PROCEEDINGS OF SCIENCE

CMS High Level Trigger Performance for Run 3

PoS(EPS-HEP2023)517

Sanu Varghese[a,*] **on behalf of the CMS collaboration**

[a] *Institute of Physics Bhubaneswar, Sachivalya Marg, Bhubaneswar, India*

E-mail: sanu.varghese@cern.ch

The CMS experiment at CERN uses a two-stage trigger system to filter and store events of physics importance: a hardware-based Level 1 (L1) trigger that uses fast electronics (based on FPGAs and ASICs) to process data in a pipeline fashion at 40 MHz with an output rate of around 110 kHz and a software-based High-Level Trigger (HLT) run on computer farms with an average output rate of around 5 kHz. Many novel trigger algorithms, coupled with technological developments such as heterogeneous computing in GPUs were developed to cope with the increased centre of mass energy, instantaneous luminosity and the physics needs of Run 3. This talk summarises the performance of the CMS HLT during the first year(s) of Run 3.

EPS-HEP 2023,
25 Aug 2023
Hamburg , Germany

*Speaker

1. Introduction

The CMS experiment [1] employs a two-level trigger system: a hardware-based Level 1 Trigger (L1T) reduces the event rate from 40 MHz to 110 kHz, and a software-based High Level Trigger (HLT) further decreases it to about 5 kHz for offline reconstruction, comprising around 2 kHz for prompt processing which usually starts within 48 hours of data collection and around 3 kHz for delayed/opportunistic processing (Data Parking). Additionally, the Data Scouting Strategy enables event collection at a higher rate of 30 kHz, with reduced content, solely relying on HLT information without offline processing.

After a shutdown of more than 3 years, the LHC resumed operations on July 5, 2022, with a higher center of mass energy of $\sqrt{s}$ = 13.6 TeV and increased pileup. Several developments were made to the trigger system to cope with these conditions and to meet the desired physics goals. This note discusses these improvements: Section 2 delves into the upgrade of the HLT farm, GPU-based reconstruction, and the new tracking algorithm for Run 3. Section 3 describes the Run 3 HLT Menu developments and composition. Section 4 provides an overview of the performance of the primary physics objects in Run 3. More information on the development and performance of CMS HLT in Run 3 can be found in [2] and [3].

2. HLT farm upgrade and Run 3 Tracking

The fully upgraded HLT farm started operation from the start of Run 3. It consists of 200 machines (nodes), each consisting of dual AMD EPYC 7763 “Milan” 64-core processors and two Nvidia T4 GPUs, amounting to 25,600 CPU cores and 400 GPUs. A significant portion of the HLT reconstruction has been ported to run on GPUs, with components such as HCAL, ECAL, Pixel Local Reconstruction, Pixel-Only Track, and Vertex Reconstruction now running on GPUs. This transition resulted in a reduction in HLT timing by around 40%, equating to a throughput gain of approximately 70% compared to CPU only HLT reconstruction. To fully harness the potential of GPUs, a new tracking algorithm called Patatrack, using optimised pixel tracks, was introduced [4] for Run 3. This allowed the tracking reconstruction process to be consolidated into a single iteration, leading to a substantial cut in reconstruction time with no loss in performance. Stress tests done in the middle of 2023 showed that the farm is able to sustain pileup upto 65 for a max L1 rate of 110 kHz. Plans for increasing the capacity of the farm by around 20% for 2024 is under consideration.

3. Run 3 HLT Menu

Starting from the baseline trigger menu at the end of Run 2, the Run 3 HLT menu has been significantly expanded to explore new and previously unexplored phase space. The thresholds for almost all triggers were either loosened or retained from Run 2. Additionally, many of the legacy reconstruction algorithms from Run 2 have been replaced with more advanced algorithms as described in the next section.

One of the main additions to the Run 3 Menu was the new triggers for Long Lived Particle (LLP) Searches. Although dedicated LLP triggers existed in the Run 2 HLT menu, they largely depended on the tracker and were hindered by high thresholds and a lack of dedicated Level 1 (L1) seeds. For Run 3, particular emphasis was given to developing new dedicated displaced signatures, such as displaced leptons and delayed jets using ECAL and HCAL timing. These enhancements aim to probe the longer lifetime signatures more effectively. Moreover, new dedicated triggers have been introduced that trigger on particle showers in the muon system, targeting the detection of very long-lived particles. These triggers provide a substantial increase in acceptance (up to 10x) for benchmark signal models compared to the triggers used in Run 2.

HLT Algorithm	**Rate (Hz)**
Isolated Muon with $p_T > 24$ GeV	250
Isolated Electron with $E_T > 30$ GeV	203
Particle Flow (PF) based $p_T^{miss} > 110$ GeV	81
4 PF Jets with $p_T > 70/50/40/35$ GeV	57
Two Isolated Tau Leptons with $p_T > 35$ GeV	54
Muon with $p_T > 50$ GeV	51
Two Electrons with $E_T > 25$ GeV	21
AK4 PF jet with $p_T > 500$ GeV	16
Two same-sign muons with $p_T > 18/9$ GeV	10

Table 1: HLT rates and thresholds of some of the generic triggers in the Run-3 HLT Menu. The rates were obtained from measurements during an LHC fill in November 2022 with an average pileup of 54 (corresponding to an instantaneous luminosity of 1.8×10^{34} cm^{-2} s^{-1}) and have been scaled to 2.0×10^{34} cm^{-2} s^{-1}.

Table 1 lists the thresholds and rates for some of the standard representative triggers in the Run 3 HLT Menu. In 2023, some of the triggers in the standard menu (e.g., the BTagging , VBF triggers and high rate LLP Jet triggers) were moved to Parking with reduced thresholds and higher rates.

The strategy of Data Parking has been significantly enhanced for Run 3. Thanks to the utilization of the decommissioned legacy nodes of the Run 2 HLT farm, 100% of the parking was promptly reconstructed for both 2022 and 2023. The Parking strategy for 2022 was entirely dedicated to B-Physics studies. An inclusive dimuon trigger with a rate of around 1.5 kHz at 2.0×10^{34} cm^{-2} s^{-1} reconstructs all opposite-sign dimuon candidates (with p_T > 4 and 3 GeV) within the invariant mass range from 0.2 to 8.5 GeV. In addition, new set of low mass di-electron triggers were developed for lepton universality tests such as R_k measurement. Since the main limitation for this trigger is the high L1 thresholds, a dynamic prescaling strategy was used, enabling the lower threshold L1 seeds as luminosity drops during the fill. This trigger, which had an HLT rate up to 1.5 kHz in 2022, was made significantly purer (5-10x) in 2023 by tightening the requirement of minimum matching pixel hits at the cost of a -10% relative efficiency loss. The spare rate, made available by this optimization, was used to diversify the parking stream to accommodate looser versions of the earlier prompt triggers: new VBF triggers, Particle Net based Di-Higgs triggers with reduced H_T and b-tag requirements, and also displaced and delayed jet triggers for long-lived particle searches.

Many developments were made to Data scouting for Run 3. The average event size is only around ~7 kB compared to ~1 MB for RAW event. The scouting rate was increased from around 5 kHz in 2018 to 30 kHz for Run 3. This huge increase in input rate was made possible due to the

PoS(EPS-HEP2023)517

special version of Particle Flow (PF) using only Patatrack pixel tracks. Since Pixel reconstruction is offloaded to GPUs, this presents a significant speedup in reconstruction and allows to run scouting at a high rate. Figure 1 shows the performance comparison of the PF Jet and Dimuon scouting objects w.r.to the ones reconstructed offline.

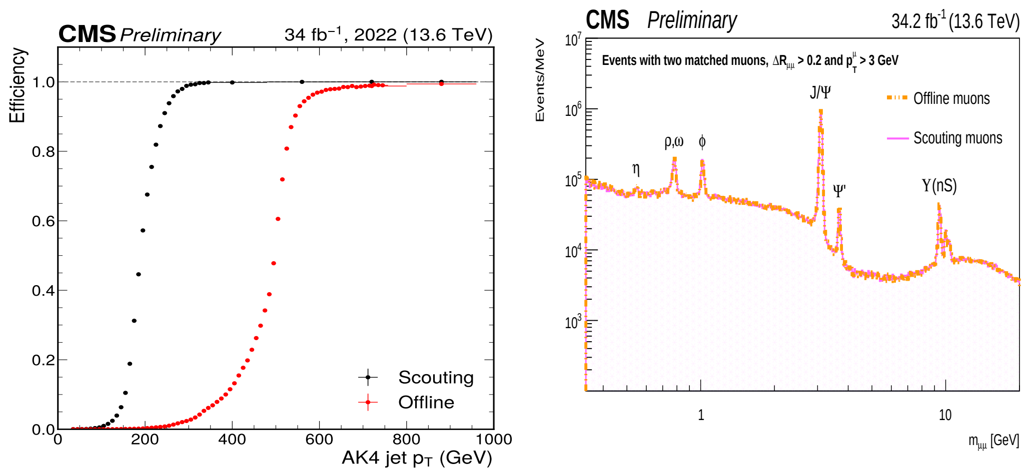

Figure 1: Performance comparison of the PF Jet and Dimuon scouting objects w.r.to the ones reconstructed offline. Left side plot signifies the huge gain in acceptance with Scouting (Single Jet p_T threshold of 180 GeV vs 500 GeV for standard trigger).Right side plot shows close to equal performance for muons

4. Object Performance

HLT Muon reconstruction is a two-step process: using hits from the muon system (L2 or standalone muon) and then integrating hits from the inner tracking system (L3). There are two main complementary tracking steps for L3: Seeding from Standalone Muons and and working inward ("outside-in") or starting from inner tracker layers and working out ("inside-out"). For Run 3, a BDT-based seeding technique improved inside-out tracking speed by 18%, without compromising on the performance. A DNN-based seeding was also introduced for outside-in tracking with slight improvement in performance. The performance of the main Single Muon trigger can be seen in Figure 2. The HLT reconstruction is almost 99% efficient with respect to offline muons.

Electron and Photon Identification begins with regional reconstruction of calorimeter deposits around the L1 candidates to form ECAL super clusters (L1T unseeded reconstruction is also used by some triggers to maximise efficiency). An energy correction, based on a BDT using only ECAL data, is applied, with additional criteria on energy deposits and the HCAL to ECAL energy ratio. Electrons are identified by matching the supercluster to a compatible track with a minimum pixel hit requirement. Once matched with pixel seeds, the electron track is reconstructed using the Gaussian sum filter (GSF) tracking method[5]. The efficiency of the Single Electron trigger in the barrel region for 2022 data is shown in Figure 3.

For tau tagging, a novel Convolutional Network-based DeepTau algorithm, adapted from offline reconstruction, replaced the earlier cut-based algorithm of Run 2. Additionally, a CNN filter was developed at an earlier stage (L2) using only calorimeter and pixel tracks, reducing reconstruction

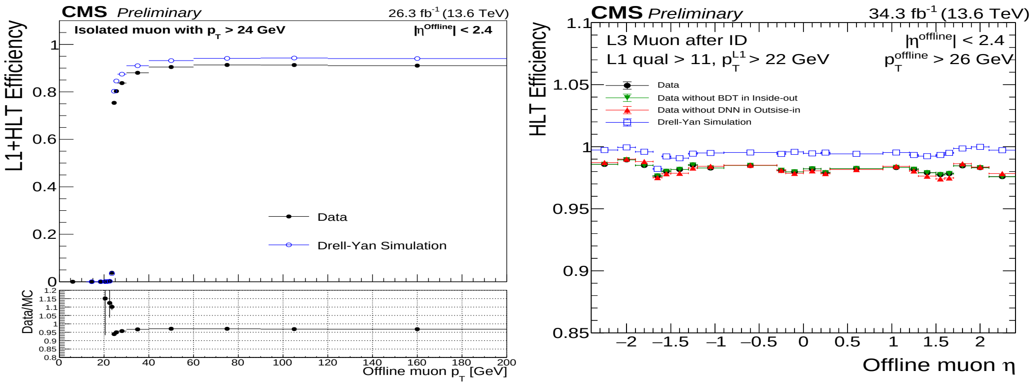

Figure 2: Efficiency of the Single Isolated Muon trigger in 2022. The HLT only efficiency is almost 99% (100% in Drell Yan MC)

time. These developments led to superior performance in Run 3 compared to Run 2 as shown in Figure 4.

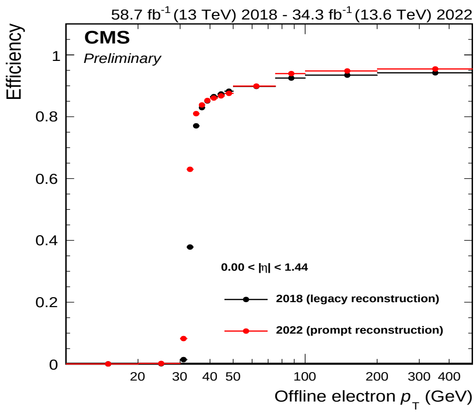

Figure 3: Efficiency of the Single Ele32 trigger in the barrel region

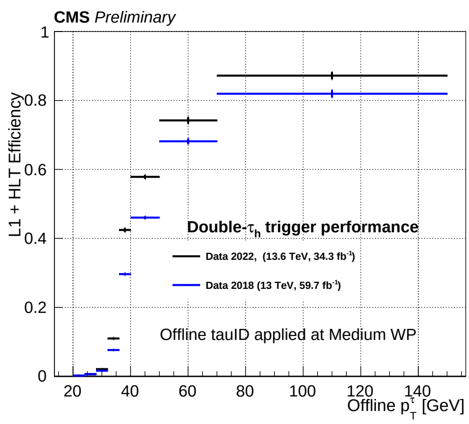

Figure 4: Comparison of the di-tau35 trigger's performance in 2022 and 2018

Jets are reconstructed at HLT starting from particle flow candidates using the Anti-kT algorithm with a distance parameter of 0.4 and 0.8 for fat jets for boosted signatures. For missing E_T triggers, a forward HCAL (HF) noise-cleaning method was incorporated, allowing the threshold to be reduced from 120 GeV to 110 GeV. Faster efficiency turn-on was observed after an HCAL response update in Oct 2022 as can be seen in Figure 5.

For B tagging, the Run 2 DNN-based DeepCSV algorithm was replaced by the CNN-based DeepJet[6] and Graph Network-based ParticleNet triggers[7]. These provide significant improvement, with much higher efficiencies for the same mistag rate as shown in Figure 6. Efforts are ongoing to completely replace the DeepJet triggers with ParticleNet.

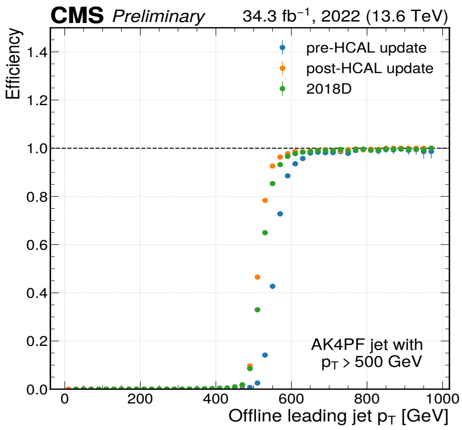

Figure 5: Comparison of high-PT jet trigger efficiency between 2018 and 2022

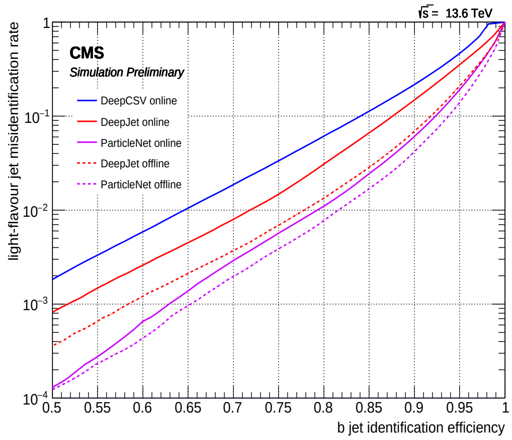

Figure 6: Performance comparison of different B-taggers used in Run 3 HLT

5. Conclusion

Many significant developments were made to the CMS HLT for Run 3 with further improvements such as transitioning to more sophisticated taggers planned for the next years. The performance in the the first years of Run 3 confirm the good performance of the CMS Trigger. These developments are expected to significantly enhance the sensitivity of many future analyses.

References

[1] CMS Collaboration, "The CMS Experiment at the CERN LHC," JINST 3, S08004 (2008).

[2] CMS Collaboration, "CMS detector for LHC Run 3," arXiv:2309.05466, CERN-EP-2023-136, 2023.

[3] CMS Collaboration, "HLT Run III Results," `https://twiki.cern.ch/twiki/bin/view/CMSPublic/HighLevelTriggerRunIIIResults`.

[4] Bocci, A. et al., "Heterogeneous Reconstruction with the CMS Pixel Tracker," Frontiers in Big Data, 3, 2020, `https://www.frontiersin.org/articles/10.3389/fdata.2020.601728`.

[5] W. Adam, R. Fruhwirth, A. Strandlie, T. Todorov, “Reconstruction of electrons with the gaussian-sum filter in the CMS tracker at the LHC”, Journal of Physics G 31 (2005) N9–N20, `https://doi.org/10.1088/0954-3899/31/9/n01`.

[6] E. Bols, J. Kieseler, M. Verzetti, M. Stoye, and A. Stakia, "Jet flavour classification using DeepJet," Journal of Instrumentation, vol. 15, no. 12, p. P12012, Dec. 2020.

[7] H. Qu, L. Gouskos, "Jet tagging via particle clouds," Phys. Rev. D, vol. 101, no. 5, p. 056019, Mar. 2020.