



OPEN ACCESS

EDITED BY

Gabriella Pugliese,
National Institute for Nuclear Physics of Bari,
Italy

REVIEWED BY

Arantxa Ruiz Martinez,
University of Valencia, Spain
Antonio Sidoti,
National Institute of Nuclear Physics of Bologna,
Italy

*CORRESPONDENCE

Pelayo Leguina,
✉ leguinapelayo@uniovi.es

RECEIVED 27 September 2024

ACCEPTED 27 December 2024

PUBLISHED 14 January 2025

CITATION

Leguina P and Folgueras S (2025) Harnessing hardware acceleration in high-energy physics through high-level synthesis techniques. *Front. Detect. Sci. Technol* 2:1502834. doi: 10.3389/fdest.2024.1502834

COPYRIGHT

© 2025 Leguina and Folgueras. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Harnessing hardware acceleration in high-energy physics through high-level synthesis techniques

Pelayo Leguina* and Santiago Folgueras

Department of Physics and Institute of Space Sciences and Technologies of Asturias (ICTEA), University of Oviedo, Oviedo, Spain

At the Large Hadron Collider, the vast amount of data from experiments demands not only sophisticated algorithms but also substantial computational power for efficient processing. This paper introduces hardware acceleration as an essential advancement for high-energy physics data analysis, focusing specifically on the application of High-Level Synthesis (HLS) to bridge the gap between complex software algorithms and their hardware implementation. We will explore how HLS facilitates the direct implementation of software algorithms into hardware platforms such as FPGAs, enhancing processing speeds and enabling real-time data analysis. This will be highlighted through the case study of a track-finding algorithm for muon reconstruction with the CMS experiment, demonstrating HLS's role in translating computational tasks into high-speed, low-latency hardware solutions for particle detection and reconstruction. Key techniques in HLS, including parallel processing, pipelining, and memory optimization, will be discussed, illustrating how they contribute to the efficient acceleration of algorithms in high-energy physics. We will also cover design methodologies and iterative processes in HLS to optimize performance and resource utilization, alongside a brief mention of additional techniques like algorithm approximation and hardware/software co-design. In short, this paper will underscore the potential of hardware acceleration in high-energy physics research, emphasizing HLS as a powerful tool for physicists to enhance computational efficiency and foster groundbreaking discoveries.

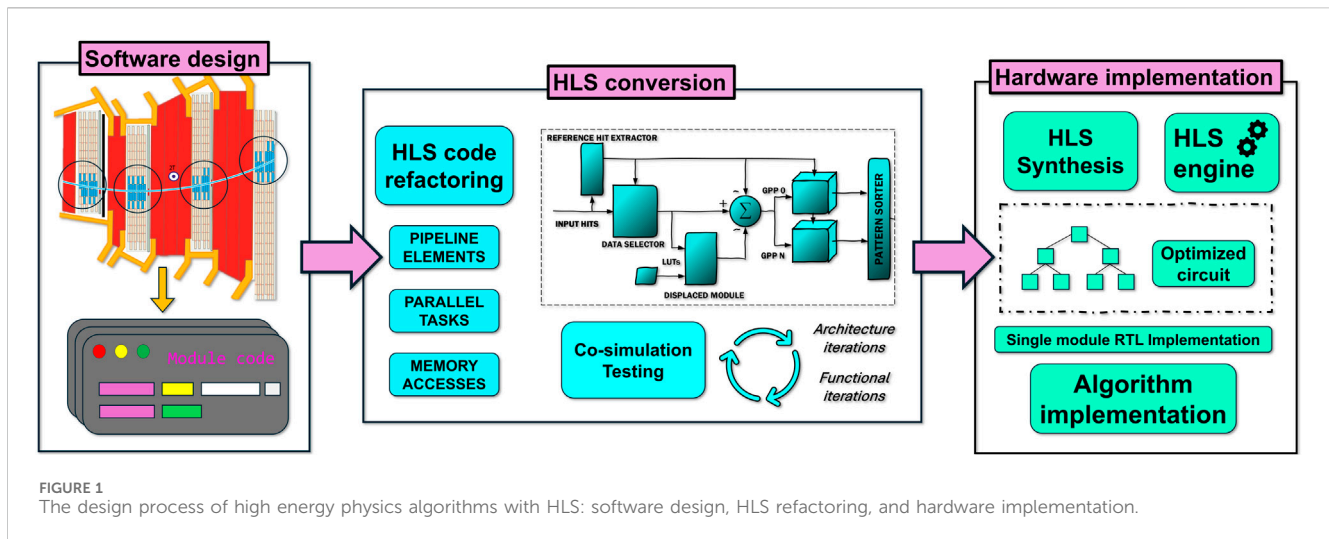
KEYWORDS

high-level synthesis (HLS), hardware acceleration, field-programmable gate arrays (FPGAs), CMS experiment, track-finding algorithm, parallel processing and pipelining

1 Introduction

The Large Hadron Collider (LHC) at CERN represents the pinnacle of high-energy physics (HEP) research, enabling scientists to probe the fundamental constituents of matter and the forces governing their interactions. Experiments like the Compact Muon Solenoid (CMS) generate an unprecedented volume of data, with collision events occurring at rates of several billions per second (Evans and Bryant, 2008). This deluge of data necessitates not only sophisticated algorithms for accurate particle detection and reconstruction but also demands substantial computational resources to process the information in real time.

Traditional software-based data processing approaches, while flexible, often struggle to meet the low-latency and high-throughput requirements of modern HEP experiments. The latency constraints are particularly stringent in trigger systems, where rapid decision-



making is crucial to determine which events are of interest and should be recorded for further analysis (CMS Collaboration, 2020). Hardware acceleration emerges as a vital solution to these challenges, offering significant improvements in processing speeds and enabling real-time data analysis.

In the context of the CMS experiment, fast and efficient particle reconstruction is essential for having a reliable trigger system and for the success of the physics program (Radburn-Smith, 2022). Implementing track-finding algorithms directly into hardware accelerators can drastically improve processing speeds and reduce latency. Previous efforts have demonstrated the feasibility of using FPGAs for real-time tracking in HEP experiments (Aad et al., 2021; Aggleton et al., 2017). However, the manual translation of algorithms into hardware description languages is time-consuming and error-prone.

High-Level synthesis (HLS) has gained traction as a powerful tool that bridges the gap between complex software algorithms and their hardware implementations. HLS allows for the description of hardware functionality using high-level programming languages like C or C++, which are then synthesized into hardware description languages suitable for implementation on Field-Programmable Gate Arrays (FPGAs) or Application-Specific Integrated Circuits (ASICs) (Nane et al., 2016). This approach significantly reduces development time and makes hardware acceleration more accessible to software engineers and physicists who may not be experts in hardware design.

With HLS, we can directly implement sophisticated track-finding algorithms into hardware, leveraging techniques such as parallel processing and pipelining to optimize performance (see Figure 1). Parallel processing allows multiple computations to occur simultaneously, significantly increasing throughput (Nane et al., 2016). Several studies at CERN have employed HLS to optimize pipelining and memory usage. Pipelining enables overlapping of operations, reducing the overall processing time per event, while memory optimization techniques enhance performance by minimizing access times and efficiently utilizing on-chip resources (Husejko et al., 2015; Ghanathe et al., 2017).

This paper presents a comprehensive study on the application of HLS for hardware acceleration in HEP data analysis. We focus on the implementation of a track-finding algorithm for muon

reconstruction within the CMS experiment as a case study. Our work demonstrates how HLS facilitates the translation of complex computational tasks into high-speed, low-latency hardware solutions.

We also discuss design methodologies and iterative processes in HLS to optimize performance and resource utilization. The importance of algorithm approximation is highlighted, where trade-offs between precision and computational efficiency are considered (Han et al., 2016). In addition, we show an automated local pipeline for HLS module building, simulation, and hardware implementation, demonstrating how automation can streamline the development process and enhance efficiency.

This paper is organized as follows: Section 2 provides an introduction to the CMS Level-1 trigger system and the software tools utilized in this work; Section 3 describes the required steps to use HLS for algorithm acceleration; Section 4 depicts the bases for an automated workflow for prototyping, validation and integration; Section 5 describes the obtained results. Finally, Section 6 summarizes our findings.

2 Experimental setup

2.1 The overlap muon track finder of the CMS Level-1 trigger system

The CMS Level-1 trigger system will undergo a significant upgrade to accommodate the increased luminosity and data rates expected from the High-Luminosity LHC (HL-LHC) (CMS Collaboration, 2020). The upgraded trigger system, also referred as Phase-2 Level-1 trigger, is designed to handle up to 750 kHz of event rate with a latency of approximately 12.5 μ s. Our hardware implementation integrates with this upgraded trigger system, requiring compatibility with its stringent latency constraints and data formats. A detailed description of the Phase-2 Level-1 trigger system can be found in (CMS Collaboration, 2020).

The Phase-2 Level-1 trigger system comprises advanced electronics, including state-of-the-art FPGAs and high-speed optical links. The system employs the Advanced

Telecommunications Computing Architecture (ATCA) standards, providing high-density, high-throughput, and low-latency data processing capabilities necessary for real-time event selection in the HL-LHC environment. As a baseline, the Xilinx VU13P FPGA is used, these FPGA offers a high number of logic cells, digital signal processing (DSP) slices, and block RAMs, making them suitable for the complex, high-throughput applications required in HEP experiments. The VU13P features over three million logic cells, 12,288 DSP slices, and 360 Mb of UltraRAM, providing ample resources for implementing deep pipelining and parallel processing techniques inherent in HLS-optimized designs.

As a case study, the track finding algorithm for the overlap muon track finder (OMTF) of the CMS experiment is used. The OMTF is designed to reconstruct muon trajectories in the barrel-endcap transition region of the detector (Zabolotny and Byszuk, 2016). This region is extremely complex due to the inhomogeneous magnetic field and the different geometric orientations of the muon detectors: Drift Tubes (DT), Cathode Strip Chambers (CSC), and Resistive Plate Chambers (RPC). The OMTF algorithm tackles this challenge by evaluating how well the detected hits, or stubs, correspond to expected patterns of muon tracks with specific transverse momenta (p_T). The algorithm matches the incoming detector stubs to a set of predefined reference patterns. Each pattern represents the typical configuration for a muon with a certain p_T and charge traversing the detector layers. By calculating a similarity score between the observed stubs and these reference patterns, the algorithm identifies the most probable track candidate, assigning appropriate kinematic parameters to the muon.

Earlier implementations of the OMTF algorithm were realized in VHDL (Zabolotny and Byszuk, 2016) and later using high-level synthesis techniques (Zabolotny, 2019). Recent developments incorporated the reconstruction of displaced muons (Leguina, 2023). Our implementation incorporates optimizations to reduce latency and resource usage. By utilizing HLS, we achieve a modular and maintainable design that can be easily adapted for future upgrades and more complex detector conditions.

In the current foreseen scenario for the HL-LHC upgrade, the OMTF system consists of 6 ATCA boards, three for each side of the detector. Each board should process the information from a 120° ϕ -sector of the detector and will receive up to 1,033 Gbps of information from the muon detectors. For each clock cycle, the algorithm can deliver up to nine muons, with a maximum bandwidth of 450 Gbps (divided into 18 optical links) that will be sent to later stages of the trigger. The whole OMTF reconstruction has to be accommodate in a 2 microsecond latency budget.

Our implementation ensures compliance with the system's timing budget and interfaces (CMS Collaboration, 2020). The hardware modules are designed to meet the physical, power, and thermal constraints of the CMS trigger crates, ensuring reliable operation under the demanding conditions of the HL-LHC.

The dataset used for this study was generated using a trigger emulator of the aforementioned algorithm written in C++ that generates a test vector formatted in XML files consisting of a small sample of 1,000 events, coming from a muon gun sample with pairs of muons with a flat p_T distribution between 1 and 100 GeV. These test vectors represent the digitized detector signals and are structured to match the input specifications of the hardware

implementation. Using XML allows for a flexible and human-readable format that facilitates debugging and verification processes.

2.2 Software tools

We employ AMD Vitis HLS version 2023.2 for converting the high-level algorithm descriptions into hardware description language (HDL) code suitable for FPGA implementation. Vitis HLS allows for the synthesis of C, C++, and SystemC code into Verilog or VHDL, facilitating rapid prototyping and optimization. The tool supports various optimization directives, such as loop unrolling, pipelining, and dataflow, which are essential for enhancing performance and resource utilization in FPGA designs.

The use of HLS accelerates the development cycle by enabling software engineers and physicists to design hardware accelerators using familiar programming languages. It also allows for quick iterations and testing of different optimization strategies to meet the stringent performance requirements of the CMS Phase-2 trigger system.

In addition, scripting languages such as bash and tcl were utilized for automating the build process, simulation, and generating an automated local pipeline for HLS module building, simulation, and hardware implementation. Python scripts were also employed for data analysis, visualization, and further automation tasks. The integration of these scripting languages streamlines the development workflow and facilitates collaboration among team members.

VHDL was used for lower-level hardware description and for integrating the HLS-generated modules into the existing hardware infrastructure. The use of VHDL allows for precise control over hardware resources and timing, which is critical in matching the performance requirements of the CMS Level-1 trigger system.

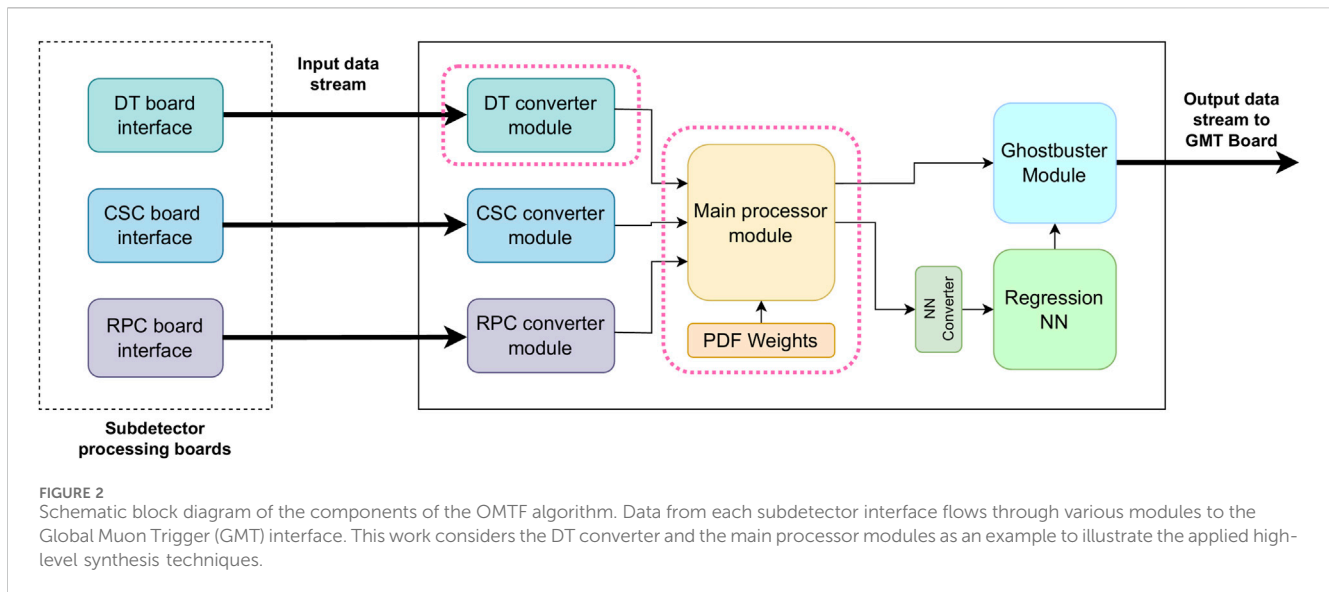
3 Algorithm acceleration using high-level synthesis in high-frequency applications

The adaptation of an algorithm for hardware acceleration using high-level synthesis involves a thorough understanding of its main components and how data flows through them. The algorithm must be adapted to operate efficiently on FPGA hardware, which requires consideration of data streaming, parallel processing, and pipelining.

3.1 Algorithm implementation

Figure 2 shows the main modules in the HLS implementation of the OMTF algorithm, which receives input data from the multiple muon subdetectors: DT, CSC and RPC. Each produces data in its own specific format and frame structure. These data are transmitted in a streamed fashion from various parts of the Level-1 trigger system to the conversion modules of the algorithm.

To handle this heterogeneous and continuous stream of data, we design HLS modules capable of processing streamed inputs using the `hls::stream` library. This library provides FIFO-based data streams that facilitate the handling of sequential data in a pipelined



manner. The input converting modules continuously read incoming data streams, process them in real-time, and ensure that no data is lost due to buffer overflows or processing delays.

The design process of the OMTF system consists of several key steps, with a focus on efficient data handling and processing in an FPGA environment:

- **Input converting modules:** These modules handle data from each subdetector using `hls::stream` objects to efficiently manage streamed data. Each module parses incoming data frames, extracts relevant hit information (primitives), and converts the data into a unified format that is compatible with the main processing module. Additionally, they merge information from multiple detector layers (DTs, CSCs, and RPCs) and ensure that the data format matches the input structure required by the main module. To minimize latency and maximize throughput, directives such as `#pragma HLS PIPELINE` and `#pragma HLS DATAFLOW` are used.
- **Main processing module:** This module receives the unified data frame through custom input ports and implements the core OMTF algorithm to infer p_T based on the spatial positions of detected primitives. In addition, the module processes multiple “golden patterns” (Zabolotny and Byszuk, 2016) in parallel, which represent the expected hit configurations for muons with specific p_T and charge. The design leverages `#pragma HLS UNROLL` to unroll loops and create multiple instances of the processing logic for parallel computation. It also employs `#pragma HLS ARRAY_RESHAPE` to ensure that array elements storing pattern (PDF) weights and parameters are available simultaneously, eliminating data dependencies and improving efficiency. A more detailed description of this module can be found in (Zabolotny and Byszuk, 2016).

This structured approach ensures that the OMTF system efficiently processes data while optimizing FPGA resources and

maintaining the necessary throughput for real-time applications that was described in Section 2.

Throughout the design process, we continuously verify the functionality and performance of each module using C++ test benches and HLS simulations. This iterative approach allowed us to identify bottlenecks and optimize the design before synthesizing it onto the FPGA hardware. In this work the DT converter and the processor modules are taken as an example to illustrate the applied high-level synthesis techniques.

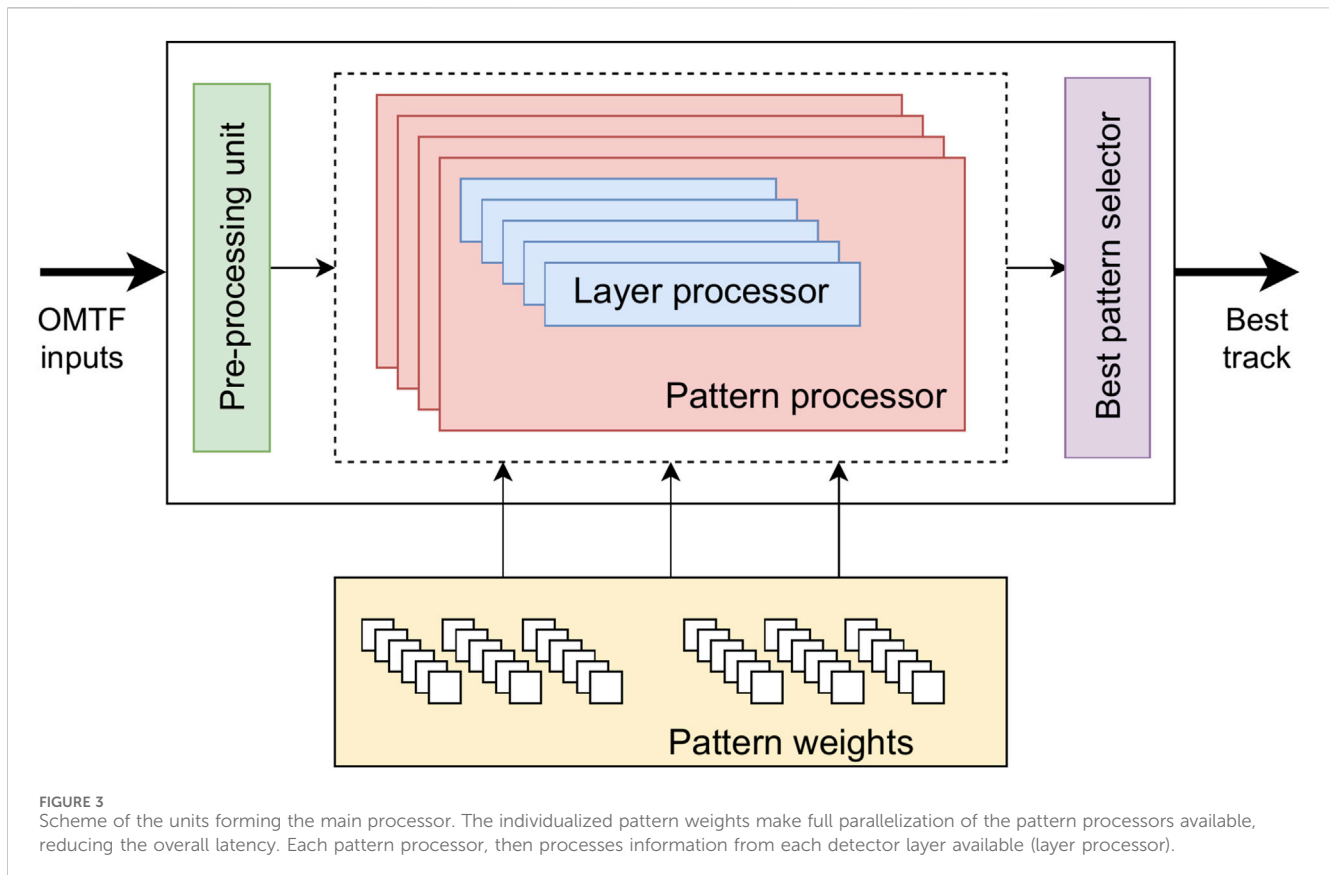
3.2 Optimization techniques

Optimization was a critical aspect of our implementation to meet the stringent performance requirements of the CMS Level-1 trigger system. Several techniques were applied, focusing on parallel processing, pipelining, and memory optimization, each tailored to the specific needs of different modules within the algorithm.

3.2.1 Parallel processing and memory optimization

Parallel processing was primarily employed in the main processing module, where the evaluation of patterns is inherently parallelizable. Each pattern represents a potential muon trajectory with a specific p_T and charge. To process all patterns simultaneously, we unroll loops that iterate over the patterns using the `#pragma HLS UNROLL` directive. By fully unrolling these loops, separate processing elements for each pattern were instantiated, allowing the module to compare incoming primitives against all patterns in a single clock cycle, significantly reducing latency as shown in Section 5.

Additionally, arrays containing pattern weights and parameters were reshaped using the `#pragma HLS ARRAY_RESHAPE` directive with the “complete” option, ensuring that all array elements were accessible in parallel without causing memory access conflicts. Constant data, such as pattern weights, were stored in read-only memory blocks, mapped to on-chip block RAMs or UltraRAMs using the `#pragma HLS RESOURCE`



directive. Array partitioning through the `#pragma HLS ARRAY_PARTITION` directive enabled multiple simultaneous read accesses, further optimizing memory access.

In the input converting modules, parallel processing was also applied by designing separate modules for each muon subdetector. Each module operates independently, handling its specific data stream and merging results into a unified data frame for the main module, ensuring seamless integration.

As shown in Figure 3, the individualized pattern weights allow full parallelization of the pattern processors, significantly reducing the overall latency of the system.

3.2.2 Pipelining

Pipelining was a crucial optimization technique applied to the input converting modules. Taking the DT converting module as an example, the DT module processes data frames representing muon trigger primitives, or stubs. Each event consists of nine data frames, each 64 bits wide. Traditionally, these frames would be deserialized, combined, and processed together, introducing latency and increasing resource usage. To counter this, we use the `hls::stream` library to process primitives as they arrived, enabling a pipelined approach.

The pipeline in the DT module consists of several stages:

1. Streaming input read: Primitives are read from the `hls::stream` input one at a time.
2. Quality and position filtering: Apply detector-specific quality criteria to filter out spurious hits and ensure the hit positions are within range for muon candidates.

3. Coordinate conversion: Use DSP units for local coordinate computations, applying fixed-point arithmetic to maintain precision while optimizing resources.

4. Data packing: Once all nine primitives are processed, the converted data is packed into the expected input format for the main processing module.

To further enhance the efficiency of the DT module pipeline, the `#pragma HLS DATAFLOW` directive was used. This directive allowed different stages of the pipeline (e.g., reading, filtering, conversion, and packing) to operate concurrently, rather than sequentially. By enabling parallel execution of the pipeline stages, we minimize the overall latency of the data processing, ensuring that each primitive is processed as soon as it arrives, without waiting for the entire event to be collected.

By applying both the `#pragma HLS PIPELINE` and `#pragma HLS DATAFLOW` directives, we ensure that each operation could process new data every clock cycle, allowing continuous data flow without waiting for the entire event to be read, thus reducing latency.

The pipelining technique has been applied in the DT module to demonstrate its effectiveness. By processing each primitive individually as it arrived, and using the `#pragma HLS DATAFLOW` directive to allow concurrent execution of pipeline stages, we avoid the latency and resource demands associated with traditional deserialization and collective processing of data frames. The pipeline stages described above enabled continuous processing, achieving high throughput with minimal delay, as shown in Figure 4. We achieve significant reductions in latency (factor 2.43,

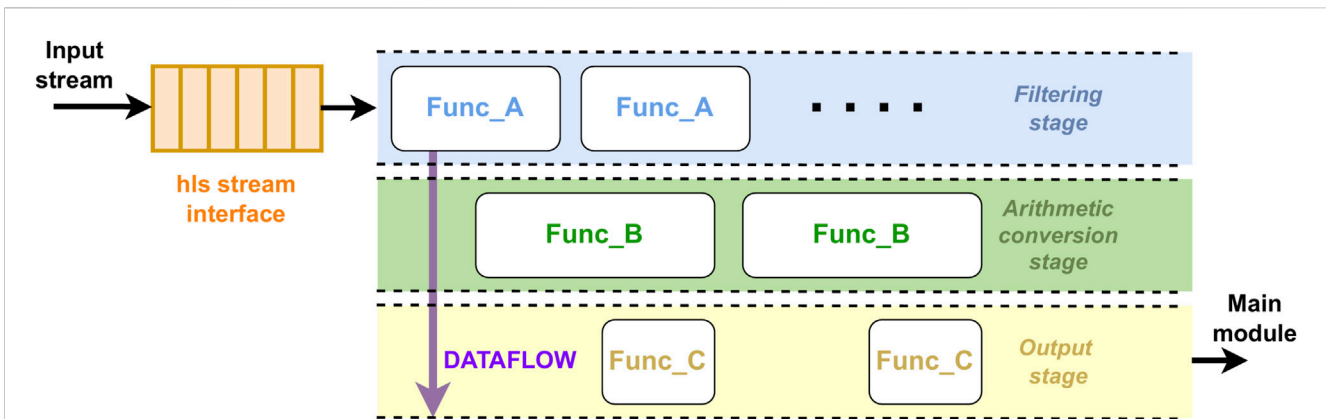


FIGURE 4 Pipeline design of the DT module converter, utilizing a streamed interface inferred by the `hls::stream` class. The functions are executed in parallel following the HLS DATAFLOW paradigm, ensuring efficient data processing and throughput.

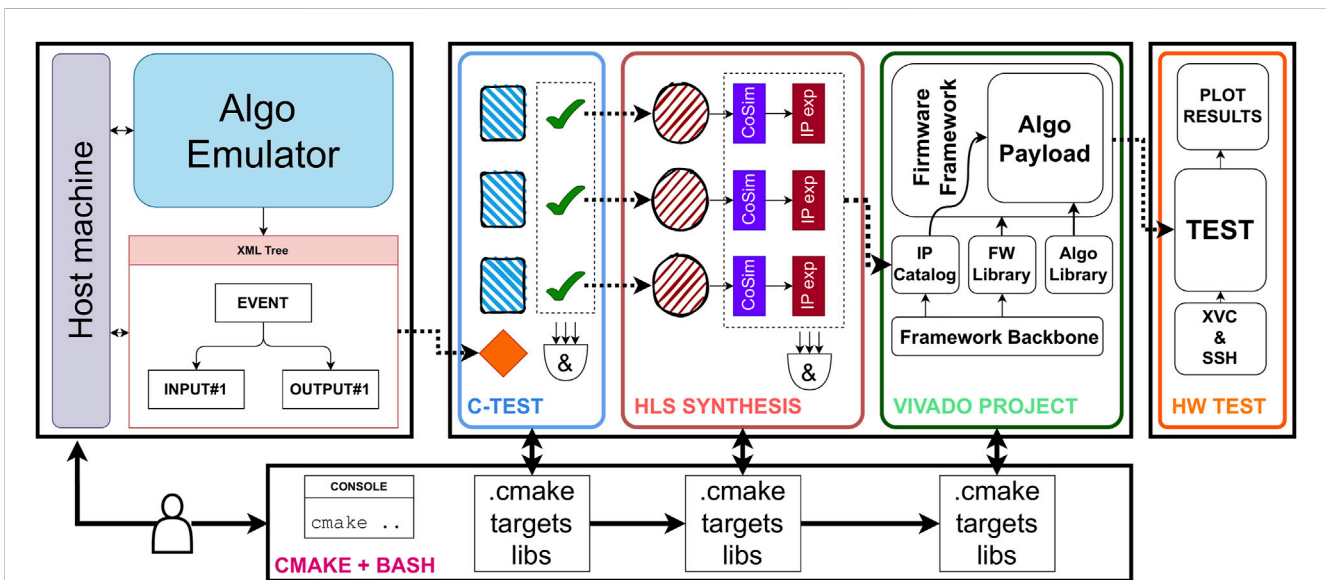


FIGURE 5 HLS build pipeline overview. The process begins with test vector generation in the algorithm emulator software. Next, the HLS modules are tested and synthesized. Following this, the board framework is built, and the top file along with the block design, including all HLS modules, is generated and implemented. Finally, after generating the bitstream, the design is tested on the target hardware.

see Section 5) while ensuring that data could be processed continuously and efficiently within the CMS Level-1 trigger system’s stringent timing constraints.

3.2.3 Pause points

Throughout the development process, we implement pause points to verify the correctness and performance of the algorithm at various stages. These pause points are strategically placed after each implementation step to ensure that any issues could be identified and addressed promptly, preventing the propagation of errors to subsequent stages.

This structured approach to validation made us realize the potential for creating a pipeline where automatic verification could be applied at each step. The concept of continuously

verifying the design at each stage naturally evolved into the idea of integrating automation throughout the entire development process.

4 Rapid prototyping workflow through automation scripting

To manage the complexity of the development process and ensure efficient verification, we implement a custom local building and verification pipeline, see Figure 5. This pipeline automates the compilation, testing, synthesis, and simulation steps, integrating them into a cohesive workflow that streamlines the development of the hardware-accelerated algorithm.

The validation of each step of the pipeline is crucial to ensure that the hardware-accelerated implementation of the track-finding algorithm functions correctly and mirrors the expected behavior of its software counterpart. The validation process involves several key steps:

1. **Functional equivalence testing:** We perform comprehensive tests to compare the outputs of the hardware-implemented algorithm with those of the software version using identical input data from the CMS experiment. This ensures that the logic and computational integrity of the algorithm are preserved during the translation from software to hardware.
2. **Bit-true verification:** Utilizing co-simulation features provided by HLS tools, we verify that the hardware design operates correctly at the bit level. This involves checking that the numerical computations and data handling in the FPGA match the software algorithm's precision and accuracy.
3. **Iterative refinement and debugging:** Any discrepancies or unexpected behaviors observed during the validation steps of the pipeline are addressed through iterative refinement. This process involves debugging the hardware design, adjusting optimization directives in HLS, and re-validating until the hardware implementation consistently produces the correct results.

Successful validation is achieved when the hardware implementation:

- Produces output data that matches the software algorithm's results within error margins below 5%.
- Operates correctly at the required high frequencies without data loss or corruption.
- Meets the performance targets for processing speed and latency improvements.
- Integrates seamlessly into the existing data processing infrastructure, facilitated by the automated prototyping workflow.

4.1 C simulation and unit testing

After designing all the modules, we initiate the C simulation phase to evaluate their behavior. To automate this process, we implement a cmake target system that first compiles all the necessary libraries and source files. Cmake scripts are used to define build targets, manage dependencies, and configure the build environment.

We develop a unit testing system to verify each module individually. The testing process involves the following steps:

1. **Extraction of test events:** We extract XML files from the emulator's algorithm simulation containing the expected test events, inputs, and outputs for each module across a number of events.
2. **Creation of test targets:** For each module, a test target is created within the cmake system. These targets utilize a test template that is customized for each module by substituting placeholder names with the actual module names.

3. **General Testing Library:** A general testing library developed in C++ is used to facilitate testing. This library consists of two main components:
 - **CLIUtility module:** Responsible for processing input arguments of each test, reading the XML events, and managing command-line interfaces.
 - **IModule library:** Contains a Module_IModule file for each module, which receives the XML events and converts them into the expected input formats required by the module. It includes testing functions such as input conversion, output conversion (from C++ data types to arbitrary-precision types used in HLS), module execution, result comparison, and logging. Additionally, the IModule library includes a threshold-based verification system. A parameterized threshold is provided for each module to account for acceptable discrepancies in the output, arising from fixed-point arithmetic and hardware constraints. The results are verified by comparing them against the expected outputs within this threshold.
4. **Running the tests:** The tests are executed, with each module being fed by the corresponding inputs and producing outputs that are compared against the expected results from the XML files.
5. **Recording results:** Test results are stored in a JSON file that records the events passed, events failed, and their respective inputs and outputs. This detailed record facilitates debugging and further analysis.

4.2 HLS synthesis and co-simulation

Once the unit tests are passed, the synthesis of the HLS modules is performed. [Figure 6](#) shows the detailed parts of the HLS synthesis step. The pipeline includes: 1) **Synthesis:** Using the cmake system and tcl scripts, the HLS synthesis process is automated for each module. If any error is detected during synthesis, the pipeline stops, and the developer is prompted to refactor the code to resolve the issues. 2) **Co-simulation:** For each module, co-simulation is performed using the passed events from the JSON files as inputs. A co-simulation template is used to generate the necessary test benches and scripts. The co-simulation compares the C++ simulation results with the RTL simulation results to verify that the synthesized hardware matches the expected behavior.

4.3 Integration

Following the successful synthesis and verification of the HLS modules, the next critical step in the development pipeline is the integration of the algorithm into the hardware framework and the generation of the FPGA bitstream. This process involves combining our custom HLS design with existing hardware frameworks used at CERN, configuring communication protocols, and automating the build process through scripting.

At CERN, several hardware frameworks are employed to facilitate communication between FPGA boards and manage essential system functions. These frameworks provide standardized interfaces, communication protocols, and

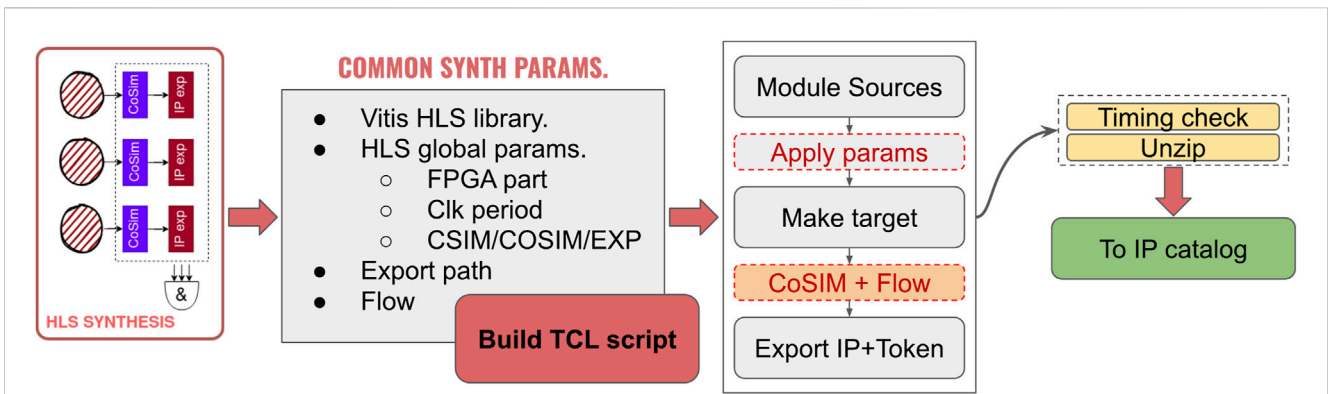


FIGURE 6 Overview of the HLS synthesis and export process. The pipeline begins by setting common synthesis parameters such as FPGA part, clock period, and simulation/export options. A tcl script is generated to automate the process, which includes compiling the module sources, applying parameters, and targeting the design for synthesis. Co-simulation is performed along with the flow execution, and the IP, along with a token, is exported. The token ensures the IP is not rebuilt if the process is repeated. After a timing check and extraction, the IP is added to the catalog.

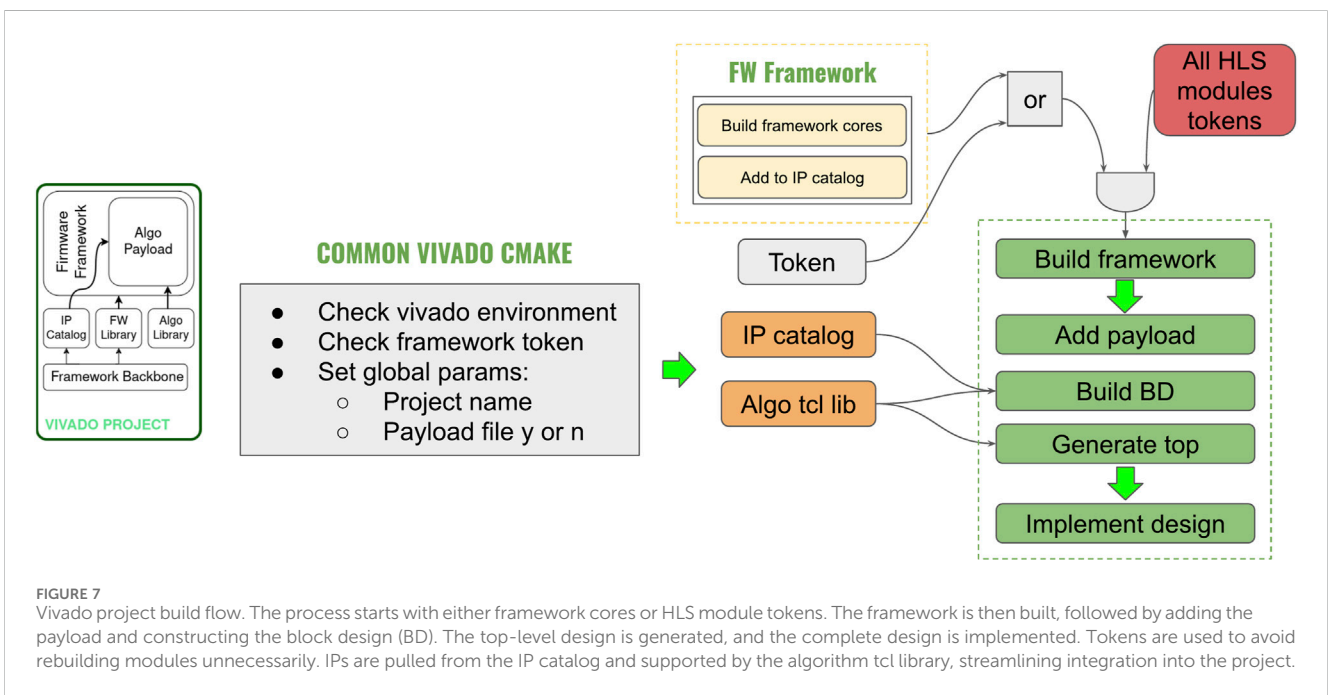


FIGURE 7 Vivado project build flow. The process starts with either framework cores or HLS module tokens. The framework is then built, followed by adding the payload and constructing the block design (BD). The top-level design is generated, and the complete design is implemented. Tokens are used to avoid rebuilding modules unnecessarily. IPs are pulled from the IP catalog and supported by the algorithm tcl library, streamlining integration into the project.

infrastructure components necessary for the operation of the boards. Key features of the hardware framework include:

- **Communication protocols:** Defines the methods for data exchange between boards, including the number of high-speed transceivers (e.g., Multi-Gigabit Transceivers or MGTs) used for high-throughput data links.
- **Control registers:** Sets up registers accessible from the processor system (e.g., embedded ARM cores), allowing for configuration, monitoring, and control of the FPGA logic.
- **Infrastructure components:** Includes clock management, reset logic, and other essential services required for stable operation.

We begin by building the hardware framework, ensuring that it is correctly configured to meet the requirements of our application.

This involves selecting the appropriate number of transceivers, setting up control interfaces, and integrating any necessary communication protocols.

4.4 Integration of HLS design

Once the hardware framework is established, we proceed to integrate our synthesized HLS modules into the design. Figure 7 shows the project building process. This process is automated using tcl scripting within Vivado, enabling us to build the entire block design and manage connections efficiently.

1. **Importing HLS IP cores:** The HLS synthesis process generates IP cores for each module, which are added to the Vivado IP

repository. These IP cores encapsulate the functionality of our algorithm modules and are ready for integration into the block design.

2. Automated block design creation: Tcl scripts are used to instantiate the HLS IP cores and other necessary blocks within the Vivado block design environment. The scripts handle the placement of modules, configure parameters, and establish connections between blocks.
3. Connecting modules: The scripts read the expected ports of both the hardware framework and the HLS payload modules, wiring them accordingly. This includes connecting data paths, control signals, and clock domains as required.

After assembling the block design, a top-level HDL file is generated to serve as the entry point for synthesis and implementation. The top-level file defines the interface of the FPGA design, including all input and output ports, and ensures that all modules are correctly interconnected.

- Tcl scripting for top file generation: We utilize tcl scripts to automate the generation of the top-level file. The scripts parse the port definitions from the hardware framework and the HLS modules, ensuring that all connections are accurately represented.
- Interface definitions: The top-level file specifies the physical interfaces, such as high-speed transceivers, GPIOs, and clock inputs, aligning them with the FPGA's pins and the board's connectors.

With the block design and top-level file prepared, we proceed to the synthesis and implementation stages:

1. Synthesis: The combined design is synthesized to convert the RTL code into a gate-level netlist. This process involves optimizing the logic, mapping the design to the FPGA's resources, and ensuring that timing constraints are met.
2. Implementation: The synthesized netlist undergoes placement and routing to assign the logic elements to specific locations on the FPGA and establish physical connections. This step is critical for meeting timing requirements and optimizing performance.

To achieve optimal performance and ensure the design meets the operational requirements, we apply specific constraints during the implementation:

- Physical constraints: Tcl scripts are used to generate a constraints file (.xdc) that specifies the placement of critical logic elements. This is particularly important when the FPGA is heavily utilized, or when certain paths require minimized propagation delays.
- Timing constraints: We define timing constraints to guide the synthesis and implementation tools in meeting the desired clock frequencies and setup/hold times.
- Resource utilization constraints: Limits are set on the usage of FPGA resources such as LUTs, BRAMs, and DSP slices to prevent over-utilization and ensure reliable operation.

The final step in the pipeline is the generation of the FPGA bitstream:

1. Bitstream file creation: The implementation process produces a bitstream file (.bit) that encapsulates the configured FPGA design.
2. Verification: Before deploying the bitstream, we perform a final verification step to ensure that all constraints are met and that there are no critical warnings or errors.
3. Deployment: The bitstream is then ready to be loaded onto the FPGA for hardware testing and deployment within the CMS Level-1 trigger system.

With the bitstream fully generated, the development pipeline concludes, and the design is ready for deployment and hardware testing. This comprehensive and automated pipeline—from algorithm preparation to bitstream generation—ensures that the hardware-accelerated algorithm is efficiently developed, thoroughly tested, and optimally implemented for operation within the CMS Level-1 trigger system. The automation of the hardware integration process and the application of the constraints, maximizes the performance and reliability of the FPGA implementation, meeting the rigorous demands of HEP data processing at the High-Luminosity Large Hadron Collider.

4.5 From local automation to future CI/CD integration

The entire hardware integration and bitstream generation process is automated within our build pipeline using tcl scripting and the capabilities of Vivado. This automation ensures consistency, reduces the potential for human error, and accelerates the development cycle.

It is important to note that the pipeline described above is developed for local development purposes. This setup allows individual developers or small teams to efficiently test and iterate on the design within their own development environments. The local pipeline provides the flexibility to make rapid changes, test new ideas, and perform detailed debugging without the overhead of a larger, more complex system.

To support repository management, scalability, and collaboration across larger teams, integrating this local pipeline into a Continuous Integration/Continuous Deployment (CI/CD) system is essential. CI/CD systems automate the building, testing, and deployment processes, ensuring that code changes are consistently integrated and validated.

- Consistency: CI/CD pipelines enforce consistent build and test procedures across all developers, reducing integration issues and ensuring that the codebase remains stable.
- Scalability: As the team grows, a CI/CD system can handle multiple developers working concurrently, managing merges, and detecting conflicts early.
- Automated Testing: Automated tests can be run on every commit or pull request, ensuring that new changes do not introduce regressions.
- Traceability and documentation: CI/CD systems provide logs and reports for each build and test cycle, aiding in traceability and compliance with development standards.

TABLE 1 Comparison of key metrics for HLS implementations.

Parameter	Zabototny (2019)	Our work
Device	XC7VX690T	XCVU13P
Frequency (MHz)	160	360
Latency (cycles)	54	50
Latency (ns)	337.5	139
LUT	123,964 (28.6%)	204,300 (11.8%)
FF	112,240 (12.9%)	198,022 (5.7%)
BRAM	720 (24.5%)	274 (10.2%)
DSP	-	204 (1.7%)

- Deployment automation: Streamlines the process of deploying updates to hardware in the field, reducing manual intervention and potential errors.

By integrating the local development pipeline into a CI/CD system, we can leverage tools such as GitLab CI/CD, Jenkins, or Travis CI to automate the entire workflow. This integration ensures that the development process is scalable, maintainable, and aligned with best practices for software and hardware development.

5 Experimental results

This document presents a comparative analysis of the HLS-based implementation of the Overlap Muon Track Finder (OMTF) algorithm on an older *Virtex-7 XC7VX690T* FPGA versus a new implementation on the *Virtex UltraScale+ XCVU13P* FPGA. The new implementation incorporates significant architectural improvements and includes additional logic due to the requirements of the Phase-II upgrade, which represent an increase in the volume and complexity of data processed. Moreover, the new implementation includes extrapolation logic leveraging DSPs for multiplication operations Leguina (2023). The comparison is summarized in Table 1.

The reductions in latency and resource usage, alongside the ability to handle more data, mark significant advancements in the algorithm's design. Despite handling more data, the new implementation achieves a real-time latency improvement by a factor of:

$$\text{GainFactor} = \frac{337.5\text{ns}(\text{Virtex} - 7)}{139\text{ns}(\text{VirtexUltraScale+})} \approx 2.43$$

This demonstrates the substantial advantage of the newer FPGA architecture combined with enhanced algorithm design.

6 Discussion

The results are expected to demonstrate that hardware acceleration using High-Level Synthesis (HLS) significantly improves processing speeds by 2.25 and reduces latency by factor 2.43 in a high-frequency application such as the muon track-finding algorithm of the CMS Level-1 trigger system. The successful

implementation confirms that HLS is an effective tool for translating complex algorithms into efficient hardware designs suitable for real-time applications in HEP. The performance gains validate the optimization techniques applied, such as parallel processing and pipelining, highlighting their impact on enhancing computational efficiency without sacrificing accuracy.

The advancements achieved through this work have substantial implications for HEP research. By enabling real-time data processing with improved accuracy, researchers can explore new physics phenomena that require rapid and precise measurements, such as the detection of rare particles or events occurring at high luminosities. The adoption of hardware acceleration expands the capabilities of trigger systems, potentially leading to more effective data collection strategies and enhancing the overall scientific output of experiments like CMS.

The methodologies developed in this study demonstrate scalability to other algorithms and detector systems within HEP. Future work may involve extending the use of HLS and hardware acceleration to additional components of the trigger system or other experiments facing similar computational challenges. Exploring hardware/software co-design approaches and integrating machine learning algorithms into hardware accelerators are promising directions. Continued development of HLS tools and techniques will further simplify the design process, making hardware acceleration more accessible to a broader range of researchers.

This study highlights the transformative potential of hardware acceleration in HEP, showcasing how high-level synthesis can effectively bridge the gap between complex software algorithms and hardware implementation. By harnessing HLS, we achieve significant improvements in processing speed and latency reduction for the muon track-finding algorithm in the CMS experiment. The successful application of these techniques not only enhances current data analysis capabilities but also sets the stage for future innovations, enabling physicists to tackle increasingly complex challenges and drive groundbreaking discoveries in the field.

Data availability statement

The raw data supporting the conclusions of this article will be made available by the authors, without undue reservation.

Author contributions

PL: Conceptualization, Data curation, Formal Analysis, Investigation, Methodology, Software, Validation, Visualization, Writing—original draft, Writing—review and editing. SF: Funding acquisition, Project administration, Supervision, Writing—original draft, Writing—review and editing.

Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. This work and the authors are partially supported by ERC grant (INTREPID, 101115353) and the Ministerio de Ciencia e Innovación (Spain) with the project PID2020-113341RB-I00. Funded by the European Union.

Acknowledgments

We would like to thank the Level-1 Trigger Project of the CMS experiment for their support to the authors. Special thanks to K. Bunkowski, P. Fokow, M Konecki, and W. Zabolotny for their insight in the HLS implementation of the OMTF algorithm, as well as to M. Batchis and his team for facilitating the validation of the algorithm.

Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that Generative AI was used in the creation of this manuscript. Generative AI has been used to help

structure the paper and the Generative AI output has been checked for factual accuracy and plagiarism.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Author disclaimer

Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- Aad, G., Abbott, B., Abbott, D., Abed Abud, A., Abeling, K., Abhayasinghe, D., et al. (2021). The ATLAS fast TracKer system. *JINST* 16, P07006. doi:10.1088/1748-0221/16/07/P07006
- Aggleton, R., Ardila-Perez, L., Ball, F., Balzer, M., Boudoul, G., Brooke, J., et al. (2017). An FPGA based track finder for the l1 trigger of the CMS experiment at the high luminosity LHC. *J. Instrum.* 12, P12019. doi:10.1088/1748-0221/12/12/P12019
- CMS Collaboration (2020). The phase-2 upgrade of the CMS level-1 trigger. *Tech. Rep.* CERN, Geneva. CERN-LHCC-2020-004, CMS-TDR-021.
- Evans, L., and Bryant, P. (2008). LHC machine. *J. Instrum.* 3, S08001. doi:10.1088/1748-0221/3/08/S08001
- Ghanathe, N. P., Madorsky, A., Lam, H., Acosta, D. E., George, A. D., Carver, M. R., et al. (2017). Software and firmware co-development using high-level synthesis. *JINST* 12, C01083. doi:10.1088/1748-0221/12/01/C01083
- Han, S., Mao, H., and Dally, W. J. (2016). "Deep compression: compressing deep neural networks with pruning, trained quantization and Huffman coding," in *4th international conference on learning representations (ICLR)*.
- Husejko, M., Evans, J., and Rasteiro Da Silva, J. C. (2015). Investigation of high-level synthesis tools' applicability to data acquisition systems design based on the cms ecal data concentrator card example. *J. Phys. Conf. Ser.* 664, 082019. doi:10.1088/1742-6596/664/8/082019
- Leguina, P. (2023). Firmware implementation of a displaced muon reconstruction algorithm for the phase-2 upgrade of the cms muon system. *J. Instrum.* 18, C12005. doi:10.1088/1748-0221/18/12/C12005
- Nane, R., Sima, V.-M., Pilato, C., Choi, J., Fort, B., Canis, A., et al. (2016). A survey and evaluation of FPGA high-level synthesis tools. *IEEE Trans. Computer-Aided Des. Integr. Circuits Syst.* 35, 1591–1604. doi:10.1109/TCAD.2015.2513673
- Radburn-Smith, B. (2022). "Overview of the HL-LHC upgrade for the CMS level-1 trigger," in *Proceedings of science (ICHEP2022)*, 639. doi:10.22323/1.414.0639
- Zabolotny, W., and Byszuk, A. (2016). Algorithm and implementation of muon trigger and data transmission system for barrel-endcap overlap region of the CMS detector. *J. Instrum.* 11, C03004. doi:10.1088/1748-0221/11/03/C03004
- Zabolotny, W. M. (2019). "Implementation of OMTF trigger algorithm with high-level synthesis,". *Photonics applications in astronomy, communications, industry, and high-energy physics experiments*. Editors R. S. Romaniuk and M. Linczuk (Wilga, Poland: International Society for Optics and Photonics SPIE), 11176. doi:10.1117/12.2536258