

The CMS High Level Trigger

Stéphanie Beauceron*

on behalf of the CMS Collaboration

Institut de Physique Nucléaire de Lyon, France

E-mail: stephanie.beauceron@cern.ch

The CMS experiment has been designed with a two-level trigger system: the Level 1 Trigger, implemented using FPGA and ASIC technology, and the High Level Trigger (HLT), running a streamlined version of the CMS offline reconstruction software on a cluster of commercial rack-mounted computers, comprising thousands of CPUs. The design of a software trigger system requires a tradeoff between the complexity of the algorithms running online, the output rate, and the selection efficiency. The complexity is limited by the available computing power, while the rate is constrained by the offline storage and processing capabilities. The main challenge during 2012 is the fine-tuning and optimisation of the algorithms, in order to cope with the increasing LHC pile-up without impacting the physics performance.

We present a review of the performance of the main triggers used during the 2012 data taking, ranging from simple single-object selections to more complex algorithms combining different objects, and applying analysis-level reconstruction and selection. We discuss how the increasing LHC pile-up is affecting their performance, and how these effects are being mitigated.

36th International Conference on High Energy Physics

4-11 July 2012

Melbourne, Australia

*Speaker.

1. CMS High Level Trigger

The trigger system of the Compact Muon Solenoid (CMS) experiment is based on a hardware level (Level 1) coupled to a software level named High Level Trigger (HLT). Level 1 is based on FPGA and custom ASIC technology reducing the 40 MHz design collision rate to 100 kHz. An overview of the system and its performance can be found elsewhere[1]. The High Level Trigger is made of software distributed on a filter farm of commercial rack-mounted computers, comprising over 13000 CPU cores. The events recorded by the trigger system of CMS correspond to a large variety of use cases. The flux of events was split into various streams depending on the event content and the use case. Most of the bandwidth of the data was for physics analysis and data quality monitoring. We were also saving partial events for calibration and trigger studies purposes. CMS had been recording events since 2009. Over the years and the progress of LHC, HLT was operated with increasingly strict selection criteria. 2012 data taking was challenging due to the high luminosity resulting in a large number of pile up events.

2. Challenges of 2012

The peak luminosity in 2011 was 3.5 Hz per nano-barn which increased to 6.5 Hz per nano-barn at the beginning of 2012. The factor two increase in luminosity was not the major issue. The average number of pile up events over 5 inverse femto-barn collected in 2011 at 7 TeV was around 6, with a tail of up to 25 events. For the same luminosity collected in 2012 at 8 TeV, the average was around 18 with the tail extended up to 45. The major challenge was to deal with an increase of luminosity together with the resulting increase in pile up. Pile up affects mainly triggers with isolation requirements. A new algorithm for computing the isolation was deployed during the first part of 2012 data taking: in the turn on curve for offline identified muons with transverse momentum around 25 GeV/ it recovered a few percent of efficiency, with a plateau at 90%. The purity for isolated muons identified at HLT with transverse momenta greater than 24 GeV/c was of the order of 85% with respect to the standard offline selection of CMS, and was stable with respect to pile up.

3. Particle Flow at HLT

Offline CMS reconstruction software moved to particle flow reconstruction: it used the fine granularity/segmentation of the detector in order to reconstruct each particle individually and to clustered them in larger objects as needed[3]. Such software relied on a very good tracking reconstruction and could be CPU intensive. In 2011, the particle flow algorithm was mostly in use for tau identification at HLT. In 2012, it was extended to jets, missing transverse energy so that most of the software used for offline reconstruction was being used on the online farm. Such improvements eased the understanding of the trigger performances and improved the efficiencies with respect to offline selection. Precise implementations of the tracking allowed for more refined techniques for pile up mitigation. Extensive studies together with optimized particle flow algorithms allowed to obtain a flat dependence of the cross section of the triggers with respect to instantaneous luminosity - and therefore to pile up - as illustrated in figure 1.




Figure 1: Overall cross section of the HLT menu with respect to instantaneous luminosity. HLT menu was divided into a “core” of trigger paths promptly reconstructed and trigger paths with delayed reconstruction (Parking). Even at high pile up the cross section was rather flat.

4. CPU Usage

Tracking at HLT had become the major CPU consumer. In 2012, the HLT farm was extended in order to cope with luminosities higher than 5 Hz per nano-barn. As illustrated in figure 2, with the current HLT menu and considering a full L1 rate of 100 kHz as input, the full farm would be able to cope with up to 8 Hz per nano-barn instantaneous luminosity, with around 36 pile up events, in line with the LHC scope.

5. HLT Object Performances

The performance of the trigger paths was constantly monitored and corrections were applied when required. During 2011, CMS Electromagnetic Calorimeter irradiation in the endcap appeared to be larger than expected. Loss of transparency of endcap ECal crystals induced a dilution of the efficiency turn on curve with an increase of the transverse momentum for which the plateau was reached. In 2012, a procedure was put in place in order to compensate for the loss of transparency in the endcap. As one can see in figure 3, the light monitoring corrections brought the endcap turn on curve closer to the one in barrel where losses of transparency can still be neglected. During the first days of running in 2012, it was also noticed that some muon+tau triggers used for searches for the standard model Higgs decaying into two taus, were having lower efficiencies than expected. By a better tuning of the track isolation variables as well as taking into account some detector




Figure 2: Time per event used by machines equipped with dual E5430 Xeon quad-core processors, running 7 HLT processes, as a function of instantaneous luminosity. The extrapolation (dotted line) was performed to indicate current limit of the farm.

differences between barrel and endcap, the trigger path efficiencies were increased from a plateau around 75% in the endcap to around 90%.

Table 5 lists the thresholds and rates for the major triggers used in the HLT menu at that time. All along the year, the thresholds for single isolated muons and electrons triggers had been kept at 24 GeV/c and 27 GeV/c, respectively, for rates of 21 Hz and 59 Hz at 6 Hz per nano-barn of instantaneous luminosity.

6. Data Parking and Data Scouting

As the LHC would stop delivering collisions in 2013, CMS was recording a higher rate of events than could be promptly reconstructed. This was called “Data Parking”[2]. The recorded events served many purposes, for example Vector Boson Fusion physics, multijet final states, low invariant mass of muon pairs (B-physics), tau pairs with 3 prong hadronic decays. 5% of these data were promptly reconstructed to ensure their quality. On top of 350 Hz of physics events immediately reconstructed, 300 Hz of events were parked for future reconstruction.




Figure 3: Turn on curve of a single electron trigger with momentum greater than 33 GeV/c for the barrel ($|\eta| < 1.5$) [triangle] and for the endcap on which light monitoring corrections have been applied [full circle] or not [open circle]. Applying such corrections brought the turn on curve closer to the barrel one (no corrections are needed).

On top of the data parking, CMS put also in place “Data Scouting”: trying quickly to notice any new physics in a regime where Standard Model processes dominate largely. A trigger path was designed to keep events with HT, the transverse energy sum of reconstructed objects in an event, greater than 250 GeV. This trigger would have a very large rate of 1 kHz and must therefore be pre-scaled. The content of the events was reduced to contain only HLT jets and not the usual raw data. With this reduced event content, the bandwidth usage (rate \times event size) was kept manageable. An analysis was performed in the data quality framework and any deviation would be a hint for CMS to dedicate the bandwidth to fully record these events and perform higher level studies. The data scouting was put in place for searches like dijet final state, razor and α_T searches.

(Unprescaled) Object	Trigger Threshold (GeV/c)	Rate (Hz)	Physics
Single Muon	40	21	Searches
Single Isolated muon	24	43	Standard Model
Double muon	(17, 8) [13, 8 for parked data]	20 [30]	Standard Model / Higgs
Single Electron	80	8	Searches
Single Isolated Electron	27	59	Standard Model
Double Electron	(17, 8)	8	Standard Model / Higgs
Single Photon	150	5	Searches
Double Photon	(36, 22)	7	Higgs
Muon + Ele cross-trigger	(17, 8), (5, 5, 8), (8, 8, 8)	3	Standard Model / Higgs
Single PFJet	320	9	Standard Model
QuadJet	80 [50 for parked data]	8 [100]	Standard Model / Searches
Six Jet	(6 x 45), (4 x 60, 2 x 20)	3	Searches
MET	120	4	Searches
HT	750	6	Searches

Table 1: Summary of threshold/rate of the HLT menu for luminosity up to 7 Hz per nano-barn.

7. Conclusion

The High Level Trigger system of CMS selected event for various use cases, from detector maintenance to specific events for physics searches. With the challenge of increasing pile up in 2012, the HLT had shown a remarkably stable behavior thanks to many improvements in code and tuning of corrections. The current HLT farm was able to cope with a processing time of 165 ms per event at 100 kHz of L1 input rate, thus able to sustain an instantaneous luminosity up to 8 Hz per nano-barn without changes to the current HLT settings. All efficiencies for the various trigger paths were high, turn-on curves were sharpened and rates were stable.

References

- [1] CMS Collaboration, “The CMS experiment at the CERN LHC”, JINST 3 S08004 (2008).
- [2] CMS Collaboration, “Data Parking and Data Scouting at the CMS Experiment”, CMS-DP-2012-022 (2012).
- [3] CMS Collaboration, “Commissioning of the Particle Flow Event Reconstruction with the First LHC Collisions Recorded in the CMS Detector”, CMS-PAS-PFT-10-001 (2010).