes a CPU-based processing farm complemented by a hardware-based co-processor system, called HTT (Hardware Tracking for the Trigger). The HTT system offers several benefits over a CPU-only farm, e. g. a lower processing latency and a more efficient use of resources. The latter allows for a reduction of the CPU farm size required for online tracking by a factor of 10 [6], which in turn leads to a lower power consumption of the overall system. This TDAQ architecture is referred to as the baseline design for the TDAQ upgrade, as in figure 2.

Given the availability of a hardware based tracking system, the TDR includes a scenario which evolves the trigger into a dual-level hardware trigger architecture. This is done by introducing tracking at Level-1 using a dedicated low-latency version of HTT. This can reduce the L0 rate to less than 1 MHz for further processing in the Event Filter. Implementing hardware tracking at Level-1 enables better control of the trigger rates by adding more margins and more flexibility, not only in the case of major changes of the experimental conditions or the trigger menu, but also as a handle to maintain stability as the processing at high pile-up does not scale linearly with luminosity. Furthermore, it creates opportunities to lower thresholds and improve acceptance on many important physics signatures.

As the HL-LHC upgrade projects progressed, the final choice of the ATLAS experiment was to design a heterogeneous system, based on commodity servers and possibly accelerator boards, excluding the possibility of a system based around custom associative memory ASICs [8]. Nevertheless, the decision process required performance studies, both in simulation and on hardware demonstrators, to assess the feasibility of such a system. In this paper we present the development of the firmware design for one mezzanine board dedicated to the track reconstruction and fitting, called the Pattern Recognition Mezzanine.

1.2 The Hardware Tracking for the Trigger system

The HTT system is designed to reconstruct quickly and efficiently the tracks of charged particles within the ITk. Depending on the Level-0 trigger decision, the ITk detector data can be processed in two different ways in the baseline HTT design: regional or full-scan track reconstruction, according to trigger configuration.

Table 1. Regional and full-scan tracking requirements.

	p_T	Detector coverage	Level-0 rate
Regional tracking	$>2 \text{ GeV}$	$< 10 \%$	1 MHz
Full-scan tracking	$>1 \text{ GeV}$	full coverage	100 kHz

For the regional tracking, as shown in table 1, tracks with transverse momentum of $p_T > 2 \text{ GeV}$ are reconstructed within the so-called Regions Of Interest (ROI), which cover less than 10 % of the detector volume. This allows for a quick initial background rejection for single high- p_T lepton triggers, as well as for multi-object triggers. Regional tracking can be done at the full Level-0 accept rate of 1 MHz.

For the full-scan reconstruction, all tracks with $p_T > 1 \text{ GeV}$ are reconstructed using the full detector data. This is especially important for background rejection of hadronic signatures, to enhance primary vertex identification, and to correct and mitigate pile-up effects. As these requests demand more bandwidth and processing power, they can be served at rates of up to 100 kHz.

For the scenario where the Level-1 Trigger is present, the low-latency version of the HTT system has to provide tracks to the next processing step within a few microseconds. In this case only the regional tracking is used. For the HTT system placed in the Event Filter, both regional and global tracking are available, as in the baseline design.

The requirements posed by the low-latency version are the main driver for the architectural choice of the HTT system. The high throughput and the tight latency budget can only be met using a dedicated hardware system. FPGAs that can sustain the required throughput have been selected to implement the clustering and track fitting algorithms. A custom-designed ASIC has been developed, the Associative Memory (AM) ASIC (described in detail in [9]), following the approach previously used by the CDF experiment [10] and intended for the ATLAS FastTracKer (FTK) project [11]. The AM ASICs are able to deal with the combinatorial complexity in an extremely parallel way, providing a fast recognition of track candidates with high efficiency. In parallel to the developments of the AM ASICs, hardware implementations of alternative pattern recognition methods using Hough transform have also been studied [12]. Related studies of non-AM based track triggers can be found in [13, 14].

The HTT design relies on a single ATCA [15] main board, the Tracking Processor (TP), that can be equipped with two different types of mezzanines: the Pattern Recognition Mezzanine (PRM) or the Track Fitting Mezzanine (TFM). The TP mounts a Xilinx Virtex UltraScale+ VU19P FPGA. The PRM and the TFM both host an Intel Stratix 10 MX FPGA. The AM ASICs are only mounted on the PRM. In the baseline design, see figure 3, 12 PRM boards and 2 TFM boards share one ATCA crate and form one HTT unit. The detector volume is split into $48 \eta \times \phi$ regions ($2.66 \times \pi/8$),

each being individually processed by one HTT unit. The segmentation has been chosen to minimize data sharing between HTT units required to account for regional overlaps.

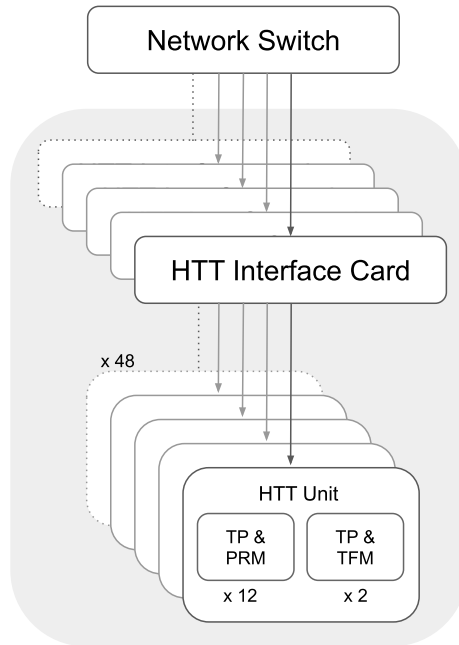


Figure 3. HTT system architecture.

The HTT system receives the detector data via optical links from the network through dedicated servers using the FELIX [16, 17] cards. In the Level-1 case, dedicated low-latency streams are necessary to transfer the data directly from the ITk to the regional tracking units.

The HTT system processes the ITk detector data in steps. The incoming ITk hits are clustered on the TP and clusters are sent to the mezzanines. The PRM reconstructs tracks from 8 ITk layers, while the TFM refines the fit adding the hits from the remaining five layers.

2 The Pattern Recognition Mezzanine

A prototype of the Pattern Recognition Mezzanine (PRM) is shown in figure 4. The PRM PCB is meant to be mounted onto its ATCA motherboard (TP), and it features a powerful FPGA, the Intel Stratix 10 MX,¹ together with 20 Associative Memory ASICs. To implement the power-up sequence of the Stratix 10 as well as for control and monitoring of the board, an Intel MAX 10 FPGA is also installed. The PRM connects to the TP via eight Samtec Z-Ray connectors of type ZA8H-24-0.33-7, that provide power and high speed serial data links, as well as SGMII [18], JTAG [19] and other control interfaces.

2.1 Functional description

The PRM performs the first track reconstruction and selection step for the HTT system. It processes data from 8 detector layers, of which a maximum of four can be silicon pixels and the rest are

¹Part number: 1SM21BHN3F53E3VG.

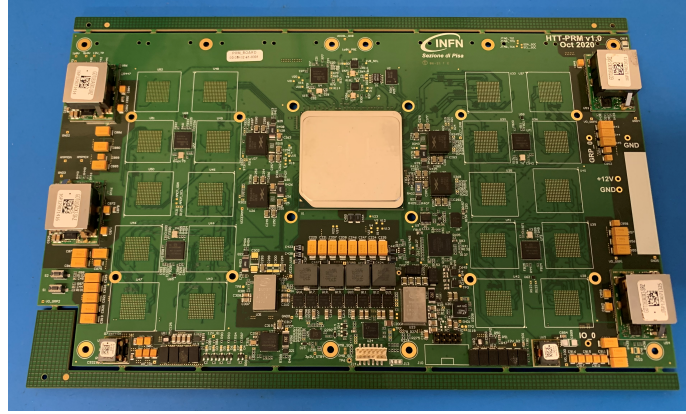


Figure 4. Photograph of a Pattern Recognition Mezzanine board equipped with an Intel Stratix 10 MX FPGA. No AM ASICs are soldered to the board.

silicon strip layers, depending on which ITk detector region is being processed [4, 5]. As in the HTT simulation studies, the configuration used for this study foresees 1 pixel and 7 strip layers. The Intel Stratix 10 MX FPGA forms the core of the PRM, handling the communication with the TP and the ASICs as well as implementing the algorithms used for track reconstruction and selection. The PRM receives clustered ITk detector data via high speed serial links from the TP using the FPGA's transceiver blocks. Up to 32 channels for the input and 16 for the output data have been foreseen.

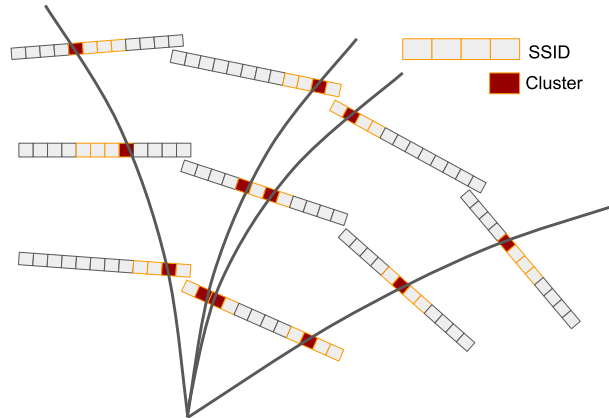


Figure 5. Simplified representation of super-strips identifiers granularity vs. single clusters granularity: one super-strip can contain multiple clusters.

Twenty AM ASICs on the PRM are arranged in four groups of five daisy-chained ASICs, that can process different sets of clusters simultaneously. All 20 ASICs combined can store a total of 7.864 million patterns. A pattern is a sequence of eight so-called Super-Strip Identifiers (SSIDs), which collect adjacent groups of silicon strip or pixel channels into a coarse position identification. Figure 5 shows a simplified representation.²

²HTT simulation studies have used the following SSID sizes: 33x402 for the pixel and 40 for the strips in the barrel, and 16x200 for the pixel and 10 for the strips in the end-caps.

All the patterns are generated from Monte Carlo simulated muons,³ which originate from the interaction region in ATLAS with a flat p_T spectrum for p_T larger than 1–2 GeV [20]. The super-strip width, in both dimensions for the pixels, is a parameter tuned in the simulation to optimize the number of patterns within the AM ASICs and to reduce the bandwidth between the FPGA and ASICs. At the same time, it impacts on the average number of fit combinations, as an increase in the super-strip width would also increase the combinatorics.

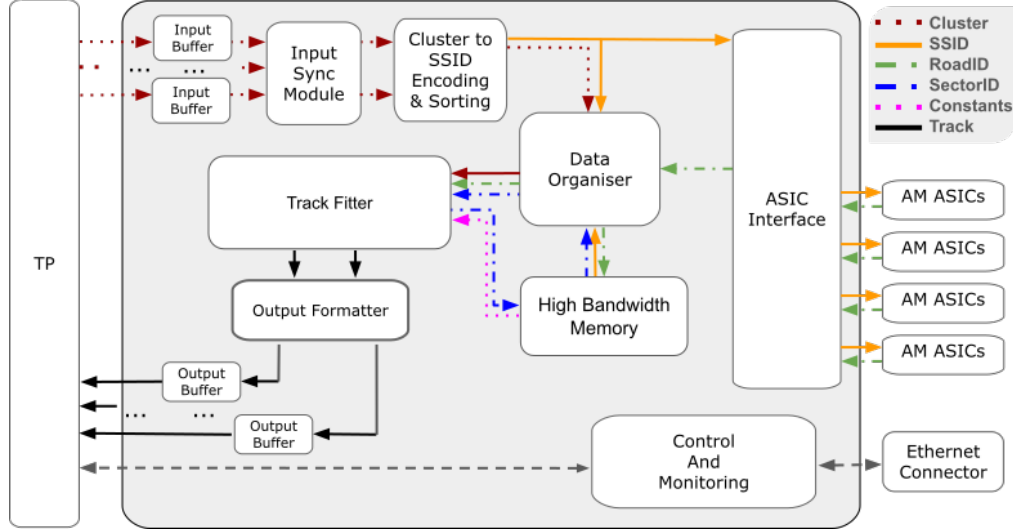


Figure 6. PRM complete block diagram. Solid lines represent the data transmission: in red the clusters, in orange the SSID, in green the roadID, in blue the sectorID, in violet the constants and in black the tracks. Dashed lines represent the firmware control and monitoring.

The complete block diagram of the Intel Stratix 10 MX FPGA design functionalities is shown in figure 6 and explained in the following paragraphs. The data processing is event based, where one event corresponds to the collision of two proton bunches. In the first processing step, the input cluster data (red) coming from the TP is synchronized and converted into SSIDs (orange), which are sent to the AM ASICs for pattern matching. The interface between the FPGA and a group of ASICs consists of source-synchronous parallel LVDS buses operated at 1 Gbps per line, which yields an ASIC input data rate of 32 Gbps and an ASIC output data rate of 8 Gbps.

Within the FPGA, all clusters are stored temporarily in an on-the-fly data base called Data Organiser, grouped under their corresponding SSIDs. When an event processing is being completed in the AM ASICs, the FPGA receives the addresses of all matched patterns, which are referred to as road IDentifiers (roadIDs, green). Using these roadIDs, the SSIDs of all matched patterns are collected in the Data Organizer from the High-Bandwidth-Memory (HBM), where a copy of the pattern bank is stored. The SSIDs and their associated clusters are then processed in the Track Fitter block.

The Track Fitter block implements a Principal Component Analysis (PCA) algorithm to determine the track parameters for a given set of clusters. A specific PCA is only applicable within a

³Efficient tracking of long living particles, i. e. largely displaced tracks, requires additional patterns and constants to be generated. Similarly, tracking efficiency of lower energy electrons can be retained with additional patterns, if no further processing, like a Gaussian Sum Filter, is available.

small detector region, referred to as a sector. The detector is therefore divided into many sectors which are identified by their sectorIDs (blue), stored in the HBM. Within a sector the track equation can then be linearized, according to equation (2.1).

$$p_i = \sum_{j=1}^{N_{\text{coo}}} (C_{ij}x_j + q_i) \quad (2.1)$$

Here, x_j are the coordinates of the clusters and p_i are the 5 expected track parameters, namely pseudorapidity η , azimuthal angle ϕ , transverse momentum p_T , and the impact parameters d_0 and z_0 . C_{ij} and q_i form the set of constants extracted from the software simulation training which are stored in the HBM (violet). The minimization of the distance between the expected and real track parameters in the PCA follows a χ^2 -distribution, allowing the use of the χ^2 -variable to evaluate the track quality.

$$\chi^2 = \sum_{i=1}^{N_{\text{dof}}} \left[\left(\sum_{j=1}^{N_{\text{coo}}} S_{ij}x_j \right) + h_i \right]^2 \quad (2.2)$$

Equation (2.2) contains the constants h_i and S_{ij} , where S_{ij} is the inverted covariance matrix between the coordinates, also stored in the HBM (violet). The number of coordinates is $N_{\text{coo}} = 2N_{\text{pixels}} + N_{\text{strips}}$, with pixels contributing two coordinates and strips only one. N_{dof} is the number of degrees of freedom, $N_{\text{dof}} = N_{\text{coo}} - N_{\text{par}}$, with the number of track parameters $N_{\text{par}} = 5$. It should be noted that the algorithm also allows for tracks with up to two missing clusters, e. g. due to detector inefficiencies. In these cases, the track parameters and χ^2 can be recovered by a minimization with respect to the coordinates of the missing clusters, or by generating additional sets of constants that could be directly applied into equations (2.1) and (2.2).

The computations in the Track Fitter block rely solely on multiply-accumulate (MAC) operations. Hence, they can be efficiently implemented with the FPGA's integrated DSPs. A cut on the χ^2 value is implemented to reject most of the wrong combinations of clusters before computing the track parameters, such that only about 8 % of the tracks are selected, as shown in table 2. This concludes the data processing on the Intel Stratix 10 FPGA. For each track (black), its track parameters, its χ^2 value, and the associated cluster information (represented by the black lines) are sent from the PRM to the TP via high speed serial links.

2.2 Requirements

Table 2 summarizes the requirements for the Pattern Recognition Mezzanine in the TDAQ baseline architecture with HTT as part of the Event Filter System. These requirements are derived from simulations of the expected data flow through the entire HTT system for pile up $\langle \mu \rangle = 200$ [20]. The main parameters that affect the processing speed are the average cluster rate, road processing rate, fit constants access rate, and track fitting rate. In the low-latency scenario, an additional constraint on the processing latency of $O(1 \mu\text{s})$ has to be fulfilled. The HBM storage requirements for a full PRM are about 73 MB to hold the patterns, and about 375 MB for the constants, estimated from simulation studies. Multiple copies of the patterns and constants can be stored in the HBM to fulfil the rate requirements.

Table 2. Main parameters for the PRM and their requirements derived from the HTT system in the TDAQ Phase-II baseline design. The values are related to an expected maximum event request rate of 240 kHz and refer to one of four firmware units working in parallel within the PRM that are connected to separate groups of 5 AM ASICs.

Parameter	Requirements per group of ASICs
Peak cluster rate per layer	95 MHz
Roads processing rate	77 MHz
Constants readout rate (χ^2 calculation)	41 MHz
Fit rate (χ^2 calculation)	714 MHz
Constants readout rate (parameter calculation)	41 MHz
Fit rate (parameter calculation)	57 MHz
Track output rate	57 MHz

3 Block description of firmware for standalone testing

The PRM firmware has been developed and tested in a “standalone” set-up, where the transceivers and the ASIC functionalities have been emulated through a Data Generator and an ASIC Emulator block. The corresponding firmware diagram is shown in figure 7. In the following, the individual blocks are described in more detail in the order of the data flow.

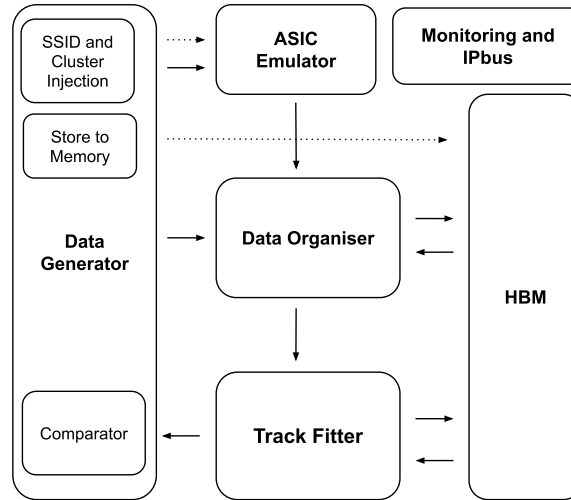


Figure 7. PRM standalone block diagram. The dashed arrows are for paths representing the storage of patterns or constants into memory. The solid line arrows are for data paths.

3.1 Data Generator

The Data Generator block role is to provide different test modes, initialize the PRM internal memories with the necessary constants, inject the data, and check the results. The block diagram is shown in figure 8.

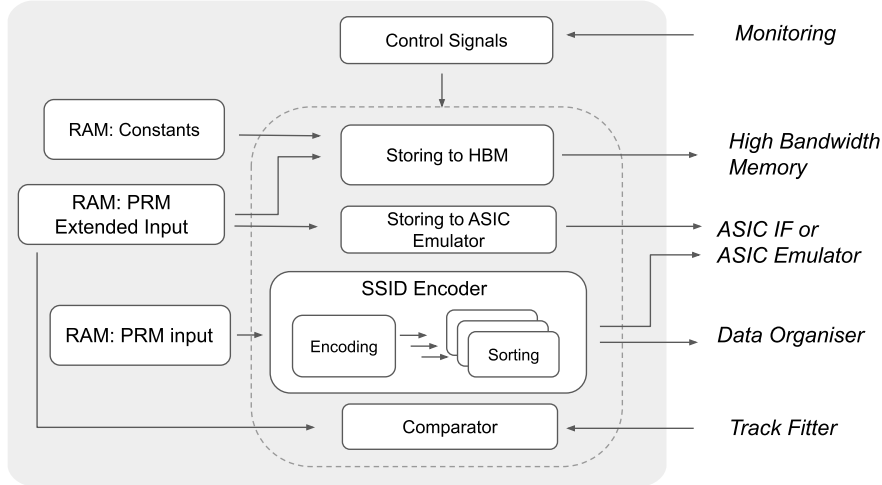


Figure 8. Data Generator block diagram with activated SSID Encoder.

For the first functionality, the Control Signal block handles the different testing modes. Two injection modes have been implemented: with and also without the SSID Encoder block, responsible for the assignment of an SSID to each cluster. The selection between the two modes is done at compilation time. In both modes various parameters of the data injection can be controlled, such as the injection frequency and the number of events. In addition, it is possible to add randomly generated fake clusters and SSIDs into the data stream, to emulate the real detector behavior in the high-pile-up environment.

The PRM internal memories are initialized when a memory initialization command is issued. The fit constants are uploaded from the Constants RAM into the High-Bandwidth Memory, presented in section 3.4, through the Storing to HBM block. The track patterns are also sent from the PRM Extended Input RAM to the ASIC Emulator block, described in section 3.2, and to the HBM via the Storing to ASIC Emulator and HBM blocks.

Following the initialization of the memories, the injection is started. The PRM Input RAM contains the clusters in the same format as expected to be computed in the TP motherboard.

For the results check, the PRM Extended Input RAM contains the simulated true track χ^2 and the helix parameters for an easy comparison that happens in the Comparator block.

3.2 ASIC Emulator

The ASIC Emulator implements the AM pattern matching by storing 16 patterns per instance, each configurable at run-time. It uses the same data and road bus input and output interfaces as the internal core of the AM ASIC [9]. To fully mimic a single group of AM ASICs on the PRM, a daisy chain of five emulator instances is implemented. Figure 9 shows a block diagram of a group of five ASIC Emulators.

The input bus is designed to supply the ASIC Emulator with three types of data: cluster data, control commands, and the idle character. The Input Decoder distributes cluster data to the Core block and commands to the Controls block. The Core block stores the patterns and implements the pattern matching functionality. The Controls block decodes the commands and implements the

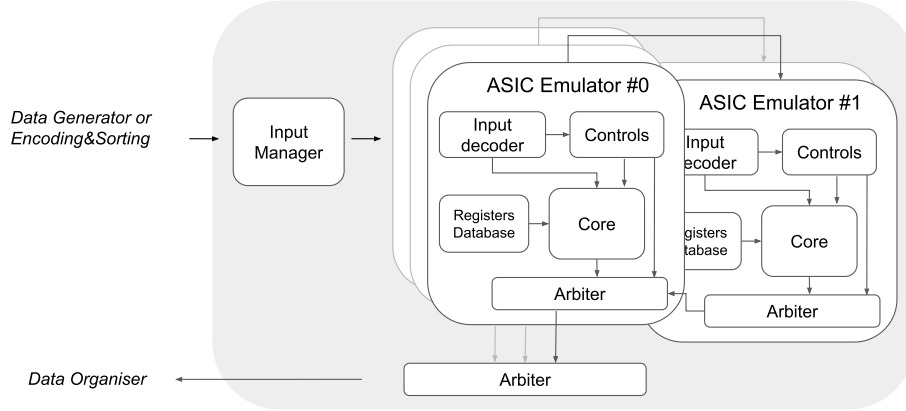


Figure 9. ASIC Emulator and Interface block diagram.

execution timing structure. Events are initialized before the first SSIDs are to be sent to the core for comparison, and ended to gather the information from all matches. All matched roadIDs are then read out in series. The roadIDs from the Core block are combined in the Arbiter block with the output from the Controls block. The output data streams of the five emulators are connected in three daisy-chains of up to two emulators, mimicking the hardware implementation on the demonstrator. The total output data is collected in the ASIC Interface via another Arbiter block and sent to the Data Organiser.

3.3 Data Organiser

The Data Organiser is an on-the-fly database that allows fast access to the cluster position and their SSID. This is critical to extract the cluster from a matched pattern for the linearized track fit. A schematic view of the Data Organiser block is shown in figure 10.

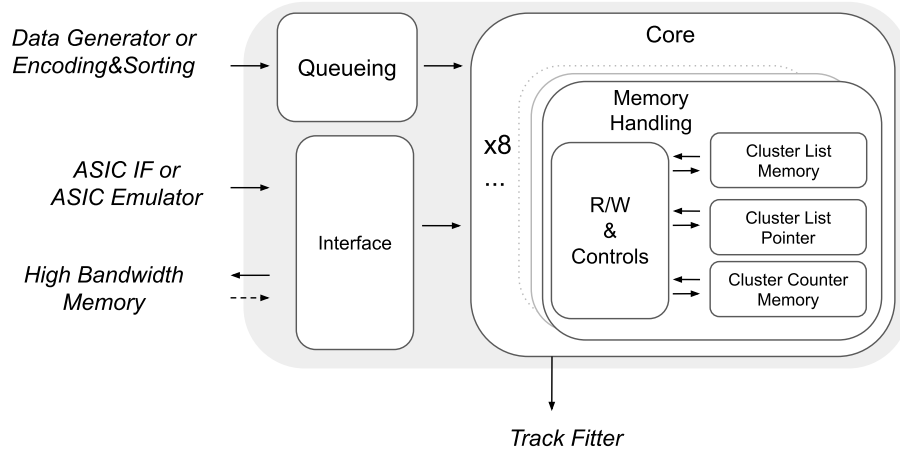


Figure 10. Data Organiser block diagram.

The clusters and SSIDs coming from the Data Generator or from the SSID Encoding and Sorting blocks, already organised per layer, are temporarily stored in FIFOs in the Queueing block. This makes sure that the readout of the previous event has been completed.

Later, a three-memory structure is used in order to better retrieve the clusters. The Data Organiser has been designed under the assumption that all clusters for a given SSID arrive in one uninterrupted sequence. The Cluster List Memory sequentially stores each cluster received. The Cluster List Pointer stores the Cluster List Memory address of the first cluster stored for each SSID. The Cluster Counter Memory stores the number of clusters in each SSID. These functions are performed in parallel for all eight layers.

The Interface block forwards the roadID from the ASIC Emulator blocks towards the High Bandwidth Memory. The SSID list returned by the HBM for a matched pattern is used to retrieve the correct clusters from the DO memories. The found roads are then sent towards the Track Fitter block.

3.4 High-Bandwidth Memory Interfaces

The High-Bandwidth Memory block, shown in figure 11, stores the patterns as well as the constants needed for the computation of the quality of the track fit (χ^2) and the corresponding track parameters.

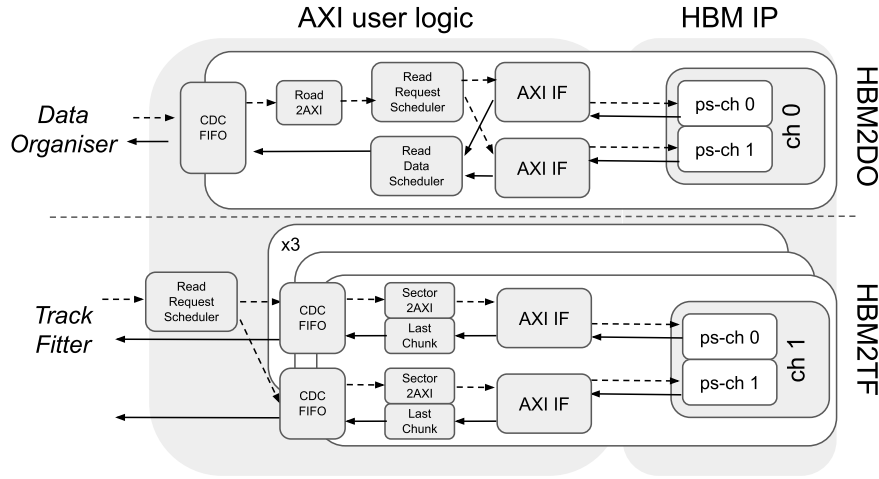


Figure 11. High-Bandwidth Memory user logic block diagram. CDC FIFO = Clock Domain Crossing FIFO, ps-ch = pseudo-channel. The functionality of the individual blocks is described in the text.

FIFO buffers are used at the interfaces between the core logic and the HBM AXI user logic for crossing the clock domains, as each HBM channel has an independent clock, and also for handling back pressure. Two distinct user interfaces are implemented to best match the required data rates and interfaces of the Data Organiser (HBM2DO) and the Track Fitter (HBM2TF).

As the individual data packets requested from the HBM are rather small (ranging from 18 Bytes for each pattern for the Data Organiser, to a maximum of 364 Bytes for each set of constants for the Track Fitter), all channels used of the Intel HBM IP [21] are configured in burst length 4 mode, returning 32 Bytes per single read request. The HBM IP internal Read-Order-Buffers are enabled to simplify the data flow. They can be bypassed to improve latency and bandwidth. The remaining settings of the IP have been optimized in simulation studies for lowest latency and highest throughput. Latency histograms have been implemented within the AXI interfaces to be able to monitor the read data turnaround time.

The Data Organiser requests the patterns from the HBM2DO interface by sending the roadID paired with a requestID. The HBM2DO translates the roadID into a valid memory address

(Road2AXI). Two pseudo-channels are used within the HBM2DO to match the maximum rate requirements. A scheduler checks which pseudo-channel can serve a read request and evens the load on the two pseudo-channels connected using round-robin arbitration. The requestID is stored temporarily matching the AXI-ID used to send out the read request to the HBM IP.

The data received from the HBM is matched to the requestID retrieved via the AXI-ID. A scheduler merges the data streams of the two pseudo-channels towards the output FIFO. The data is split into sectorID and SSIDs for the Data Organiser to continue processing.

The Track Fitter requests constants from the HBM2TF for χ^2 and/or parameter computations specifying the sectorID. The request gets distributed to one of 6 pseudo-channels connected using round-robin arbitration, taking into account if a pseudo-channel is currently busy. Within a pseudo-channel interface, a number of consecutive AXI memory read requests is issued to retrieve a full set of constants (Sector2AXI). Upon return of the data, the last chunk of a set of constants is flagged which allows the Track Fitter to easily identify when a set of constants has been fully retrieved. This allows operation of the HBM even without the Read-Order-Buffers being enabled.

In order to write the patterns and constants data into the HBM, an IPbus [22] compatible interface has been implemented. For online monitoring and configuring of the HBM, an IPbus to Advanced Peripheral Bus (APB) translator interface has been developed.

3.5 Track Fitter

The Track Fitter block, shown in figure 12, assesses the quality of the track candidates by computing the χ^2 value, and then extracts the track parameters for those tracks that pass the χ^2 test. Upon arrival of a road data set, the Track Fitter retrieves the constants for the χ^2 computation from the HBM. All constants are encoded as signed 32 bit floating point numbers. Hence, all DSPs implemented are used in floating point mode.

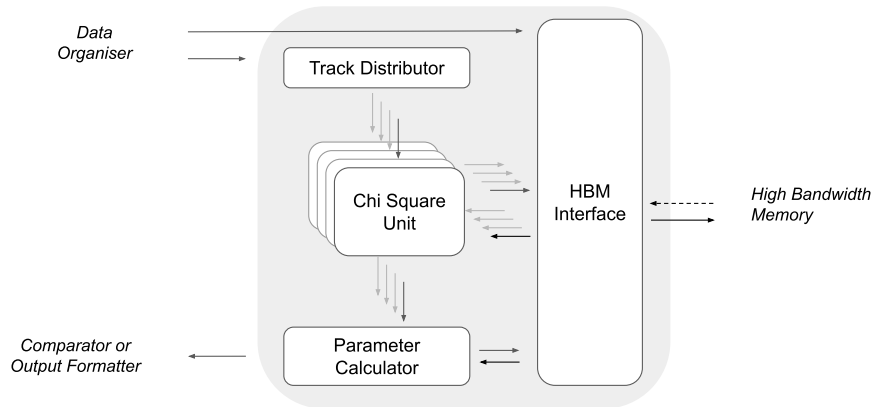


Figure 12. Flow diagram of the Track-Fitter. Buffers and other minor entities are not displayed.

In the Track Distributor, all possible combinations of clusters from the matched patterns are generated as candidate tracks for the fit. To match the required data rates, this is performed through four parallel sub-entities, each associated with a Chi Square Unit block. For each track candidate, the HBM interface retrieves the corresponding constants. All constants requests are buffered while one set of constants is retrieved, and a supporting entity aligns each set to the corresponding track data.

The Chi Square Unit receives a track candidate (x_j), and its IDs and constants ($S_{i,j}$ and h_i) simultaneously. Here, the scalar products per number of coordinates are executed in parallel for all degrees of freedom, according to the inner sum of equation (2.2), using a cascaded pipeline. Likewise, the scalar product over all degrees of freedom is performed, according to the outer sum of the equation (2.2). This design can perform a calculation at every clock cycle with a fixed latency, and originates from the works of [23] and [24]. Finally, the track parameters for this track candidate will be computed if the resulting χ^2 is below a predefined threshold.

Since only 8% of the track candidates are expected to pass the selection in the studied configuration, see table 2, only one Parameter Calculator block is instantiated. Here, the HBM interface provides again the constants, and a scalar product per parameter is executed according to equation (2.1). The resulting parameters, the χ^2 of the track fit, and the clusters are converted into 16 bit fixed point numbers and transferred to the comparator block of the data generator.

3.6 Monitoring

The Monitoring block handles the communication between the PRM FPGA and an external PC, via Ethernet using IPbus protocol over UDP. The block diagram is shown in figure 13.

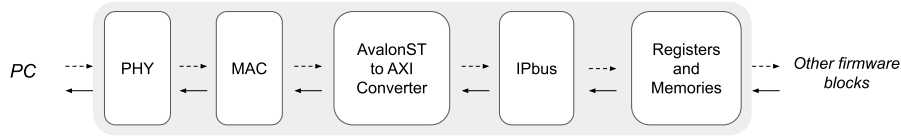


Figure 13. Monitoring block diagram.

The Physical (PHY) and the Medium Access Control (MAC) layer blocks are Ethernet-dedicated IP cores provided by Intel. To allow the communication between the MAC block (Avalon ST interface) and the IPbus block (AXI Interface), a dedicated converter module has been designed. Information about the functionality of the IPbus core can be found in [22]. Hardware checks (e.g. voltages, temperatures), firmware configuration (e.g. data generator setup), and data flow monitoring (e.g. FIFO full and empty, counters) can be all executed via this block. The software part of the communication is provided by the same group which developed the IPbus core [22].

4 Implementation and results

The firmware described above has been implemented and tested on an Intel Stratix 10 MX development kit, which features a nearly-identical FPGA to the one chosen for the PRM prototype.⁴ Simulation studies have been conducted using QuestaSim version 2019.4 [25]. The Intel Quartus Prime Pro software version 19.2 [26] has been used for hardware synthesis and fitting. The results of the initial firmware implementation are presented. It does not include further optimizations that would allow to meet all the requirements of the HTT system; these are presented as proposals in the final section of the paper.

⁴The FPGA on the development kit (part number 1SM21BHU2F53E1VG) and the one on the PRM (part number 1SM21BHN3F53E3VG) differ only in transceiver count, transceiver speed grade, and FPGA fabric speed grade.

4.1 Resource usage

The resource utilization after the final fitting step for each entity of the standalone implementation is presented in table 3. The block-by-block as well as the total occupancy are presented in absolute numbers and in % fraction of the full device resources. In addition to the individual components, the resource usage of the top level (where all blocks are interconnected) is shown with the addition of SignalTap [27] debug monitoring. All blocks that are expected to consume significant logic resources are implemented and considered in this measurement. More advanced features, in particular the extension to include tracks with missing clusters, are not implemented.

Table 3. Resource usage of the PRM standalone firmware on the Intel Stratix 10 MX 2100. Absolute values are given (in multiples of 1000 = k), accompanied by the corresponding percentage of the device in brackets. For considerations on the final design the Data Generator and the ASIC Emulator blocks should not be taken into account, as they are used only for standalone testing. Abbreviations are explained in the text.

	ALMs k (%)	Logic Registers k (%)	Block Memory kb (%)	DSPs # (%)	HBM GB (%)
Data Generator	17.36 (2.47)	26.67 (0.95)	206.03 (0.15)	0 (0.00)	0 (0.00)
ASIC Emulator	7.06 (1.08)	16.42 (0.58)	45.28 (0.03)	0 (0.00)	0 (0.00)
Data Organisier	4.78 (0.68)	6.20 (0.22)	11,625.48 (8.29)	0 (0.00)	0 (0.00)
High Bandwidth Memory	14.15 (2.01)	39.85 (1.42)	188.91 (0.13)	0 (0.00)	2.00 (25.00)
Track Fitter	112.06 (15.95)	235.68 (8.38)	541.84 (0.39)	424 (10.7)	0 (0.00)
Monitoring	7.90 (1.12)	12.97 (0.46)	934.15 (0.67)	0 (0.00)	0 (0.00)
Top level + Debug	8.19 (1.17)	17.14 (0.61)	321.06 (0.23)	0 (0.00)	0 (0.00)
TOTAL	172.04 (24.48)	354.92 (12.63)	13,862.75 (9.89)	424 (10.71)	2.00 (25.0)
TOTAL w/o Data Generator and ASIC Emulator	154.69 (20.93)	328.25 (11.09)	13,611.44 (9.71)	424 (10.71)	2.00 (25.0)

The Track Fitter, where the track-dedicated algorithms are implemented, is the block using the most logic resources, i. e. Adaptive Logic Modules (ALMs), registers, and Digital Signal Processing (DSP) blocks. The Data Organisier, where the data are stored during the pattern recognition step, is the block consuming the majority of the memory resources. The FPGA occupancy does not exceed one quarter of the available resources. Considering that optimizations could further reduce the resource utilization significantly, it would be feasible to fulfill the PRM system requirement of implementing four independent instances in parallel, including the missing extensions. For operation on the PRM board, the ASIC emulator would be replaced by the actual ASIC interface and the Data Generator by the input/output buffers, synchronization modules,⁵ and the output formatter. Hence, the total resource utilization excluding the contribution of these two entities should be considered.

4.2 Clock and power studies

In order to satisfy the system requirements presented in table 2, the PRM firmware was planned with the goal of running the Track Fitter block at 200 MHz, and the rest of the logic at 250 MHz.

The hardware tests with the standalone firmware have been performed constraining the design for a core clock speed of 100 MHz, aiming to verify the hardware implementation rather than

⁵The Cluster to SSID Encoding block is already included in the Data Generator.

optimizing for highest processing speed. A single clock domain has been used for the entire design, with the exception of the HBM that is operated at a memory clock frequency of 600 MHz and a core clock frequency of 300 MHz. Without violating internal setup and hold time requirements, the maximum core clock frequency f_{\max} reported by the Timing Analyzer (part of the Intel Quartus Prime Pro software) is $f_{\max} = 105$ MHz in the worst case (Slow 900 mV Model @ 100 °C).

Under these conditions, the Quartus Power Analyzer tool estimated the total thermal power dissipation to be about 11 W, where about 6 W are due to the device static thermal power dissipation. Unfortunately, the Quartus tool does not include a power model for the HBM component. Its power consumption can hence not be properly estimated in this stage of the development.

In order to better understand the most critical blocks, dedicated 100 MHz clock lines are used for each block by splitting the core clock with a PLL. The results per block are reported below in table 4. It is important to note that in these studies clock domain crossings have not been properly implemented. Although this should not affect the performance and functionality of the individual blocks, treating these crossings correctly is important for their interconnection, and could potentially further improve the results. The following operating conditions are used for this study. The V_{cc} minimum supply voltage is fixed at 900 mV, and the temperature is set to 0 °C and 100 °C, which represent the two extreme corner cases for the f_{\max} value.

Table 4. Per block f_{\max} derived from static timing analysis for a target clock of 100 MHz (300 MHz for the HBM). For considerations of the final design only the Data Organiser, HBM and Track Fitter should be taken into account.

Slow 900 mV Model	Target Clock Speed (MHz)	Data Organiser f_{\max} (MHz)	HBM f_{\max} (MHz)	Track Fitter f_{\max} (MHz)	Data Generator f_{\max} (MHz)	ASIC Emulator f_{\max} (MHz)
100 °C	100	141.60	301.39	114.94	141.26	136.69
0 °C		154.63	322.37	125.20	152.39	148.08

A test has been made to evaluate the work needed to fulfill the HTT system requirement to run at higher frequencies. With the current design constraints and targeting an input clock speed of 200 MHz, an $f_{\max} = 174$ MHz (Slow 900 mV Model @ 100 °C) was achieved, not yet meeting our requirements. The Track Fitter block yields the lowest f_{\max} of all the individual blocks. An algorithm rework would be required for a substantial speed optimization, as will be discussed in the final section.

For this clock speed, the Quartus Power Analyzer tool estimated the total thermal power dissipation to be about 15 W, where about 6 W are due to the device static thermal power dissipation, excluding the HBM contribution.

4.3 Performance studies

The performance of the standalone firmware, with a focus on the processing latency, has been evaluated in simulation using dedicated test vectors derived from the software simulation of the HTT system [8]. A test vector is a collection of information of cluster positions and track parameters, generated by the simulation that reproduces the algorithms of the linearized track fit, and converted into 32 bit words readable by the Data Generator interface. Simulated events with single muon tracks crossing the ITk layers are used.

Larger event sizes have been studied by repeating the same track multiple times as an intermediate step towards more realistic conditions. Each track consists of one cluster per layer. This was done to study the constants request rate and latency, as every track implies a request of constants towards the HBM. Test vectors with few additional fake clusters, i. e. clusters not belonging to any real track, have also been used during hardware verification, but have not been part of the performance evaluation. Studies trying to mimic realistic HL-LHC conditions by injecting events which fulfill approximately the number of roads and fits per event, as specified in table 2, have been started but could not be concluded for this publication. For the performance studies reported here, the SSID Encoding and Sorting block within the Data Generator has not been included. However, both Data Generator modes, with and without the SSID Encoding and Sorting block, have been verified in hardware. Despite the limitations in clock speed seen in the hardware implementation, the simulation has been run at a clock speed of 250 MHz, with the exception of the HBM user logic running at 300 MHz, as explained in section 4.2.

For events with single and multiple tracks, the processing latency per block has been measured as a function of the event rate. Single-track events are filled at a rate of 886 kHz. For events with 16 tracks we chose an event rate of 541 kHz, leading to road, track, cluster, constants readout and fit rates of 8.65 MHz. These event rates are chosen to keep well separated single-event processing, since the handling of multiple concurrent events is not implemented.

Figure 14 shows the histograms of the memory latency between the HBM and the Data Organiser, between the HBM and the Track Fitter, and also the full processing latency of each individual track from the output of the ASIC emulator to the output of the Track Fitter. About 4000 events have been processed to capture the effect of the HBM on the latency. The majority of the read requests are served in a short time of about 250–300 ns; there are tails of up to about 1 μ s due to the dynamic nature of the HBM, being a DRAM, which requires refreshing of the stored data every few microseconds. During the refresh, the HBM pseudo-channel cannot serve any read requests. This behaviour is inherently different compared to internal/external SRAM or the FPGA’s internal block RAM. Compared to single track events, no significant increase on the average HBM latency per pattern or set of constants is observed for events with 16 tracks, because the overall HBM request rates are still well below the bandwidth limits.

The accumulated HBM latency accounts for the biggest part of the total latency as there are three separate accesses per track: to retrieve the pattern, the constants for the χ^2 computation, and the constants for the track parameter calculation. Considering the mean values of the latency distributions, about 66 % (60 %) of the latency is due to the HBM for events with single tracks (16 tracks). Moreover, we have an increase of latency in the total processing pipeline. This can be linked to the Data Organiser, which in its current state needs 3 clock cycles to transfer a single road to the Track Fitter, and to the Track Fitter itself, whose processing can be temporarily interrupted when the requested constants are delayed. The average total processing latency is 1.3 μ s (1.5 μ s) for events with single tracks (16 tracks), with increasing tails in the distribution for higher track rates.⁶ These tails could be reduced by optimizing the design of the Track Fitter and the output of the Data Organizer. No significant increase in latency due to the HBM is expected at the required rates (table 2). More details on design optimizations are given in the following section.

⁶In the presence of high energy jets and pileup 200 we would expect an average of about 320 roads and 80 tracks per event per group of ASICs.

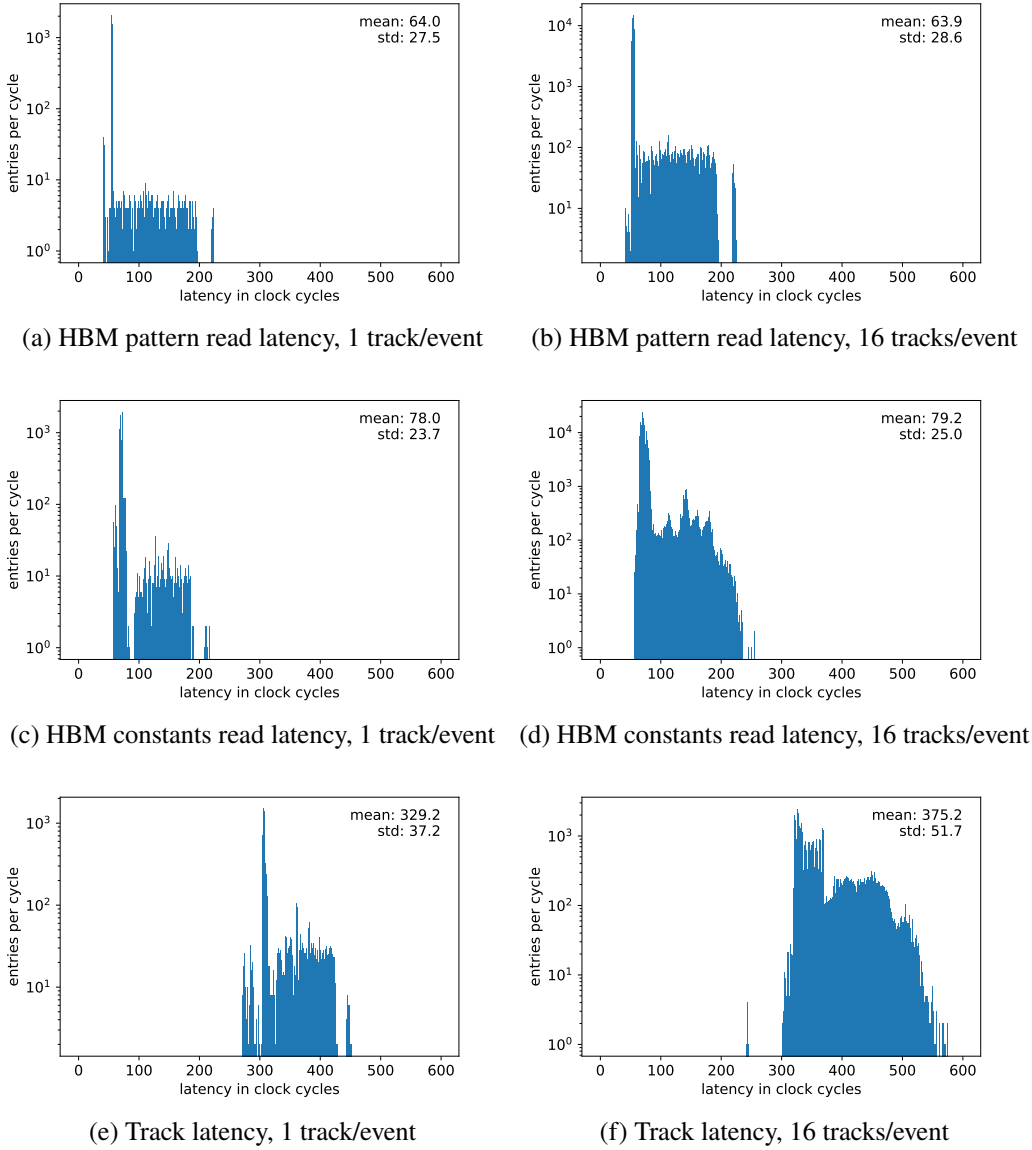


Figure 14. Processing latencies of different parts of the PRM firmware for events with a single track (left) and events with 16 tracks (right) in clock cycles of $f_{\text{core}} = 250$ MHz. (a) and (b) show the latency of patterns being returned from the HBM when requested by the Data Organiser. (c) and (d) show the latency of full sets of constants returned from the HBM when requested by the Track Fitter. (e) and (f) show the latency of a track through the processing pipeline, between its appearance as a road at the output of the ASIC emulator and the final track at the output of the Track Fitter. For all distributions, mean and standard deviations are given.

4.4 Considerations for a final design

These studies show a possible implementation of the main functionalities of the firmware for a linearized track fit with 8 detector layers, associated to AM pattern matching. Not all requirements from the HTT system design, outlined in table 2, are achieved at this stage of our development. For example, the χ^2 fit rate requirement of 714 MHz can be achieved with a clock speed for the Track

Fitter of 200 MHz, but hit a limit of $f_{\max} = 175$ MHz in our design. Because of this, and in order to lower its resource usage, the Track Fitter block would be the very first candidate for optimizations. This is despite the expectation that it will consume the largest fraction of resources compared to all other blocks. For instance, the interface between the Track Fitter and the HBM could be simplified such that there is one dedicated Chi Square Unit per HBM pseudo-channel, instead of the current implementation where all four Chi Square Units can retrieve their constants from all of the six HBM pseudo-channels. This may have a positive impact on the achievable f_{\max} within the Track Fitter, and therefore increase the overall throughput. The mean processing latency measured in simulations exceeds the target of 1 μ s for the low-latency scenario. Several options could be adopted to reduce latency:

- From simulation studies we know that the Read-Order-Buffer within the HBM Controller yields a large latency penalty, especially when considering higher request rates. Removing this buffer, however, comes at the cost of increasing complexity of the data processing at the Track Fitter and Data Organiser.
- The constants for the χ^2 and parameter calculations within the Track Fitter could be requested in parallel, as opposed to the current implementation where the second set of constants is only requested when the track has passed the χ^2 cut. This comes at the cost of increased HBM bandwidth.
- The constants for the Track Fitter could be requested as soon as the HBM has retrieved the pattern for the Data Organiser in parallel to the processing pipeline. This is instead of waiting for the data being transferred from the Data Organiser to the Track Fitter, which then requests the constants. Moreover, the sectorIDs could be stored in the FPGAs internal eSRAM, which has a much shorter read latency than the HBM. That would allow for truly parallel retrieval of the patterns and the fit constants from the HBM. From simulation studies of the HBM model at the required roads processing rates and constants readout rates specified in table 2, it is expected that 99 % of patterns and constants would be returned within 650 ns after their request.
- The Track Fitter could be adapted to start calculations as soon as a single chunk of constants has arrived, instead of waiting for the full set of constants being returned from the HBM. This would however require a drastic re-design of the Track Fitter block.
- The overall number of constants could be reduced by adopting an algorithm as it is reported in [14]. This makes use of hit position transformations to address the non-linearities of the system in order to reduce the number of sectors. This could allow the removal of the need for an HBM to retrieve the constants at all. However, this would not only require a complete redesign of the blocks following the Data Organiser within the PRM firmware, but also major changes to the simulation software of the full HTT system.
- Lastly, clock speed may be increased if a better speed grade version of the FPGA is used. This change would also impact the HBM's read latency, as the core and memory frequency could be raised as well. This change would not require any design changes on the hardware or firmware, but increase the board's production costs.

5 Conclusion

Within the HTT project, a custom-hardware tracking system has been designed. It is relying on associative memories for pattern matching and FPGAs for data organization, communication, and track fitting. Even if the project is abandoned in ATLAS, many studies have been performed to assess its feasibility, both on hardware demonstrators and with simulations. The Pattern Recognition Mezzanine as part of the HTT system features the pattern matching and track fitting capabilities. Key challenges to the design of the PRM firmware are the high processing rates, e. g. the fit rate of 4×714 MHz, and the low-latency target of order $1 \mu\text{s}$ requested for a short-latency L1 trigger scenario.

This paper describes the implementation of the firmware for the PRM board, its validation in simulation, and tests on an Intel Stratix 10 MX development kit. The major processing blocks, i. e. Data Organiser, Track Fitter, and HBM interfaces, have been implemented and tested using an AM ASIC Emulator and Data Generator developed specifically for this task. Resource utilization of these blocks is in agreement with the requirements for a full PRM board. Even if the requirements imposed by the HTT system design are not met, its feasibility with modern FPGAs with High-Bandwidth Memories is assessed. Although High-Bandwidth Memories are optimized for highest throughput at a cost of higher latencies compared to SRAM or internal block memory, they can be utilized in low-latency ($O(1 \mu\text{s})$) applications that require large data storage. Many improvements are possible that could reduce the latency and resource utilization, leaving space for optimizations for custom hardware systems in future experiments.

Acknowledgments

We acknowledge support by the Federal Ministry of Research and Education (BMBF), the Swiss National Science Foundation (SNSF), the University of Geneva, Department of Nuclear and Particle Physics, the Italian Istituto Nazionale di Fisica Nucleare and the Aristotle University of Thessaloniki, School of Physics. We would like to thank Francesca Pastore, Naoki Kimura, and Mel Shochet for valuable feedback that helped shape this article.

References

- [1] G. Apollinari et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical Design Report V. 0.1*, [CERN-2017-007-M](#), CERN, Geneva (2017) [[DOI:10.23731/CYRM-2017-004](#)].
- [2] ATLAS collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, [2008 JINST 3 S08003](#).
- [3] ATLAS collaboration, *ATLAS Phase-II Upgrade Scoping Document*, [CERN-LHCC-2015-020](#), CERN, Geneva (2015) [[DOI:10.17181/CERN.7CRX.AJHP](#)].
- [4] ATLAS collaboration, *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*, [CERN-LHCC-2017-021](#) CERN, Geneva, (2017) [[DOI:10.17181/CERN.FOZZ.ZP3Q](#)].
- [5] ATLAS collaboration, *Technical Design Report for the ATLAS Inner Tracker Strip Detector*, [CERN-LHCC-2017-005](#) CERN, Geneva, (2017).
- [6] ATLAS collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System*, [CERN-LHCC-2017-020](#) CERN, Geneva, (2017) [[DOI:10.17181/CERN.2LBB.4IAL](#)].

- [7] ATLAS collaboration, *Expected tracking and related performance with the updated ATLAS Inner Tracker layout at the High-Luminosity LHC*, [ATL-PHYS-PUB-2021-024](#) CERN, Geneva, (2021).
- [8] ATLAS, collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS Trigger and Data Acquisition System-Event Filter Tracking Amendment*, [CERN-LHCC-2022-004](#), CERN, Geneva (2022) [[DOI: 10.17181/CERN.ZK85.5TDL](#)].
- [9] A. Stabile et al., *Phase-II Associative Memory ASIC Specifications*, [CERN-OPEN-2018-003](#), CERN, Geneva (2018).
- [10] CDF collaboration, *The CDF silicon vertex trigger*, *Nucl. Instrum. Meth. A* **518** (2004) 532 [[physics/0306169](#)].
- [11] ATLAS collaboration, *The ATLAS Fast TracKer system*, [2021 JINST 16 P07006](#) [[arXiv:2101.05078](#)].
- [12] ATLAS collaboration, *Fast pattern recognition of ATLAS L1 track trigger for HL-LHC*, *PoS Vertex2016* (2017) 069.
- [13] I. Tomalin et al., *An FPGA based track finder for the L1 trigger of the CMS experiment at the High Luminosity LHC*, [2017 JINST 12 P12019](#).
- [14] E. Clement et al., *A High-performance Track Fitter for Use in Ultra-fast Electronics*, *Nucl. Instrum. Meth. A* **935** (2019) 95 [[arXiv:1809.01467](#)].
- [15] PICMG, PICMG 3.0 Revision 2.0, *AdvancedTCA Base Specification*, <http://www.picmg.org>.
- [16] ATLAS collaboration, *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System*, <https://cds.cern.ch/record/1602235> CERN-LHCC-2013-018 CERN, Geneva (2013).
- [17] ATLAS TDAQ collaboration, *FELIX: the new detector interface for ATLAS*, *EPJ Web Conf.* **245** (2020) 01037.
- [18] Cisco Systems, *Serial-GMII Specification*, (2005) <https://archive.org/details/sgmii>.
- [19] *IEEE Standard Test Access Port and Boundary Scan Architecture*, in IEEE Std 1149.1-2001, pp.1-212, 23 July 2001 [[DOI: 10.1109/ieeestd.2001.92950](#)].
- [20] ATLAS TDAQ collaboration, *The ATLAS Hardware Track Trigger design towards first prototypes*, *PoS Vertex2019* (2020) 049.
- [21] Intel®, *High Bandwidth Memory (HBM2) Interface Intel®FPGA IP User Guide*, <https://www.intel.com/content/www/us/en/docs/programmable/683189/21-3-19-6-1/about-the-high-bandwidth-memory-hbm2.html>.
- [22] C. Ghabrous Larrea et al., *IPbus: a flexible Ethernet-based control system for xTCA hardware*, [2015 JINST 10 C02019](#).
- [23] M. Braagaard et al., *Implementation of a linearized χ^2 -Square [sic] method for the ATLAS Pattern Recognition Mezzanine*, Technical University of Denmark, Department of Applied Mathematics and Computer Science, Lyngby Denmark (2020).
- [24] M. Jensen et al., *Implementation of Particle Track Parameter Calculations for the ATLAS Pattern Recognition Mezzanine*, Technical University of Denmark, Department of Applied Mathematics and Computer Science, Lyngby Denmark (2020).
- [25] Siemens®, *Questa Advanced Simulator*, <https://eda.sw.siemens.com/en-US/ic/questa/simulation/advanced-simulator/>.
- [26] Intel®, *Intel® Quartus® Prime — FPGA Design Software*, <https://www.intel.com/content/www/us/en/software/programmable/quartus-prime/overview.html>.
- [27] Intel®, *About the SignalTap II Logic Analyzer*, https://www.intel.com/content/www/us/en/programmable/quartushelp/13.0/mergedProjects/program/ela/ela_view_using.htm.