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

Thiago R. F. P. Tomei, for the CMS Collaboration

SPRACE, Unesp — Thiago.Tomei@cern.ch

Overview

The High Luminosity LHC (HL-LHC) experimental conditions entail an ultimate instantaneous luminosity of $7.5 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and up to 200 simultaneous collisions per bunch crossing (pileup). In order 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 (DAQ) and of the High-Level Trigger (HLT) will have to address a much higher event rate and more complex events. We discuss the HLT upgrade: the development of the online reconstruction, the construction and timing/rate measurement of a simplified HLT menu, the role of heterogeneous architectures in the HLT, and the development outlook until the beginning of Phase-2.

Triple Challenge of the Trigger

CMS will continue to feature a two-level trigger system for the HL-LHC. A synchronous, hardware-based Level-1 Trigger (L1T), operates on dedicated data streams; it is followed by an asynchronous, software-based High-Level Trigger (HLT) that leverages the full detector granularity whilst running on a computing farm composed of heterogeneous commodity hardware. The HLT has to simultaneously optimise for three goals.

- ☐ Select events of interest with high efficiency, by deploying both generalist and more specialised selection algorithms.
- ☐ Control the data acquisition rate, fitting within the limits envisioned in the CMS computing model.
- ☐ Process the data stream selected by the L1T in its entirety, within the capacity of the computer cluster where it runs.

	LHC	HL-LHC	
CMS detector	Phase-1	Phase-2	
Peak ⟨PU⟩	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

HLT Objects and Menu Construction

Goal: achieve or exceed the Run-2 (⟨PU⟩ = 34) performance. Online reconstruction of basic physics objects (leptons, jets, energy sums) with:

- ☐ Optimised identification criteria, algorithm speed constraint.
- ☐ Good response and resolution with respect to truth level.

Simplified HLT menu: paths based on simple object combinations.

- ☐ Comprises the highest-rate algorithms in the Run-2 HLT menu, adding up to roughly 50% of its total rate.
- ☐ Target rate achieved with reasonable kinematic thresholds.

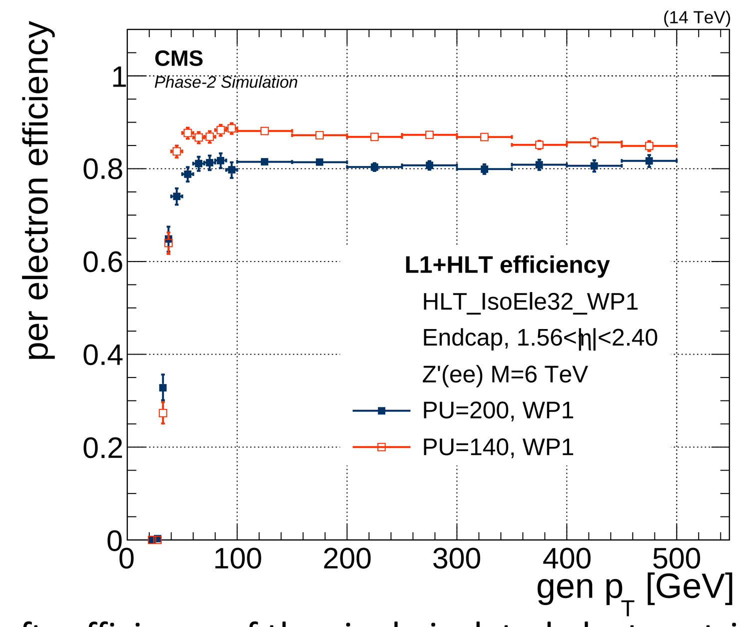

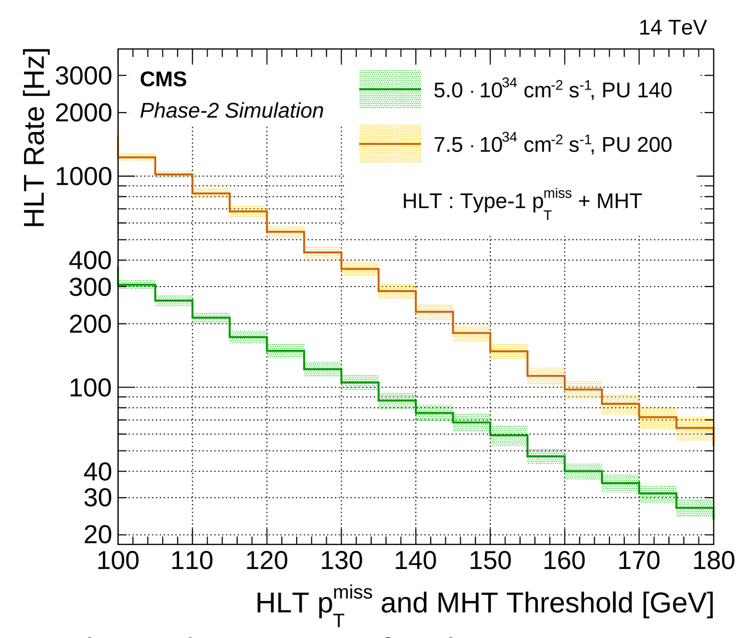

Left: efficiency of the single isolated electron trigger, in the endcap region, for the same $p_T > 32$ GeV threshold used in 2018. Right: rate of the p_T^{miss} trigger, as function of the threshold. The equivalent trigger in 2018 had a rate of 53 Hz at $\mathcal{L} = 1.8 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ for a threshold of $p_T^{\mathrm{miss}} > 120$ GeV.

Simplified Menu Timing Measurement

- ☐ Realistic test conditions:
  - ☐ Benchmark node: two AMD EPYC™ 7502 processors.
  - ☐ Full usage: all cores + multithreading, total 128 logical cores.
  - ☐ Input: L1T-accepted minimum-bias events.
- ☐ Timing reduction: $\mathcal{O}(100\,\mathrm{s})$ (offline) $\Rightarrow \sim 5$ s (HLT).

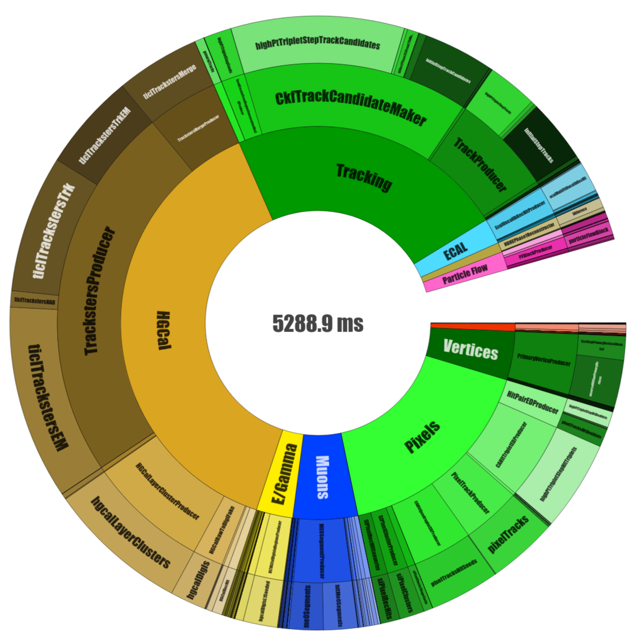

Element	Time	Fraction
B tagging	0.4 ms	0.0 %
E/Gamma	158.4 ms	3.0 %
ECAL	110.9 ms	2.1 %
Framework	0.0 ms	0.0 %
HCAL	41.6 ms	0.8 %
HGCal	2030.5 ms	38.4 %
HLT	0.7 ms	0.0 %
I/O	0.4 ms	0.0 %
Jets/MET	32.1 ms	0.6 %
LIT	2.5 ms	0.0 %
Muons	280.9 ms	5.3 %
other	232.8 ms	4.4 %
Particle Flow	78.9 ms	1.5 %
Pixels	902.3 ms	17.1 %
Tracking	1204.5 ms	22.8 %
Vertices	211.9 ms	4.0 %
total	*5288.9 ms*	*100.0 %*

Average processing time per event, and composition of the HLT reconstruction weighted by the relative processing time, for the HLT menu running over L1T-accepted minimum-bias events with ⟨PU⟩ = 200.

Heterogeneous Computing

Increase performance and energy efficiency not by merely increasing the number of the same-kind processors, but by employing different co-processors specifically designed to handle specific tasks in parallel.

- ☐ CMS software framework extended to be able to offload asynchronous work outside of the framework scheduler.
- ☐ Automatic job configuration based on node capabilities (e.g. presence of a GPU card).
- ☐ Minimise cost of memory transfers and data formatting: usage of *Structures of Arrays* pattern.

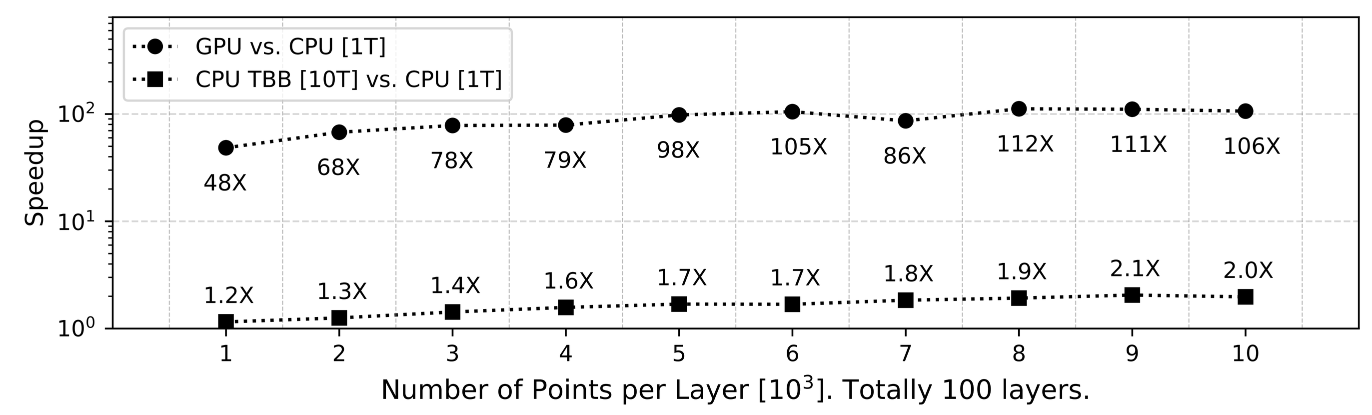

Speedup factors of the GPU and for multi-threaded CPU with Threading Building Blocks (TBB) for the the CMS endcap calorimeter local reconstruction algorithm CLUE, compared to single-threaded CPU operation.

Conclusions and Outlook

The CMS collaboration has demonstrated a first implementation of realistic software reconstruction algorithms, able to cope with the challenges of the Phase-2 environment and detector and optimised to be able to run online as part of the HLT. Further improvements in algorithm speed, usage of heterogeneous environments, as well as the extension of the simplified menu, will be achieved in the coming years to ensure the system readiness for the HL-LHC era.

Acknowledgements

This material is based upon work supported by the São Paulo Research Foundation (FAPESP) under Grant No. 2018/25225-9. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of FAPESP.

References

The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger, Technical Design Report. `https://cds.cern.ch/record/2759072/` (2021).