Available on CMS information server

CMS CR -2023/022

The Compact Muon Solenoid Experiment

Conference Report

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

27 February 2023 (v4, 19 March 2023)

The Level-1 Global Trigger for Phase-2 Algorithms, configuration and integration in the CMS offline framework

Gabriele Bortolato, Maria Cepeda, Jaana Heikkila, Benjamin Huber, Elias Leutgeb*, Dinyar Rabady and Hannes Sakulin on behalf of the CMS collaboration

Abstract

For the High-Luminosity Large Hadron Collider (HL-LHC) operation, the Compact Muon Solenoid will undergo a significant upgrade and redesign. An upgraded Level-1 Trigger system, based on multiple types of custom processing boards equipped with Xilinx Ultrascale+ Field Programmable Gate Arrays (FPGAs), will exploit fine grained information from the detector subsystems (calorimeter, muon systems and the silicon-strip tracker). The final stage of the Level-1 Trigger, the Phase-2 Global Trigger (P2GT), will receive more than 20 different trigger object collections from upstream systems and will be able to evaluate about 1000 cut-based and machine learning algorithms distributed over up to twelve boards. The P2GT is designed as a modular system with an easily re-configurable firmware, designed to meet the demand of high flexibility required for adapting trigger strategies during operation of the HL-LHC. The algorithms are kept highly configurable and tools are provided to allow their study, verification, and emulation from within the CMS offline software framework (CMSSW) without the need for knowledge of the underlying firmware implementation. A tool has been developed that converts the Python-based configuration used by CMSSW into VHDL for use in the hardware trigger. A prototype firmware for a single Global Trigger board has been developed, which includes de-multiplexing logic, conversion to an internal common object format and distribution of the data over the FPGA. In this framework, 197 algorithms are implemented at a clock speed of 480 MHz. The prototype has been thoroughly tested and verified using the CMSSW emulator. The P2GT is presented with the novel integration within CMSSW and streamlined translation into VHDL code.

Presented at *ACAT2022 21st International Workshop on Advanced Computing and Analysis Techniques in Physics Research*

The Level-1 Global Trigger for Phase-2: Algorithms, configuration and integration in the CMS offline framework

Gabriele Bortolato[1,2], Maria Cepeda[3], Jaana Heikkilä[4], Benjamin Huber[1,5], Elias Leutgeb[1,5], Dinyar Rabady[1] and Hannes Sakulin[1] on behalf of the CMS collaboration

1) CERN, Esplanade des Particules 1, 1211 Geneva, Switzerland
2) Università degli Studi di Padova, Via VIII Febbraio 2, 35122 Padova, Italy
3) CIEMAT, Avda. Complutense 40, 28040 Madrid, Spain
4) University of Zürich, Winterthurerstrasse 190, 8057 Zürich, Switzerland
5) Technical University of Vienna, Karlsplatz 13, 1040 Wien, Austria

E-mail: `elias.leutgeb@cern.ch`

Abstract. For the High-Luminosity Large Hadron Collider (HL-LHC) operation, the Compact Muon Solenoid will undergo a significant upgrade and redesign. An upgraded Level-1 Trigger system, based on multiple types of custom processing boards equipped with Xilinx Ultrascale+ Field Programmable Gate Arrays (FPGAs), will exploit fine grained information from the detector subsystems (calorimeter, muon systems and the silicon-strip tracker). The final stage of the Level-1 Trigger, the Phase-2 Global Trigger (P2GT), will receive more than 20 different trigger object collections from upstream systems and will be able to evaluate about 1000 cut-based and machine learning algorithms distributed over up to twelve boards. The P2GT is designed as a modular system with an easily re-configurable firmware, designed to meet the demand of high flexibility required for adapting trigger strategies during operation of the HL-LHC. The algorithms are kept highly configurable and tools are provided to allow their study, verification, and emulation from within the CMS offline software framework (CMSSW) without the need for knowledge of the underlying firmware implementation. A tool has been developed that converts the Python-based configuration used by CMSSW into VHDL for use in the hardware trigger. A prototype firmware for a single Global Trigger board has been developed, which includes de-multiplexing logic, conversion to an internal common object format and distribution of the data over the FPGA. In this framework, 197 algorithms are implemented at a clock speed of 480 MHz. The prototype has been thoroughly tested and verified using the CMSSW emulator. The P2GT is presented with the novel integration within CMSSW and streamlined translation into VHDL code.

1. Introduction

The Level-1 Trigger (L1T) of the Compact Muon Solenoid (CMS) experiment consists of a pipeline of Field Programmable Gate Array (FPGA) based triggers, which receive low-latency data from the CMS detector. Based on a reconstruction of the final state particles in the event, the L1T decides if an event should be read out or discarded. For the High-Luminosity Large Hadron Collider (HL-LHC) operation [1], the CMS detector will be equipped with a completely new L1T system. With the increased latency budget of $\approx 12.5\,\mu s$, compared to

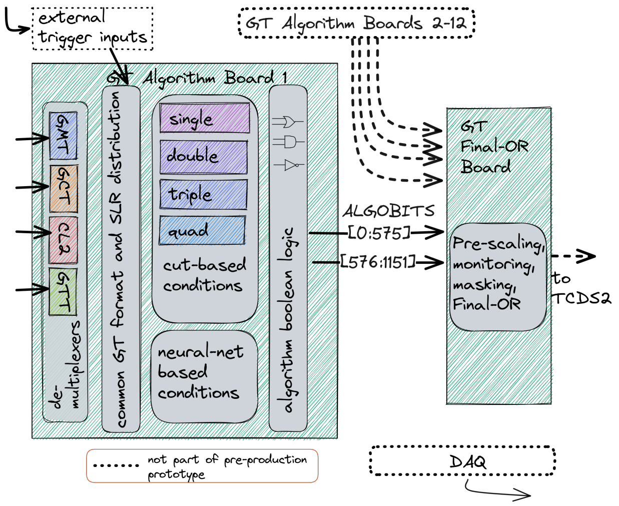

Figure 1. Simplified diagram of the P2GT pre-production prototype architecture logic

the $\approx 3.5\,\mu s$ of the current system, the inclusion of tracker information and high-granularity calorimeter information, the Phase-2 L1T aims to tackle an ambitious physics program at unprecedented levels of luminosity. The final stage of the Phase-2 L1T, the Phase-2 Global Trigger (P2GT), receives high precision inputs from all detector subsystem specific triggers, as well as the Correlator Trigger, which combines information from all detectors taking part in the trigger. The P2GT (Figure 1) consists of up to twelve algorithm- and one Final-OR processing board. It outputs the final trigger decision to the Phase-2 Timing Control and Distribution System (TCDS2). This decision, triggers the readout of the event, which is then further processed by the software based High Level Trigger (HLT), using CPUs and GPUs [2]. Each of the algorithm boards will receive a copy of all upstream system inputs. The P2GT is planned to evaluate up to $\approx$ 1000 configurable cut-based and machine learning algorithms, the so called *Menu*, within a latency budget of $1\,\mu s$ [3]. The outputs of all algorithm boards are sent to the Final-OR board, where they are combined, pre-scaled and monitored [4]. To assess and test the functionality of the P2GT algorithms, a bit-wise compatible emulator has been developed within the CMS software framework (CMSSW). This emulator is going to be used for simulation, firmware emulation, and potentially to seed the HLT algorithms with the L1T results during data taking. To allow the use of the *Menu* configuration in both emulation and firmware, an automated VHDL Writer has been developed which translates the *Menu* written in the CMSSW configuration language into VHDL. The output of the VHDL Writer can be used to directly build the P2GT firmware.

2. The Phase-2 Level-1 Global Trigger Hardware

The P2GT will be implemented on generic ATCA processing boards, developed by the Serenity collaboration [5], equipped with a Xilinx Virtex Ultrascale+ VU13P FPGA. These boards offer $\approx$ 120 input and output links at 25 Gb/s, which are made available using optical transceivers. The FPGAs are produced using stacked silicon interconnection technology [6], allowing for the transfer of signals between stacked silicon dies (referred to as Super Logic Regions (SLRs)). This results in devices with higher performance and density, but creating routing bottlenecks at the SLR boundaries. Designing firmware for such devices imposes additional challenges to meet timing, routing and resource requirements. The boards come with a firmware framework that manages the 25 Gb/s links using a common protocol and interfaces with TCDS2 for control and monitoring. The P2GT firmware is created as a payload and integrated into the framework.

2.1. The Algorithm Boards

In each of the algorithm boards, data arrive in the form of collections usually containing twelve physical objects (e.g. muons) in a time-multiplexed format [7]. Since the P2GT works in a non time-multiplexed way [3], data have to be de-multiplexed and stored in parallel memories (one per collection). After all data of a single event are present, the clock domain is crossed from 360 MHz, the frequency at which data arrive, to 480 MHz, the frequency at which the algorithms are applied. In the same step, the upstream objects are converted to the internal P2GT format. After the conversion, all objects share common bit-widths and scales on all available parameters and the event data are streamed over all SLRs. Each SLR contains an algorithm unit which hosts a subset of the cut-based and machine learning algorithms. After evaluation of the algorithms, a single bit per algorithm, is routed to the output, stating whether the event contains data that passes all conditions in the algorithm. At the output, the results of all algorithms on the chip are re-serialized and transmitted over two links to the Final-OR board. Each P2GT algorithm board evaluates a fraction of the total number of algorithms.

2.2. The Final-OR Board

The results (algorithm bits) of each P2GT algorithm board arrive at the Final-OR board where all algorithm results are combined and the final step of processing is applied. This involves determining the trigger sub-type, applying pre-scales and bunch masks as well as issuing the final trigger decision and trigger sub-types to TCDS2.

3. The Phase-2 Global Trigger Emulator

The P2GT Emulator aims to provide accurate emulation of the hardware, while simplifying the workflow from simulation to hardware implementation in the CMSSW framework. CMSSW provides a modular framework for Event Data (ED) analysis. Individual modules can be interconnected to "paths", which process ED. The modules are: ED Producers (adding data to the event file), ED Filters (reading event data and deciding weather to stop or continue processing the path) and ED Analyzers (reading event data, used to study event properties). The P2GT Emulator is designed to achieve bit-wise compatibility with the P2GT hardware and enables the simulation and emulation of events. The P2GT emulator produces P2GT Objects (P2GT Candidate Producer) from the objects provided by the upstream systems. P2GT algorithms are formed by applying single conditions (ED Filters) to these objects or by performing logical

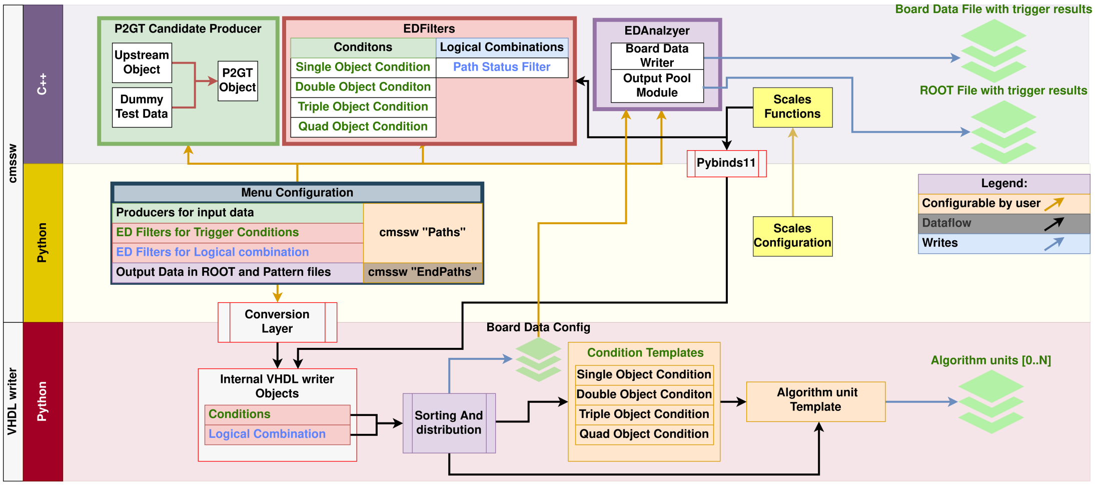

Figure 2. Overview of the modules and interfaces of the Phase-2 Global Trigger Emulator

combinations on the condition results. The result (if an event contains object that passed a condition) is stored (ED Analyzers) with references to the objects that passed the condition in the Event Data file (ROOT file). For hardware evaluation the results can also be stored in a Board Data File which can be compared with the output of the hardware. Python is used to configure the P2GT emulator modules (ED Producers, ED Filters, ED Analyzers). The same configuration can also be utilized by the VHDL Writer to configure the P2GT hardware.

3.1. The Phase-2 Global Trigger Object

The P2GT Object serves as a generic representation for all particle-like objects in the P2GT environment, containing a super-set of all possible parameters. This common P2GT Object allows the P2GT algorithms to be written in a generic way and stores bit-wise compatible values of the parameters available in hardware. Furthermore, it enables the VHDL Writer (discussed in Section 3.3) to process them uniformly, eliminating the need for a specialized translation to handle each particle type.

3.2. Phase-2 Global Trigger Algorithms

The P2GT Algorithms offer a set of generically written conditions (single, double, triple, and quadruple object conditions) which can be configured for all particle-like objects with various cuts. The quantities to cut on can be simple object parameters (like $\mathrm{p_T}, \eta, \phi$) or correlational parameters such as the angular separation between two objects or the invariant mass of a potential mother particle calculated from two detected particles.

3.3. The VHDL Writer

The VHDL Writer provides an automated translation of the *Menu* expressed in the CMSSW Python configuration into VHDL code used by the P2GT hardware. While the CMSSW configuration is written using physical units (i.e GeV), the firmware needs information on the actual hardware resolution of individual parameters. This is achieved by using bit-wise compatible scaling functions which are also used by the emulator. The Python configuration is translated into VHDL by using templates for each condition type as well as the algorithm unit. This is achieved by keeping a similar structure of conditions in hardware as their CMSSW counterparts. Algorithms with correlational conditions require special resources in hardware such as Digital Signal Processors (DSPs) or Block Ram (BRAM) to perform mathematical operations in the tight timing budget. To achieve a high success rate in the implementation of the algorithms in hardware, the VHDL Writer has knowledge of the resource consumption of each individual algorithm. The program not only estimates slice logic used in the FPGA, but also calculates the exact resource consumption of BRAMs and DSPs. These special blocks are a scarce resource on the chips compared to the available slice logic.

4. Results

A test *Menu* with 197 algorithms (77 using correlational cuts which require DSPs) has been written and the workflow from simulation to implementation on a single FPGA with four SLRs has been tested with various distribution strategies. The three chosen strategies were:

- Random: The individual algorithms are distributed at random with no balance in resource consumption of the individual algorithm or number of algorithms on the individual SLR.
- Equal number of algorithms: The algorithms are placed following the order of definition in the CMSSW configuration and distributed equally (by number) over the SLRs.
- Balanced by resource consumption: The algorithms are distributed over the SLRs trying to balance the usage of DSPs.

Type	Random	Equal num. of algos.	Balanced DSP usage
SLR 0			
Algorithms	37	49	49
DSP	216	564	(612) 612
RAMB	36	138	(162) 162
LUT	26097	43976	(44310) 44308
SLR 1			
Algorithms	51	49	49
DSP	708	780	(624) 624
RAMB	162	150	(120) 120
LUT	46121	47273	(44630) 41805
SLR 2			
Algorithms	54	49	49
DSP	1008	528	(624) 624
RAMB	264	168	(192) 192
LUT	54017	41400	(44630) 46653
SLR 3			
Algorithms	55	50	50
DSP	588	648	(660) 660
RAMB	162	168	(150) 150
LUT	17175	42494	(44810) 42853
Succes in HW impl.:	Random	Deterministic	Yes

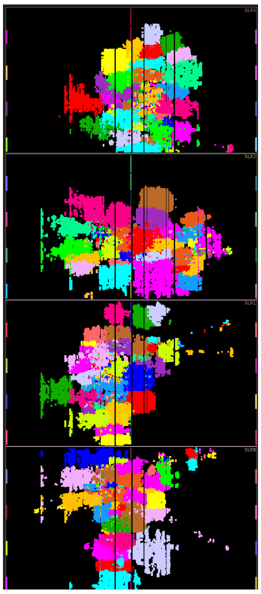

Figure 3. Left: Resource usage of various distributions strategies (in brackets estimated consumption) and their success in hardware implementation, for "Random" a successful example has been chosen; Right: floorplan of the balanced DSP usage strategy with colored algorithms

It is shown that various distribution strategies yield different rates of success in implementing the CMSSW *Menu* in firmware. The random distribution serves as a benchmark, where no optimization is applied. It yields the most inconsistent results (three out of five implementations of the same *Menu* succeeded). Balancing by the number of algorithms yields consistent failure or success depending on the order of algorithms and complexity of the *Menu*. Balancing the DSP usage is a good strategy, as the amount of DSPs needed directly corresponds to the complexity of the algorithm. This results in the highest success rate of individual builds (Figure 3).

5. Summary

The Phase-2 Global Trigger for CMS has been presented along with the functionality of its Emulator in the CMSSW framework. The *Menu*, written in Python, natively serves as the configuration of the emulator. A VHDL Writer tool has been developed to translate the *Menu* to VHDL code that can directly be used to generate the firmware. Various strategies to distribute algorithms over the Super Logic Regions of the FPGA have been compared for a test *Menu* containing 197 algorithms. A distribution strategy balancing the DSP usage yielded consistent firmware build success.

References

[1] Aberle, O *et al* 2020 *High-Luminosity Large Hadron Collider (HL-LHC)* URL `cds.cern.ch/record/2749422`

[2] Jeitler, M 2020 *Journal of Instrumentation* **15** C09009

[3] Sakulin, H *et al* 2023 *Journal of Instrumentation* **18** C01034

[4] CMS collaboration 2020 *The Phase-2 Upgrade of the CMS Level-1 Trigger* URL `https://cds.cern.ch/record/2714892`

[5] Rose, A W *et al* 2019 Serenity: An ATCA prototyping platform for CMS Phase-2 *Proceedings of Topical Workshop on Electronics for Particle Physics — PoS(TWEPP2018)* vol 343 p 115

[6] Xilinx 2015 Xilinx multi-node technology leadership continues with ultrascale+ portfolio "3d on 3d" solutions URL `https://docs.xilinx.com/v/u/en-US/wp472-3D-on-3D`

[7] Hall, G 2016 *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **824** 292–295