


Hunting long-lived particles with intelligent triggers at the High-Luminosity LHC

The Higgs boson discovery at the Large Hadron Collider (LHC) closed a major chapter in particle physics but left profound questions. What is dark matter? Why are neutrinos so light? And why does matter dominate antimatter? One possibility is new long-lived particles (LLPs) that travel measurable distances before decaying, producing signatures easily missed by traditional triggers.



At the High-Luminosity LHC (HL-LHC), collisions will occur at unprecedented rates, requiring ultra-fast trigger systems to decide which events to store. Current triggers favour prompt particles and struggle with displaced tracks or delayed signals typical of LLPs. INTREPID addresses this by redesigning trigger architecture, embedding machine-learning algorithms and high-speed hardware into early decision stages to capture rare signatures.

Background and objectives

The need for a trigger system

Only a handful of the collisions produced by the LHC contain potentially interesting physics. The trigger system allows real-time decisions on whether a given event of interest should be saved for data analysis or discarded. In the CMS experiment, this selection is achieved by a two-level system: the Level-1 trigger (L1T), consisting of custom hardware processors that receive data from the calorimeter and muon systems, generating a trigger signal with a latency of a few microseconds. The High-Level Trigger is implemented in software and reduces the rate down to ~1 kHz. In view of the HL-LHC, the L1T system of the CMS experiment will be fully replaced to greatly extend the current system's throughput and capabilities, despite the harsher environment. The L1T system (CMS Collaboration, 2020) has been designed to process 63 Tb/s of input bandwidth with state-of-the-art commercial field-programmable gate arrays (FPGAs) and high-speed optical links reaching up to 28 Gb/s, using generic processing cards based on Advanced Telecommunications Computing Architecture (ATCA) technology.

The HL-LHC will dramatically increase the number of proton-proton interactions, delivering far higher luminosity and thus many more opportunities to produce rare processes, including candidates for physics beyond the standard model (BSM). This intense environment comes with challenges—such as extreme pileup

and much higher data rates—but it also offers a unique chance to probe very small interaction probabilities, provided that the trigger can recognise and retain the key events. For projects like INTREPID, the HL-LHC is therefore both a stress test and a major opportunity: by developing smarter, more flexible trigger strategies, it becomes possible to turn the increased luminosity into genuine discovery potential for LLPs and other unconventional signatures.

INTREPID aims first to design and evaluate new trigger strategies that are intrinsically sensitive to displaced and other non-standard signatures, so that long-lived particles are no longer filtered out by assumptions tailored to prompt events. A second core objective is to embed advanced, resource-efficient machine-learning algorithms into ultra-fast hardware platforms, enabling real-time pattern recognition and classification within the strict latency

and bandwidth constraints of the HL-LHC trigger. A third goal is to validate these concepts on realistic benchmark scenarios and integrate them into the trigger chains of major LHC experiments, thereby demonstrably extending their discovery reach for physics beyond the standard model.

Approach and innovations

The projected L1T upgrade muon architecture (CMS Collaboration, 2020) is the starting point of INTREPID. Such architecture consists of 3 layers of increasing complexity, each running sequentially more complex pattern recognition algorithms. The first layer collects the point-like signal (hits) from different subdetectors to generate stubs that are then processed in the second layer, the track finders, which build standalone muon tracks. The last layer, the global muon trigger, correlates

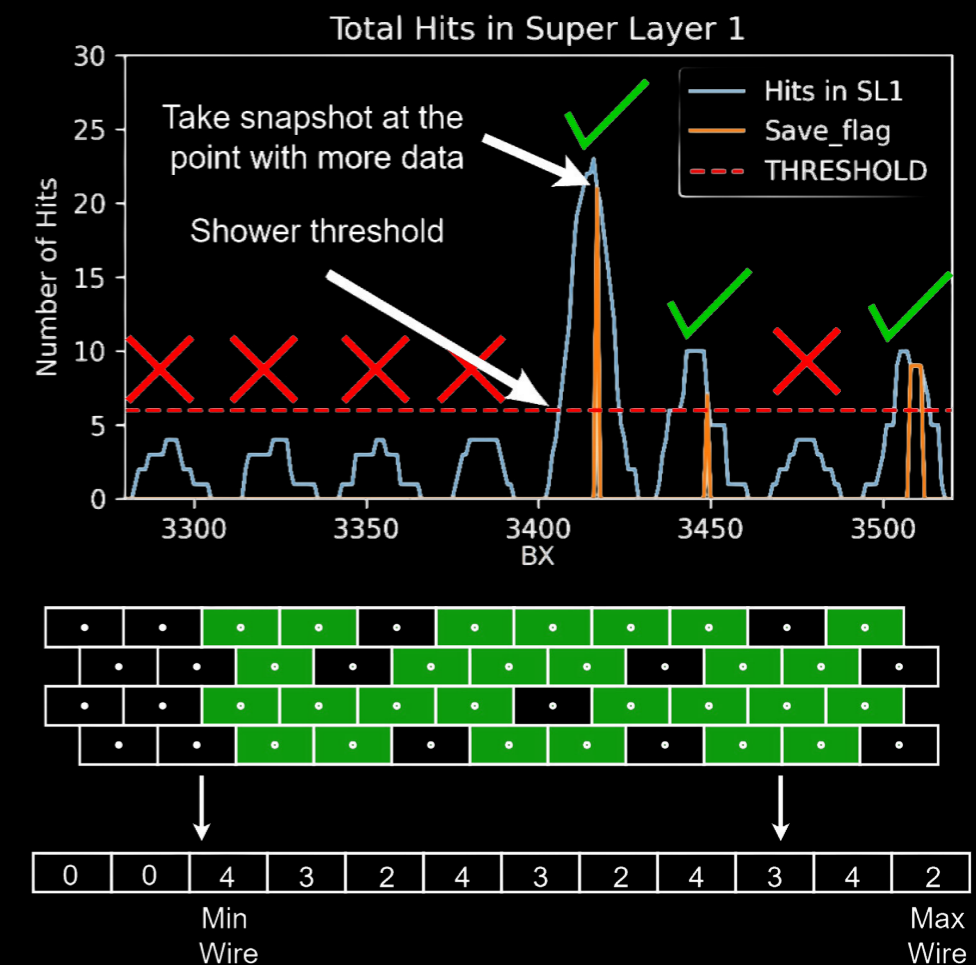


Figure 1: Muon hit count evolution for events with and without showers (above); if a shower is detected, a snapshot of the chamber occupancy is saved and transmitted to the next level (below).




Figure 2: High-momentum muon filter logic.

the muon-only information (standalone tracks and stubs) with tracker tracks for ultimate precision. The first 2 are critical for the quests of LLPs.

A low-latency muon-shower tagging for the L1T of the CMS experiment

Being able to detect, reconstruct, and tag electromagnetic or hadronic showers in the muon system would be a key (Prado, Vico, and Folgueras, 2025) feature of the trigger system for several BSM searches. Such signals may appear both on highly energetic muons that could be produced in the decay of heavy resonances or in the decay of very LLPs that decay to hadrons in the muon system.

To detect such signals, we've designed an approach consisting of 2 stages that can be integrated into the upgraded trigger chain. First, a 'shower detection' stage continuously counts how many hits appear in a region of the barrel muon chambers over a short time window, and when this count exceeds a chosen threshold, it records a compact description of the shower—when it occurred, how many hits it produced and how widely they are spread.

A second stage then combines this information with the usual muon

trigger signals from neighbouring detector stations and checks whether a reconstructed muon track points towards the shower region, promoting consistent patterns to dedicated 'high-momentum muon shower' candidates, see Figure 2.

Tests on simulated data enriched with energetic muons show that this strategy can identify showers with good efficiency while strongly rejecting background activity. In typical settings, more than half of genuine showers are correctly tagged, and a large fraction of spurious shower candidates are discarded once the directional filter is applied, meaning that many more high-momentum muons can be retained at trigger level without overwhelming the system with noise.

Accelerating trigger algorithms with High-Level Synthesis

Our team has also recently shown how High-Level Synthesis (HLS) can bridge the gap between complex C/C++ algorithms and their real-time implementation on modern FPGA platforms (Leguina and Folgueras, 2025). Using the Overlap Muon Track Finder (OMTF) as a case study, we describe how muon trajectories in the challenging barrel-endcap transition region can be reconstructed by matching detector hits from several muon subsystems to pre-defined reference patterns corresponding to different transverse momenta and charges.

The paper presents this as a complete design and development workflow that is accessible to physicists. The OMTF algorithm is refactored for streaming data, heavily parallelised and pipelined using HLS directives, and supported by careful memory optimisation, enabling many pattern evaluations to run in parallel without violating the stringent latency budget of the Phase-2 trigger.

An automated local pipeline connects C++ emulation, HLS synthesis, co-simulation and FPGA integration, enabling rapid iteration and robust validation against software reference implementations. In practice, the move from a previous Virtex-7 design to an UltraScale+ implementation, combined with these HLS techniques, increases the clock frequency from 160 MHz to 360 MHz and reduces the OMTF latency from 337.5 ns to 139 ns, a factor of roughly 2.4, while handling more complex inputs and using resources efficiently.

Porting event generator to FPGAs accelerators

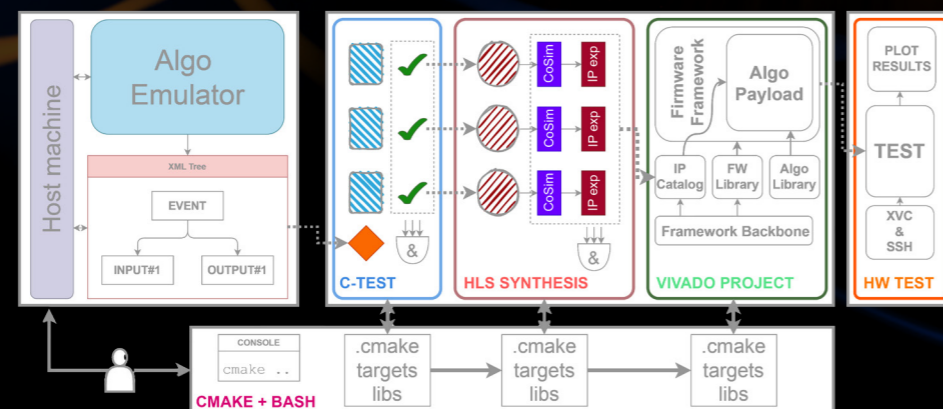
This by-product of the ERC project explores how hardware acceleration can speed up one of the most computationally demanding steps in particle physics simulations: the calculation of matrix elements in event generators such as MadGraph5. We've worked on a proof-of-concept implementation where the matrix element for a single, relatively simple process, $e^+ + e^- \rightarrow \mu^+ + \mu^-$, is ported to an FPGA using HLS, and compare its performance with conventional CPU execution and a GPU implementation (Gutiérrez et al. 2025). Because matrix-element evaluation can dominate the runtime of multi-jet and beyond-the-standard-model simulations, especially at future facilities like the HL-LHC, this step is an ideal target for acceleration.

Adapting MadGraph to FPGAs requires substantial code refactoring: dynamic memory allocations and complex C++ structures are replaced with fixed-size arrays, static buffers, and HLS-friendly data streams, while

loops are unrolled and pipelined to exploit fine-grained parallelism at the hardware level. In the test setup, the FPGA kernel receives random seeds from the host, generates phase-space points internally, computes the matrix elements and returns the results over PCIe, achieving numerical agreement with the CPU reference at the level. For small event batches, transfer and configuration overheads largely cancel the hardware advantage, but for large batches, the FPGA clearly outperforms both CPU and GPU: the FPGA is about 51 times faster than the CPU and 19 times faster than a GPU, while using only a fraction of the available FPGA resources and maintaining double-precision arithmetic. Although this first implementation is limited to one simple process, it demonstrates that FPGAs, programmed through HLS, can play a strong role in accelerating large-scale Monte Carlo production, complementing the trigger- and reconstruction-focused developments within INTREPID.

Impact and future outlook

INTREPID is laying the groundwork for a new generation of trigger systems that can handle the data deluge of the HL-LHC while remaining sensitive to rare signatures such as LLPs and high-momentum muon showers. By combining pattern-recognition algorithms in the muon system, intelligent shower tagging, and hardware-accelerated trigger logic implemented with HLS on modern FPGAs, the project strengthens the discovery potential of CMS precisely where existing triggers are weakest.



References

Gutiérrez, H. et al. (2025) 'Porting MADGRAPH to FPGA using High-Level Synthesis (HLS)', *EPJ Web of Conferences*, 337, 01173. Available at: <https://doi.org/10.1051/epjconf/202533701173>.
 CMS Collaboration (2020) *The Phase-2 Upgrade of the CMS Level-1 Trigger*. CERN-LHCC-2020-004; CMS-TDR-021. CERN. Available at: <https://cds.cern.ch/record/2714892>.
 Leguina, P. and Folgueras, S. (2025) 'Harnessing hardware acceleration in high-energy physics through high-level synthesis techniques', *Frontiers in Detector Science and Technology*, 2, 1502834. Available at: <https://doi.org/10.3389/fdest.2024.1502834>.
 Prado, J., Vico, C. and Folgueras, S. (2025) 'Triggering on muon showers', *Journal of Instrumentation*, 20(03), C03024. Available at: <https://doi.org/10.1088/1748-0221/20/03/C03024>.

PROJECT SUMMARY

The discovery of the Higgs boson at the LHC closes a central chapter of the standard model of particle physics but raises questions about dark matter, neutrino masses, and baryon asymmetry. Answers may involve long-lived particles (LLPs) with non-standard signatures, requiring advanced identification algorithms. Enhancing trigger capabilities to discover LLPs and evidence of beyond-the-standard-model (BSM) physics is therefore a crucial task. INTREPID uses advanced machine learning and ultra-fast processing to improve future trigger systems, potentially revolutionising high-energy physics and benefiting industries that need real-time data processing.

PROJECT LEAD PROFILE

I'm a professor at the University of Oviedo, and I've been working at the CMS experiment since 2009. I'm the author of more than 60 peer-reviewed scientific publications. I've been invited to present my research results in almost 20 international conferences, as the CMS collaboration has recognised me as a key expert in SUSY searches, muon performance, and trigger development. During my research career I've participated in 7 research projects funded under competitive national calls.

PROJECT CONTACTS

Santiago Folgueras
 Prof. of Physics
 Department of Physics
 Universidad de Oviedo
 ✉ folguerassantiago@uniovi.es
 🌐 intrepid.uniovi.es

FUNDING



This project has received funding from the European Research Council (ERC) under the European Union's research and innovation programme (Grant agreement No. 101115353).
 Funded by the European Union. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the ERC. Neither the European Union nor the granting authority can be held responsible for them.