

The High-Level Trigger for the CMS Phase-2 Upgrade

Thiago Rafael Fernandez Perez Tomei* on behalf of the CMS Collaboration

Universidade Estadual Paulista,

São Paulo, SP, Brazil

E-mail: Thiago.Tomei@cern.ch

The High-Luminosity LHC (HL-LHC) will usher in a new era in high-energy physics. The HL-LHC experimental conditions entail an instantaneous luminosity of up to $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and up to 200 simultaneous collisions per bunch crossing (pileup). To cope with those conditions, the CMS detector will undergo a series of improvements, in what is known as the Phase-2 upgrade. In particular, the upgrade of the Data Acquisition and of the High-Level Trigger (DAQ-HLT) will have to address a much higher event rate and more complex events. In this paper, we will discuss the aspects of the HLT upgrade, detailing the development of the online reconstruction, the construction, characterisation and timing/rate measurement of a simplified HLT menu, the role of heterogeneous architectures in the HLT and the plan of work and milestones until the beginning of Phase-2.

ARXIV EPRINT: [2211.03684](https://arxiv.org/abs/2211.03684)

41st International Conference on High Energy physics - ICHEP2022

6-13 July, 2022

Bologna, Italy

*Speaker

1. Introduction

The Compact Muon Solenoid (CMS) detector [1] will undergo the so-called Phase-2 upgrade [2] and this will allow the experiment to address the conditions of the High-Luminosity LHC (HL-LHC). All CMS subdetectors will be upgraded, and the Trigger and Data Acquisition system is no exception. The experiment will continue using a two-level trigger system, similarly to what was done for the first LHC era. The Phase-2 Level-1 Trigger (L1T) will be implemented in custom-made electronics and dedicated to analyse the detector information with a maximum latency of 12.5 μ s; the Phase-2 High Level Trigger (HLT) will be implemented as a series of software algorithms, running in a heterogeneous-architecture computing farm. The complete information on the Phase-2 HLT is available in the DAQ–HLT Technical Design Report (TDR) [3].

The Phase-2 HLT will analyse the full L1T output of up to 750 kHz. It will need high selection efficiency for events of interest whilst staying within the envelopes of acceptable output rate / bandwidth and computing farm size dictated by the experimental budget. Table 1 summarises the CMS Phase-2 trigger and DAQ projected running parameters.

CMS detector Peak \langle PU \rangle	LHC	HL-LHC	
	Phase-1	Phase-2	
	60	140	200
L1 accept rate (maximum)	100 kHz	500 kHz	750 kHz
Event Size at HLT input	2.0 MB	6.1 MB	8.4 MB
Event Network throughput	1.6 Tb/s	24 Tb/s	51 Tb/s
Event Network buffer (60 s)	12 TB	182 TB	379 TB
HLT accept rate	1 kHz	5 kHz	7.5 kHz
HLT computing power	0.7 MHS06	17 MHS06	37 MHS06
Event Size at HLT output	1.4 MB	4.3 MB	5.9 MB
Storage throughput	2 GB/s	24 GB/s	51 GB/s
Storage throughput (Heavy-Ion)	12 GB/s	51 GB/s	51 GB/s
Storage capacity needed (1 day)	0.2 PB	1.6 PB	3.3 PB

Table 1: CMS Phase-2 trigger and DAQ projected running parameters, compared to the design values of the current (Phase-1) system.

2. Physics Objects and HLT Conceptual Structure

In the HLT, the reconstruction of the physics objects – leptons, hadronic jets and global event quantities – uses the same algorithms and framework as the offline reconstruction, the CMS Software Framework (CMSSW), the difference being that the online reconstruction has added emphasis in execution speed. The reconstruction algorithms are encapsulated in modules, which are organised in logical sequences called *HLT paths*. An HLT path targets a given final state and comprises both *producer* modules, which calculate a given quantity of interest, and *filter* modules, which interrupt further processing of the event if it does not achieve a certain selection criteria. The HLT paths are then designed around the concept of *early filtering* – slower, more accurate modules are run only if the event is pre-accepted by faster, less accurate ones. The *HLT menu* is then the set of all HLT paths used to collect data. An event is saved for offline analysis if any path accepts it. It is integrated in such a way that computed variables can be reused amongst paths. Since CMSSW has had multithreading capabilities since 2016, the HLT menu allows both parallel event processing and simultaneous execution of modules within a single event.

2.1 Online Reconstruction

The greatest difference of Phase-2 online reconstruction with respect to the current one is the reconstruction of energy deposits and charged particle tracks in the High-Granularity Calorimeter (HGCAL) and the Phase-2 Tracker, which are respectively completely new and redesigned and for the upgraded detector. These subdetectors have very high granularity, and their objects reconstruction is usually based on iterative procedures: Kalman filter with deterministic annealing for the tracking and a pair of two algorithms, CLUsters of Energy (CLUE) and The Iterative CLustering (TICL) [4] for the HGCAL reconstruction. Those algorithms have to be finely tuned for the constraints imposed by the online environment; the online Phase-2 baseline tracking is based on only two iterations, whilst the online HGCAL reconstruction runs four iterations.

In the Phase-2 HLT, the reconstruction of the physics objects follows along the same lines as in Phase-1. All objects are seeded from the results of the L1T reconstruction. Electrons and photons are seeded by energy deposits in both the electromagnetic calorimeter and in HGCAL; extensive identification (ID) techniques are employed to reduce backgrounds, based on shower shapes and E/p (for electrons only). Since the Phase-2 L1T has access to Tracker information, the “L1TkMuon” that seeds the HLT muon reconstruction is more precise than the Phase-1 equivalent, which helps increase the overall reconstruction efficiency. Hadronic jets and energy sums (missing p_T and total jet transverse momentum H_T) are calculated with inputs from the particle flow algorithm in the barrel and from TICL in the endcap. Extensive pileup mitigations are needed for these objects, in view of the HL-LHC conditions; the PUPPI technique [6] was retuned and deployed for the algorithms that reconstruct jets and sums. Finally, hadronically-decaying tau leptons and b-tagged jets are identified using machine learning techniques, the DeepTau [7] and DeepCSV [8] algorithms. Figures 1 and 2 show efficiency curves for different physics objects in the Phase-2 HLT.




Figure 1: Efficiency curves for the physics objects: charged particle tracks (left), muons (middle), and electrons (right). Efficiencies are given as function of the true object parameter, in this case the object p_T . In the case of double-object paths, the efficiency for a single object reconstruction is shown. Efficiencies are measured in adequate signal samples. Full details are available in the DAQ–HLT TDR [3].

3. Simplified HLT Menu

The full Phase-1 HLT menu, as deployed during the 2018 run, comprised approximately 600 HLT paths. The majority of those paths had very low rate, $\mathcal{O}(1)$ Hz; on the other hand, a few paths, like those dedicated to the collection of single, isolated electrons and muons had a much higher




Figure 2: Efficiency curves for the physics objects: missing p_T (left), b-tagged jets (middle) and hadronically-decaying tau leptons (right). Efficiencies are given as function of the true object parameter – usually the object p_T , except for the b tagging path, where it is the invariant mass of 4 b quarks. In the case of double-object paths, the efficiency for a single object reconstruction is shown. Efficiencies are measured in adequate signal samples. Full details are available in the DAQ–HLT TDR [3].

rate allocation, $O(100)$ Hz. Currently, CMS has prepared a *simplified menu* that targets 50% of the Phase-2 rate. It comprises 15 single and double-object based paths, and follows the same overall structure as the Phase-1 menu.

The characterisation of an HLT menu involves measuring its execution time and its output rate. For the rate estimation, we use simulated samples of multijet QCD and electroweak boson events. Rates are calculated as function of the p_T or ID variable threshold, and the overall rate of the menu is calculated summing over all samples. This procedure double-counts events that are selected by more than one path but is adequate for the level of precision intended at this point. Table 2 shows the paths that comprise the simplified menu, together with their p_T thresholds, their L1T seeding, and their rate percentages (for Phase-1) and estimated rates (for Phase-2).

To estimate the timing of the full menu, we extrapolate the measured results of the simplified menu. We do this by comparing the timing of the equivalent of the simplified menu in 2018 to that of the full one; after verifying that the timing distribution is consistent between the two, we derive from that comparison a +50% correction factor to be applied to the Phase-2 simplified menu timing measurement. For the measurement itself, we run the HLT menu over a sample of minimum bias samples skimmed with a selection consistent with that performed by the Phase-2 L1T. We use as reference hardware a node with two AMD EPYC 7502 processors, totalling 64 (128) physical (logical) cores and a computing power of ~ 1680 HS06. The results are shown in Figure 3; the total timing for the simplified HLT menu is measured to be ~ 5.3 s. Two additional considerations have to be made: first, we consider improvements to the algorithms which could lead, with no performance loss, to a faster HLT by a factor of $1.5\times$ – $2\times$. Second, these results have to be extrapolated to consider the computing hardware that will be available during the Phase-2 era.

4. Heterogeneous Computing and Extrapolation to Phase-2

We use the term *heterogeneous computing* to refer to the usage of computing nodes equipped with different boards that are better adapted to a given task. The prototypical example is that of a node equipped with a general-purpose GPU that leverages its high throughput for efficient parallel

Trigger type	Phase-1		Phase-2		
	Threshold [GeV]	% rate	Threshold [GeV]	Rate at $\langle \text{PU} \rangle = 140$ [Hz]	Rate at $\langle \text{PU} \rangle = 200$ [Hz]
Single μ	50	3%	50	155 \pm 6	213 \pm 8
Single μ (isol.)	24	14%	24	943 \pm 32	1 111 \pm 29
Double μ	37, 27	1%	37, 27	27 \pm 1	40 \pm 1
Double μ (isol.)	17, 8	2%	17, 8	113 \pm 11	143 \pm 13
Triple μ	5, 3, 3	0.5%	10, 5, 5	39 \pm 8	48 \pm 8
Single e (isol.)	28	13%	32 (WP1)	609 \pm 27	1 005 \pm 33
			26 (WP2)	664 \pm 47	1 012 \pm 33
Double e	25, 25	1%	25, 25	46 \pm 4	82 \pm 6
Double e (isol.)	23, 12	1%	23, 12	52 \pm 5	104 \pm 9
Single γ	200	1%	187	32 \pm 1	56 \pm 6
Single γ (isol.)	110, EB only	1%	108, EB only	35 \pm 9	52 \pm 7
Double γ	30, 18	2%	30, 23	123 \pm 12	179 \pm 14
Double τ	35, 35	3%	22, 22	106 \pm 18	159 \pm 27
Single jet	500	1%	520	53 \pm 1	76 \pm 1
H_T	1050	1%	1 070	53 \pm 1	74 \pm 1
Missing p_T	120	3%	140	79 \pm 7	228 \pm 20
Multijets	$H_T = 330$	1%	$H_T = 330$	32 \pm 4	48 \pm 5
with b tagging	jets = 75, 60, 45, 40		jets = 75, 60, 45, 40		
Total rate		49%		2 525 \pm 57	3 621 \pm 62

Table 2: Rates and thresholds of the HLT paths that comprise the simplified menu. Paths for collection of events with muons, electrons, photons, and multiple b-tagged jets are very similar to those of Phase-1, whilst those dedicated to regular jets, H_T and missing p_T have only small threshold increases. The path dedicated to two tau leptons has had its thresholds reduced, thanks to the improved capabilities of the L1T.




Figure 3: Measurement of the HLT timing in the reference hardware. HGCAL reconstruction (in dark yellow), and pixel and silicon tracking (in green) comprise the majority of the time spent during the online reconstruction.

processing of data. It must be noted that both the data formats and the algorithm design must be adapted in order to leverage the capability of the heterogeneous hardware.

The usage of heterogeneous hardware is a keystone of the CMS Phase-2 computing strategy, both in the offline and online environments. To start building the expertise needed for the Phase-2 deployment, CMS has adapted part of its online reconstruction code for heterogeneous computing and has commissioned an HLT farm with nodes equipped with GPUs starting from 2022. For the Phase-2 heterogeneous HLT, the algorithms for the HGAL local reconstruction and the Patatrack pixel tracking reconstruction [5] are already under development. Assuming that the price/performance ratio of GPUs follows an evolution similar to that of CPUs, and that CMS will be able to adapt 50–80% of the HLT code to heterogeneous computing, we estimate that the computing power needs in Table 1 will be met with an effective cost of < 1 CHF/HS06 from 2028 onwards.

5. Conclusions

The CMS Phase-2 online reconstruction is at an advanced stage, and a simplified HLT menu based on basic single-object paths has been developed, integrated and characterised. The performance of the HLT paths is very close to that of Phase-1, keeping output rates under control without the need for large increases in p_T thresholds. The timing structure of the menu is understood, and we show it is possible to meet the overall constraints of the Phase-2 system under two conditions: improving the speed of the HLT algorithms by a small factor and adapting the majority of the code to run on heterogeneous hardware. We now focus on those two improvements, as well as expanding on the foundation provided by the simplified menu into the real HLT menu to be deployed in Phase-2.

Acknowledgments

The author would like to thank the São Paulo Research Foundation (FAPESP) for the travel funding for this conference, under Grant No. 2018/25225-9.

References

- [1] CMS Collaboration, *The CMS Experiment at the CERN LHC*, *JINST* **3**, S08004 (2008).
- [2] CMS Collaboration, *Technical Proposal for the Phase-II Upgrade of the CMS Detector*.
- [3] CMS Collaboration, *The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger*.
- [4] A. Di Pilato *et al.*, *Reconstruction in an imaging calorimeter for HL-LHC*, *JINST* **15** (2020) no.06, C06023 [arXiv:2004.10027].
- [5] A. Bocci *et al.*, *Heterogeneous Reconstruction of Tracks and Primary Vertices With the CMS Pixel Tracker*, *Front. Big Data* **3** (2020), 601728 [arXiv:2008.13461].
- [6] CMS Collaboration, *Pileup mitigation at CMS in 13 TeV data*, *JINST* **15**, no.09, P09018 (2020) [arXiv:2003.00503].
- [7] CMS Collaboration, *Identification of hadronic tau lepton decays using a deep neural network*, *JINST* **17**, P07023 (2022) [arXiv:2201.08458].
- [8] CMS Collaboration, *Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV*, *JINST* **13**, no.05, P05011 (2018) [arXiv:1712.07158].