

CMS trigger performance

Silvio Donato^{1,a} on behalf of the CMS Collaboration

¹University of Zurich

Abstract. During its second run of operation (Run 2), started in 2015, the LHC will deliver a peak instantaneous luminosity that may reach $2 \cdot 10^{34} \text{ cm}^{-2}\text{s}^{-1}$ with an average pileup of about 55, far larger than the design value. Under these conditions, the online event selection is a very challenging task. In CMS, it is realized by a two-level trigger system: the Level-1 (L1) Trigger, implemented in custom-designed electronics, and the High Level Trigger (HLT), a streamlined version of the offline reconstruction software running on a computer farm.

In order to face this challenge, the L1 trigger has been through a major upgrade compared to Run 1, whereby all electronic boards of the system have been replaced, allowing more sophisticated algorithms to be run online. Its last stage, the global trigger, is now able to perform complex selections and to compute high-level quantities, like invariant masses. Likewise, the algorithms that run in the HLT have been greatly improved; in particular, new approaches for the online track reconstruction lead to a drastic reduction of the computing time, and to much improved performances. This document will describe the performance of the upgraded trigger system in Run 2.

1 Introduction

1.1 The CMS trigger

The Compact Muon Solenoid (CMS) is a multi-purpose experiment built at the Large Hadron Collider (LHC). It is made of a silicon tracker subdetector, a homogeneous electromagnetic calorimeter, and a hadron calorimeter immersed in a magnetic field of 3.8 T produced by a superconducting solenoid. Muons are detected by gas-ionization subdetectors positioned in the flux-return yoke outside the solenoid. Figure 1 shows the scheme of the CMS experiment and of the event reconstruction. More details about the CMS detector can be found in [1].

At the LHC, protons circulate in bunches with a time interval of 25 ns in two beams with opposite directions. As a consequence, proton bunches collide with a rate of 40 MHz, but only about 1 kHz of collisions is actually registered for data analysis. This limit is mainly determined by the offline computing resources involved in the offline reconstruction. The events are selected online by a two-level trigger system that reduces the rate from 40 MHz to 1 kHz.

The Level-1 trigger (L1) is made of customizable hardware and reduces the rate from 40 MHz to 100 kHz. The L1 trigger decision is taken in less than $4 \mu\text{s}$ using a simplified readout of the calorimeter and muon subdetectors. As soon as an event is accepted by the L1 trigger, the full readout of the CMS

^ae-mail: silvio.donato@cern.ch

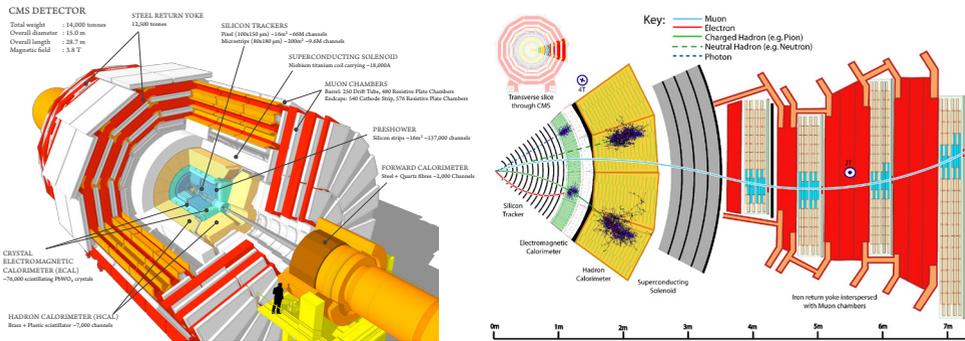


Figure 1. The subdetectors of the CMS experiment (left) and a schematic view of the different particles' signature inside the detector (right).

detector takes place. The maximum L1 trigger output rate (100 kHz) is defined by the readout of the detector.

The High Level Trigger (HLT) is a computing farm made of 26 thousands of cores that selects events using software algorithms from a rate of 100 kHz down to ~ 1 kHz. The limit on the average computing time is 260 ms. Differently from L1 trigger, the HLT has access to the full detector information.

2 The L1 trigger

The L1 trigger has been entirely upgraded between 2013 and 2015, as part of the CMS Phase-1 upgrade [2]. Figure 2 shows a working scheme of the new L1 trigger. The L1 trigger event reconstruction starts from a simplified readout performed in the calorimeter (HCAL, ECAL, HF) and in the muons chambers (CSC, DT, RPC), called trigger primitives. No tracker information is accessible to the L1 trigger. Trigger primitives are then combined to form calorimetric towers and to link compatible muon hits together. Afterwards, they are used to form the L1 trigger objects: muons, electrons/photons, τ , jets, missing transverse energy and total transverse energy. Eventually, the event is accepted if the L1 trigger objects fulfill at least one of the conditions defined by the L1 menu, eg. one muon with $p_T > 15$ GeV and missing transverse energy larger than 100 GeV.

The next paragraphs will show the performance in the reconstruction of the L1 trigger objects.

2.1 The L1 electron/photon

In the L1 trigger system photons and electrons cannot be distinguished. Many improvements have been achieved with the Phase-I upgrade in the electron/photon reconstruction, and some of them have been fully exploited recently in 2017. The left plot in Fig. 3 compares the efficiency curves obtained in 2016 and 2017 for two single electron/photon triggers having the same rate: in 2017 we lowered the energy threshold by 2 GeV and improved the energy resolution of the L1 objects.

Furthermore, in 2017 a loose version of the electron/photon isolation has been introduced to be mainly applied to the multi-object triggers in order to increase their efficiency. Figure 3 (right plot) shows the efficiency of the single electron/photon trigger with standard and loose isolation.

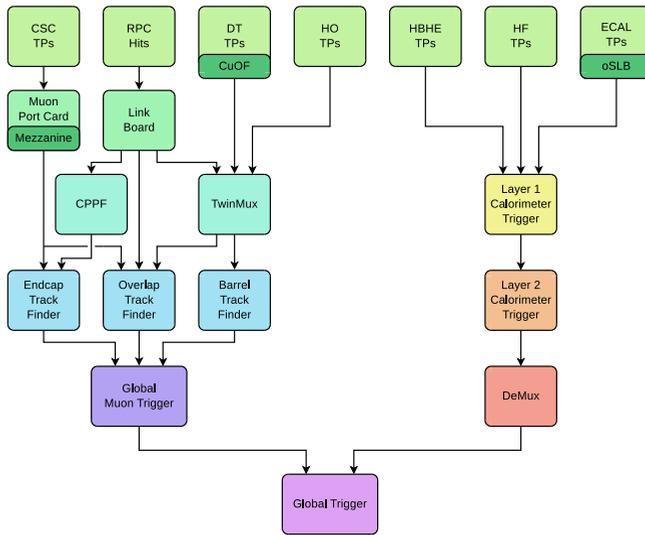


Figure 2. Working scheme of the Phase-1 L1 trigger system. [2]

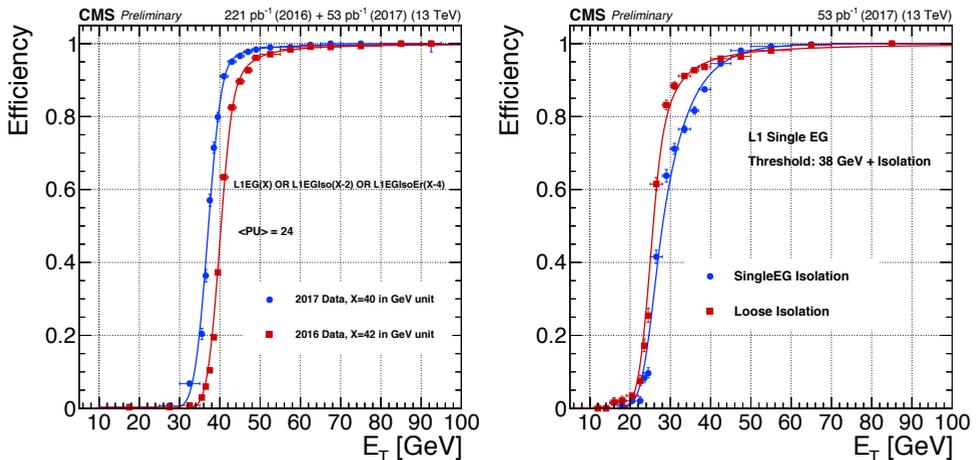


Figure 3. On the left, L1 trigger efficiency curve for an electron/photon object as a function of the offline reconstructed supercluster transverse energy E_T with the 2017 algorithm (blue) and the 2016 algorithm (red) measured using tag-and-probe method on data. A geometrical matching between the electron supercluster and the L1 candidate is applied. The efficiency is drawn for the combination of the following triggers: e/γ with $p_T > X$ GeV without isolation requirements (EG) OR e/γ with $p_T > (X - 2)$ GeV and isolation requirement (EGIso) OR e/γ with $p_T > (X - 4)$ GeV, isolation requirement and pseudorapidity $|\eta| < 2.1$ (EGIsoEr). The two curves are drawn for the thresholds X ensuring the same trigger rate. On the right, L1 trigger efficiency curve for an electron/photon object as a function of the offline reconstructed supercluster transverse energy E_T with 2017 tight EG isolation requirements and a p_T threshold of 25 GeV (blue) and emulated events with loose EG isolation requirements and p_T threshold of 25 GeV (red). The loose isolation allows to increase the efficiency for cross triggers with isolated e/γ such as isolated e/γ plus muons, jets or missing transverse energy. [3]

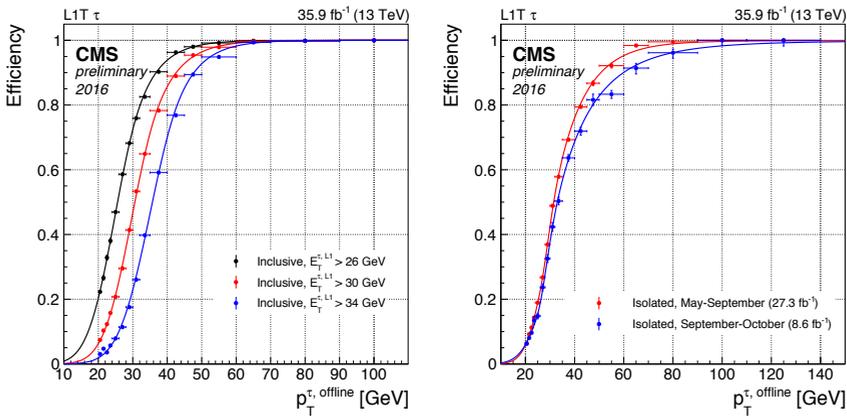


Figure 4. On the left, L1 trigger efficiency of inclusive τ -seeds (i.e. without any requirement on the isolation at L1) as a function of the offline τ transverse momentum. On the right, L1 trigger efficiency of isolated τ -seeds (i.e. requiring the L1 τ candidate to pass a cut on its isolation transverse energy) as a function of the offline τ transverse momentum, for a L1 E_T threshold of 30 GeV. The isolation criterion was tightened in September 2016 to cope with the increasing LHC luminosity. In the period May-September, the peak instantaneous luminosity was lower than $1.34 \times 10^{34} \text{ cm}^{-2} \cdot \text{s}^{-1}$, while it has reached $1.53 \times 10^{34} \text{ cm}^{-2} \cdot \text{s}^{-1}$ in the September-October period. [4]

2.2 The L1 τ

The identification of hadronically-decaying τ has been significantly improved with the Phase-I upgrade. The higher granularity of the calorimeters allowed to tag τ in a more sophisticated way, requiring isolation too. L1 τ energy is calibrated as a function of the energy and pileup. The isolation cut depends on the τ energy, η and pileup, in order to get a flat τ efficiency. Figure 4 shows the efficiency curve of the single object trigger, for different working points, for the isolated (left) and inclusive (right) trigger. The resolution of L1 τ p_T is about 20 – 25% at 30 – 50 GeV.

2.3 The L1 vector boson fusion

Another important achievement obtained with the Phase-I upgrade is a more powerful L1 trigger able to compute the invariant mass between objects. It has been exploited in a trigger dedicated to the search for the $H \rightarrow \tau\tau$ produced through the vector boson fusion. The new trigger requires the presence of two τ jets plus two forward/backward jets with a large invariant mass between them. Figure 5 shows the large acceptance gained by this new trigger as a function of the offline forward/backward jet invariant mass, for a sample of simulated $H \rightarrow \tau\tau$ events produced through the vector boson fusion.

2.4 The L1 jet/MET

At L1 trigger level, jets are reconstructed using only the calorimeters. Their reconstruction has been significantly improved with the Phase-I upgrade, especially thanks to the higher granularity of the calorimetric trigger primitives and a more advanced pileup correction. Similar improvements have been obtained for the missing and total transverse energy. Figure 6 shows the trigger efficiency curves

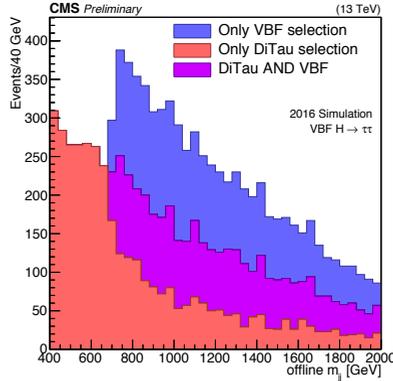


Figure 5. Gain from the L1 VBF seed (with at least one jet with $E_T > 110$ GeV and at least two jets with $E_T > 35$ GeV and $m_{jj} > 620$ GeV) with respect to the baseline Double Tau seed (at least 2 isolated taus with $E_T > 32$ GeV and $|\eta| < 2.1$) as a function of the offline m_{jj} . The offline selections are chosen to be tighter than online. The events in each of the categories “Only VBF selection” and “Only Di-Tau selection” are events that fired only one of the triggers and pass only the offline selection corresponding to that trigger. The events in the category “Di-Tau AND VBF” fire both triggers and pass both offline selections. [4]

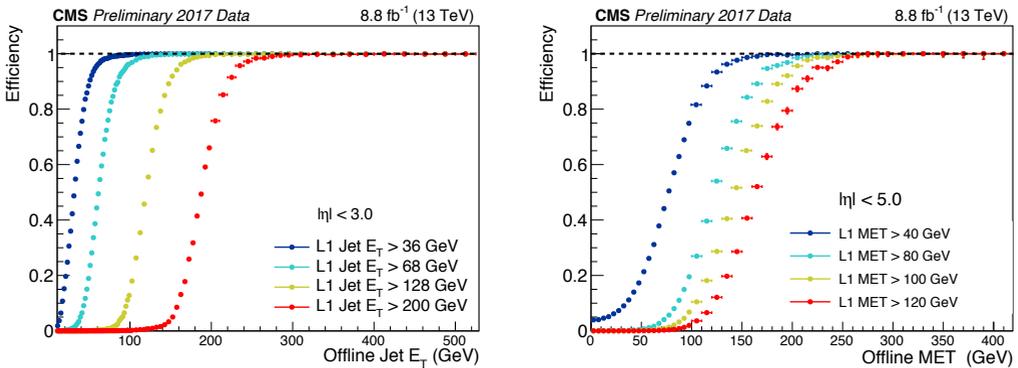


Figure 6. On the left, efficiency curves for the L1 single jet triggers using events with an offline muon as a function of offline jet transverse energy reconstructed with the anti-kt ($R=0.4$) algorithm and $|\eta| < 3$. Offline and L1 jets are matched with $\Delta R < 0.4$. On the right, efficiency curves for the L1 missing energy trigger using events with an offline muon as a function of offline missing transverse energy. [5]

as a function of the offline jet transverse energy E_T for several E_T thresholds of the single jet trigger (on the left), and as a function of the offline missing transverse energy (MET) for several MET thresholds triggers (on the right).

2.5 The L1 muon

The upgraded L1 muon system has a higher granularity, more subdetectors with the addition of the fourth layers of CSCs and RPCs, and better usage redundancy.

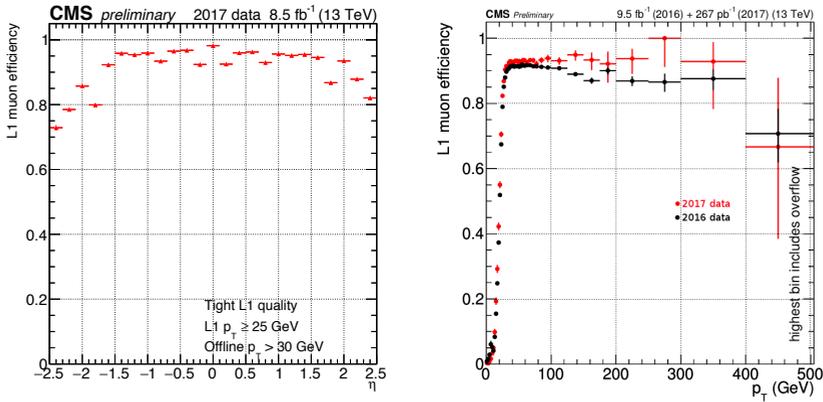


Figure 7. Trigger efficiency for the single muon L1 trigger with $p_T > 25$ GeV. On the left, efficiency as a function of the offline reconstructed muon η with $p_T > 30$ GeV [3]. On the right, efficiency as a function of the offline reconstructed muon p_T [6]. A more recent version of the plots is available at [7].

The muon reconstruction has been improved by a more sophisticated track finder that allowed to increase the efficiency. Figure 7 shows the L1 single muon efficiency as a function of η (left) and p_T (right). The latter plot shows also a further improvement that has been obtained in 2017 for high p_T muons.

3 The HLT trigger

3.1 The HLT path

The output rate of the L1 trigger is about 100 kHz. The HLT trigger reduces this rate further down ~ 1 kHz running software algorithms. The limit on the average computing time is fixed by the number of cores and the input rate: ~ 26000 cores/100 kHz ~ 260 ms.

The HLT is composed of hundreds of HLT paths, that are sequences of reconstructing and filtering modules. As soon as an event is rejected by a filter, the subsequent modules are not run. The idea is to reconstruct firstly quick objects, cut on them, and then reconstruct more precise/slow objects, to keep a low average computing time.

In order to explain in some details the functioning of an HLT trigger path, we consider as an example a hypothetical trigger that requires MET > 120 GeV plus one b tagged jet [8].

A typical sequence of selections used in a HLT path with b tagging uses the following information: L1 seed, calorimetric objects, b tagging with regional tracking, and PF objects (the numbers shown are purely indicative since they depend on many variables like pileup, CPU and so forth).

- **L1 seed** [Input rate ~ 100 kHz; Time < 1 ms]. Each HLT path runs over every event accepted by the L1 trigger (~ 100 kHz). The first cut of a HLT path is the selection of a specific L1 trigger seed, i.e. a specific trigger in the L1 menu or a logical combination of them. In the example, the L1 seed used is MET > 70 GeV. The cut on the L1 seed reduces the rate by a factor ~ 10 .
- **L2 - Calorimetric selection** [Input rate ~ 10 kHz; Time ~ 50 ms]. Since the reconstruction of the calorimetric objects (e.g. calorimetric MET) is very fast, it takes place usually just after the L1 seed selection. The resolution of these objects is worse compared to the PF objects, but they can

be reconstructed in only ~ 50 ms, 20 times faster than the corresponding PF objects. Calorimetric objects can be reconstructed for any event accepted by the L1 trigger. When the selection criteria are applied to calorimetric objects tracks are not available, and hence, for instance, we cannot tag pileup jets. In the example, the L2 cuts used are: calorimetric MET > 70 GeV and two calorimetric jets with $p_T > 30$ GeV and $|\eta| < 2.6$. Usually, the L2 cuts reduce the rate by a factor 5–10.

- **L3 - Regional tracking** [*Input rate* ~ 2 kHz; *Time* ~ 100 ms]. As tracking is a slow process, the tracker information is not used in the L1 and L2 stages. Still, tracks are essential for many tasks like distinguishing between photons and electrons, identify b and τ jets, evaluate the track-based lepton isolation, and measure precisely the lepton momentum. To take advantage of the tracking information at HLT, two strategies are used: run tracking over a small fraction of events and only in the regions of interest (cones with $\Delta R = 0.4$ around the calorimetric objects). In the example, the regional tracking is used to perform b tagging and it is used to reduce the rate down to ~ 500 Hz.
- **PF sequence** [*Input rate* ~ 500 Hz; *Time* ~ 1000 ms]. The Particle Flow (PF) sequence is the most accurate and slow reconstruction sequence available at HLT. Usually, the tightest cuts at HLT are applied to PF objects. The PF sequence used at HLT is similar to the offline version, but it uses a simplified version of tracking. The HLT PF tracking reconstructs only tracks close to leading primary vertices and with $p_T > 500$ MeV. Then, tracks are linked to objects reconstructed with the calorimeters and the muon chambers. The PF sequence reconstructs any high level objects: jets, MET, electrons, photons, muons, taus, etc. In some HLT paths, the PF tracks are also used to perform b tagging on PF jets. In the example, the PF objects are used for the final cut MET > 120 GeV bringing the rate down to ~ 10 Hz.

The input rate, time, and average-per-event time of each HLT step is reported in Table 1. Note that even if the PF sequence needs ~ 1000 ms, it increases the average path time by only ~ 5 ms because it is run at ~ 500 Hz.

Figure 8 shows the computing time distribution of the HLT in a run of 2016, when an instantaneous luminosity of $1.5 \cdot 10^{34}$ Hz cm $^{-2}$ was delivered by the LHC. In 2016, the number of cores was 22 k, and it has been increased in 2017 to 26 k. Moreover, the HLT reconstruction has been improved and the global timing decreased.

Table 1. Time and rate of the object reconstruction of an hypothetical MET + b tagging trigger. The numbers are purely indicative since they depend on many variables. The numbers reported are the running time, the input rate, and the average time spent per event. The last column can be calculated as the running time multiplied by the input rate and divided by the global HLT input rate (~ 100 kHz). For reference, the limit at HLT on the average global latency is ~ 260 ms and on the rate is ~ 1 kHz.

Object reconstruction	\sim Time (ms)	\sim Rate (Hz)	\sim Av. time (ms)
L1 variables	< 1	10^5	< 1
Jet/MET	50	10^4	5
Regional tracking (b tagging)	100	2000	2
Particle flow	1000	500	5
Data storing	–	10	–

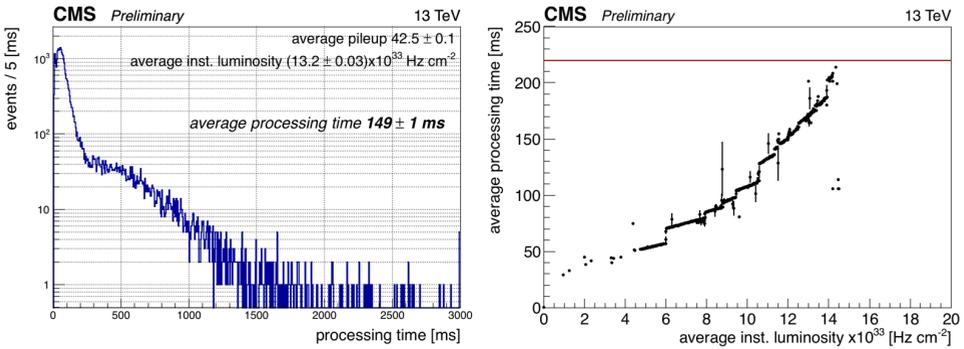


Figure 8. On the left, HLT computing time in a run with average instantaneous luminosity equal to $1.3 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and pileup of 42.5. The processing time was measured on an "Ivy Bridge" Intel Xeon E5-2650v2 processor, running one single job. On the right, average HLT computing timing vs instantaneous luminosity. Both plots have been obtained with 2016 data. [9]

3.2 The HLT rate

The average HLT rate limit over the year is about 1 kHz. This limitation comes from the computing resources involved in the offline reconstruction, while the data acquisition system is able to transfer up to 5 – 6 kHz and the space on tape would allow us to record up to 2 kHz. Since this limit applies to the average HLT rate, we can have short periods with peak rate above 1 kHz.

We can overcome the 1 kHz limit in two ways. For several exotic searches, we store only online objects with a rate of $\sim 5 \text{ kHz}$ (“data scouting”) that will be used directly in data analyses, skipping the offline reconstruction. We can also save more data on tape without running any offline reconstruction (“parking”). During Run-2, about 10% of scouting data (rate $\sim 500 \text{ Hz}$) has been “parked” saving the full event content. The rationale is to allow in the future to confirm/reject a possible hint of New Physics found by data scouting analyses. Furthermore, we record temporary additional data with no offline reconstruction for trigger commissioning, e.g. for trigger rate estimate.

3.3 The HLT tracking

In 2017, the pixel subdetector has been upgraded. The new subdetector has one more layer, both in the barrel and in the endcap, and can work efficiently even at high luminosity. The HLT tracking has been updated to exploit the four layers, especially in the track seeding. Figure 9 shows the online track seeding efficiency, on the left side, and fake rate, on the right side, as a function of η . The efficiencies and fake-track rates have been evaluated for the old pixel subdetector (“2016 HLT tracking”), for the new pixel subdetector with the old pixel track seeding based on triplets (“Triplet propagation”), and for the new pixel subdetector with the new track seeding based on quadruplets (“Cellular Automaton”). The plots show that with the new pixel subdetector we have a fake rate decrease by more than a factor ten and an efficiency increase, especially in the endcaps. The new track seeding shows a similar efficiency and a lower fake rate. The main difference between the two algorithms is in computing time. Table 2 shows the computing time expected for the different tracking algorithms: the new version is five time faster than the old tracking with the new pixel subdetector. Moreover, the new algorithm is ready to be ported to GPU. This would further decrease timing by a factor ten, as shown in the last row of the table.

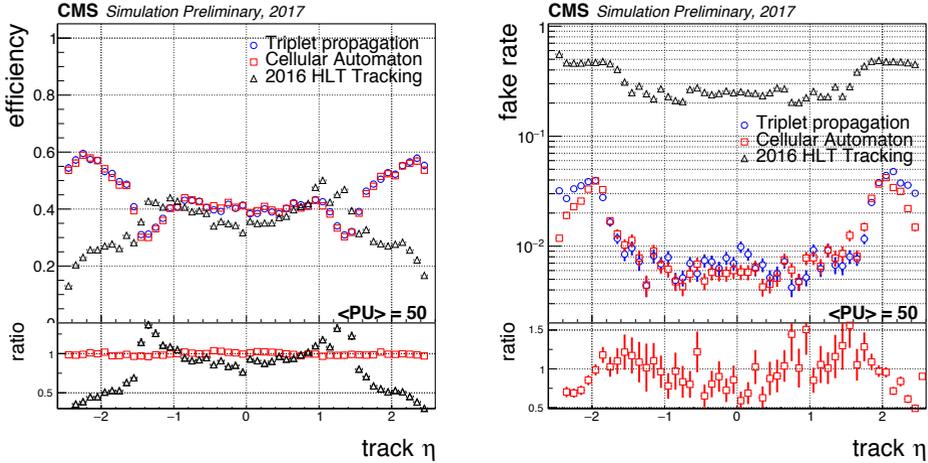


Figure 9. HLT Pixel Tracking efficiency (left) and fake-rate (right) as a function of η for 1000 $t\bar{t}$ simulated events with superimposed pileup collisions of 50, for three different algorithms. The Cellular Automaton and the Triplet Propagation algorithms with Phase-I pixel subdetector are compared to the pixel tracking running at HLT during 2016 with the Phase-0 pixel subdetector. The ratio between the 2016 HLT Pixel Tracking and Triplet Propagation with respect to the Cellular Automaton plot is also shown. The inefficiency in the interval $1.0 < |\eta| < 1.5$ is due to the non hermetic four-hits coverage transitioning from the barrel to the forward region. The asymmetry of the distribution in eta is due to the realistic position of the beamspot, slightly displaced with respect to the origin. [10]

3.4 The HLT jet and MET/HT

At HLT, tracking is used to run a simplified version of the Particle Flow (PF) algorithm to improve the objects reconstruction, especially on jets and MET. As a reference, the p_T relative resolution of a simulated jet with $p_T \sim 30$ GeV is of 28% for the online calorimetric jets and of 18% for online PF jets. Figure 10 shows the efficiency as a function of the offline reconstructed quantity for single jet triggers with different p_T thresholds, on the left side, and for missing transverse energy triggers with different MET thresholds, on the right side.

Table 2. Timing performance of different quadruplets generators at HLT PixelTracks for 1000 events $t\bar{t}$ simulated for the Phase-I subdetector with superimposed pileup collisions of 50. Tracking was executed on every event. The CPU used for tests is a Intel Core i7-4771 CPU 3.50 GHz, 4 cores, 8 simultaneous multi-threading threads. The GPU used is a NVIDIA GTX 1080. In all the cases the CMSSW framework was configured to use 8 threads.

Quadruplet Algorithm	Time per event	speedup wrt. 2016
2016 Pixel Tracks	(29.3 ± 13.1) ms	1 \times
CPU Cellular Automaton	(14.0 ± 6.2) ms	2.1 \times
Triplet Propagation	(72.1 ± 25.7) ms	0.4 \times
GPU Cellular Automaton	(1.2 ± 0.9) ms	24.4 \times

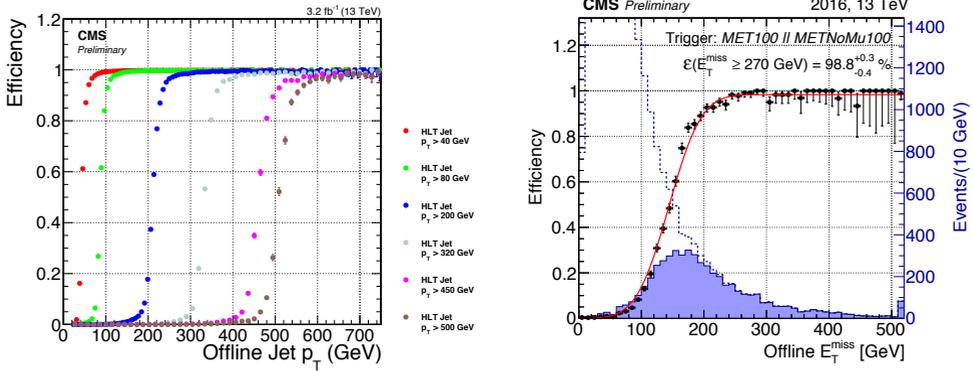


Figure 10. On the left, trigger efficiency for jets at High Level trigger (HLT) as a function of the offline jet transverse momentum. The events are triggered by a single muon. The jets are restricted to a pseudorapidity range $|\eta| < 3.0$. In each efficiency curve, the numerator is the number of offline jets which are matched to the online jets inside a $dR < 0.2$ and of which the transverse momentum is larger than trigger threshold. The denominator is the number of offline jets. On the right, efficiency of the logical OR of the HLT PFMET[NoMu]100 PFMHT[NoMu]100 IDTight triggers as a function of offline missing transverse energy measured in 13 TeV data selected with online requirements of a 25 GeV electron. [11][12]

3.5 The HLT b tagging

At HLT, online b tagging is performed in two different ways. In the first one, b tagging is evaluated using tracks reconstructed in the PF sequence. It is used to b tag PF jets. The second one is used to b tag calorimetric jets before running the PF sequence. It exploits a dedicated regional tracking around the leading calorimetric jets using a primary vertex constraint, allowing to reduce the computing timing from 356.7 ms (PF tracking) to 86.9 ms. More details are available in [8]. Figure 11 shows the b tagging efficiency vs light-jet fake-rate for the offline and online b tagging of calorimetric jets.

3.6 The HLT electron

Electrons are reconstructed at HLT using a dedicated tracking algorithm - the Gaussian-sum filter - that was developed for particles with a strong bremsstrahlung energy loss. The track information is then combined with the energy deposited in the electromagnetic calorimeters. The online electron reconstruction efficiency is very high ($> 99\%$). Tight isolation/identification cuts are applied to reduce rate originated from jets and photon conversions, allowing to reduce the trigger threshold in p_T . After tight cuts, the online electron efficiency is of $\sim 90\%$ in barrel and $\sim 80\%$ in endcaps, as shown in Figure 12. The efficiency has been measured using a tag-and-probe method on $Z \rightarrow ee$ events.

3.7 The HLT τ

Differently from the L1 trigger, HLT hadronically-decaying τ are reconstructed exploiting also tracks. At HLT, we look for τ in either one or three charged particles decay and up to two $\pi_0 \rightarrow \gamma\gamma$, reconstructed as a ϕ -long strip in the electromagnetic calorimeter. Left plot in Figure 13 shows the HLT-only efficiency for a trigger requiring $\tau p_T > 50$ GeV. A good agreement between data and simulation is found. Right plot shows the combined L1 and HLT efficiency for trigger requiring an

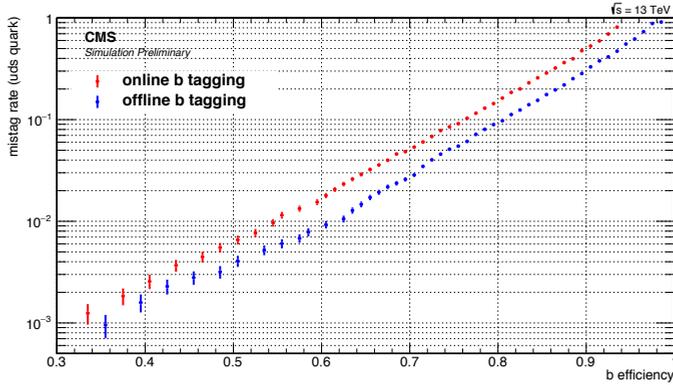


Figure 11. Performance of the online (red) and offline (blue) b jet identification efficiency demonstrating the probability for uds jets to be misidentified as b jet as a function of the efficiency to correctly identify b jets. Both online and offline b tagging are performed using CSVv2 algorithm. The curves are obtained for the eight leading online and offline jets with $p_T > 30$ GeV and $|\eta| < 2.4$ in simulated $t\bar{t}$ events. The plot is obtained using the 2016 subdetector conditions with 35 pileup collisions. Only events having a sum of online jet $p_T > 250$ GeV are considered. The performance in this figure serves as an illustration since the b jet identification efficiency depends on the p_T and η distribution of the jets in the topology as well as the amount of b jets from gluon splitting in the sample. [13]

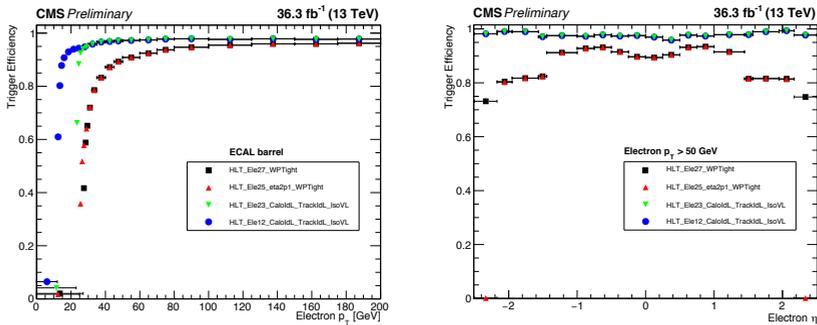


Figure 12. Electron trigger efficiency as a function of the offline electron p_T with $|\eta| \leq 1.479$ (left) and of the offline electron η with $p_T > 50$ GeV (right) in the full 2016 dataset. The efficiency is measured by a $Z \rightarrow ee$ tag-and-probe method with respect to cut-based tight electron identification. HLT Ele27 WPTight and HLT Ele25 eta2p1 WPTight are the lowest transverse energy threshold unrescaled single electron triggers that require $E_T > 27$ GeV and 25 GeV respectively, the latter also being restricted to $|\eta| < 2.17$. HLT Ele23 CaloIdL TrackIdL IsoVL and HLT Ele12 CaloIdL TrackIdL IsoVL are the two legs of the lowest threshold unrescaled dielectron trigger HLT Ele23 Ele12 CaloIdL TrackIdL IsoVL DZ that requires two electrons with $E_T > 23$ GeV and 12 GeV originated from the same vertex. [14]

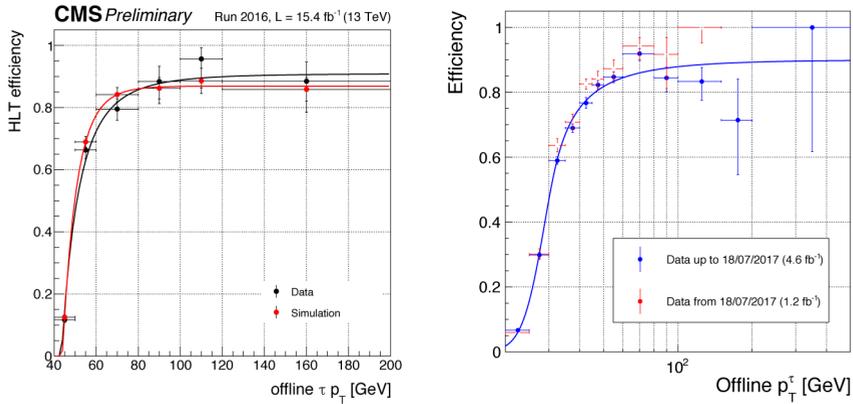


Figure 13. On the left, HLT efficiency in 2016 of a trigger requiring a single τ_h with $p_T > 50$ GeV and a track with $p_T > 30$ GeV. On the right, combined L1 and HLT efficiency of the τ -leg of the $\mu\tau_h$ trigger (loose isolation, $p_T > 27$ GeV) trigger in 2017 data; efficiency in data is shown split between data collected before and after 18/07/2017, when an improved tracking reconstruction at HLT has been deployed. Both efficiency plots have been computed through the tag-and-probe method in $Z \rightarrow \tau_\mu\tau_h$, and measured as a function of the offline $\tau_h p_T$. [15][16]

isolated τ with $p_T > 27$ GeV (L1 τ , $p_T > 24$ GeV), before and after an improvement has been deployed into the online tracking for τ reconstruction. In both cases, the efficiency has been measured using tag-and-probe method selecting events with $Z \rightarrow \tau\tau \rightarrow \mu\nu + \text{hadr.}$

3.8 The HLT muon

At HLT, muons are reconstructed exploiting both the muon subdetectors and the tracker, improving significantly the muon p_T resolution compared to the L1 trigger. The HLT reconstruction efficiency of muons is larger than 99%. In 2016, the muon unprescaled trigger with the lowest threshold required an isolated muon with $p_T > 24$ GeV. The isolation reduces the rate of more than a factor two, adding a small inefficiency of 2 – 3%, as shown in Fig. 14 reporting the HLT-only efficiency as a function of p_T and η . The main muon trigger inefficiency is originated by the L1 trigger, as shown in Fig. 15.

4 Conclusions

The LHC luminosity is steadily increasing and reached pileup 60 in 2017. Many significant improvements have been achieved to cope with such high luminosity, both with the Phase-1 upgrade of the L1 trigger and with new L1 and HLT trigger algorithms. The upgrade is not fully exploited yet and we expect to get other improvements to take on the current and future challenges of taking data at the high luminosity LHC.

References

[1] The CMS Collaboration, “The CMS experiment at the CERN LHC”, *JINST* **3** S08004 (2008).

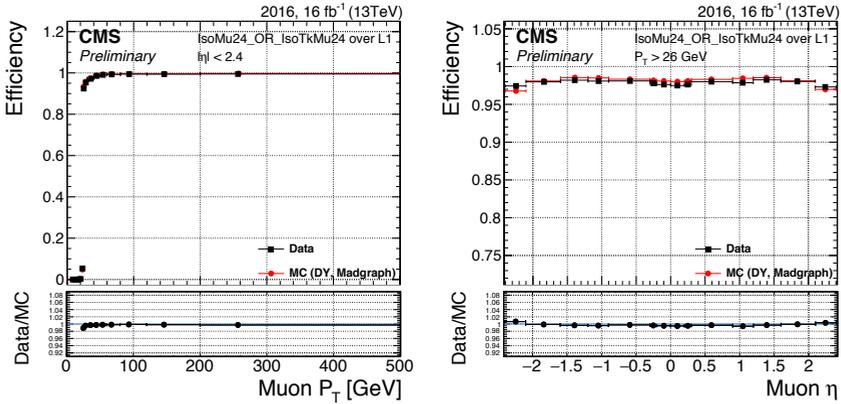


Figure 14. HLT-only trigger efficiency as a function of offline muon p_T with $|\eta| < 2.4$ (left) and of offline muon η with $p_T > 25$ GeV (right) for the OR of the HLT IsoMu24 and HLT IsoTkMu24. These two trigger paths require a muon reconstructed online with $p_T > 24$ GeV as well as passing an isolation selection. The HLT-only efficiency has been computed as the trigger efficiency for events passing the L1 trigger. [17]

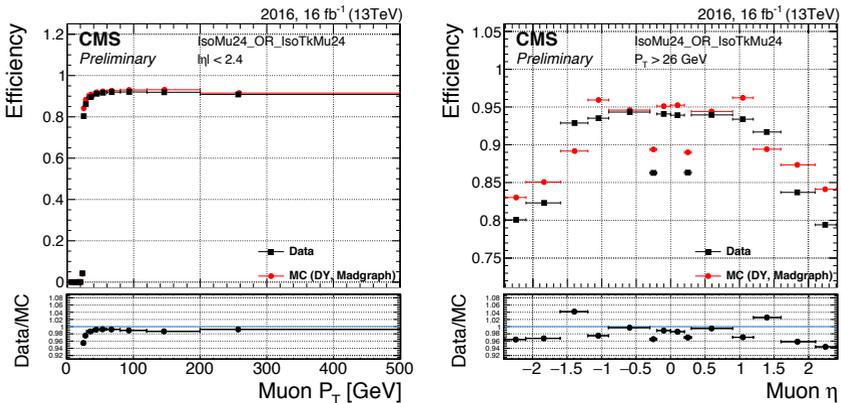


Figure 15. Combined HLT and L1 trigger efficiency as a function of offline muon p_T with $|\eta| < 2.4$ (left) and of offline muon η with $p_T > 25$ GeV (right) for the OR of the HLT IsoMu24 and HLT IsoTkMu24. These two trigger paths require a muon reconstructed online with $p_T > 24$ GeV as well as passing an isolation selection. [17]

[2] CMS Collaboration, “CMS Technical Design Report for the Level-1 Trigger Upgrade,” CERN-LHCC-2013-011, CMS-TDR-12, <https://cds.cern.ch/record/706847>

[3] CMS Collaboration, “Level-1 E/Gamma and Muon performance on 2017 data”, CMS-DP-2017-024, <https://cds.cern.ch/record/2273270>

[4] CMS Collaboration, “Level 1 Tau trigger performance in 2016 data and VBF seeds at Level 1 trigger”, CMS-DP-2017-022, <https://cds.cern.ch/record/2273268>

[5] CMS Collaboration, “CMS L1 Calorimeter Trigger performance in 2016 data”, CMS-DP-2017-044, <https://cds.cern.ch/record/2202966>

- [6] CMS Collaboration, “L1 muon trigger performance - ICHEP16 dataset”, CMS-DP-2016-050, <https://cds.cern.ch/record/2202986>
- [7] CMS Collaboration, “Level-1 muon trigger performance in 2017 data and comparison with the legacy muon trigger system”, CMS-DP-2017-041, <https://cds.cern.ch/record/2286327>
- [8] S.Donato, L.Rolandi, A.Rizzi, “Search for the Standard Model Higgs boson decaying to b quarks with the CMS experiment” CMS-TS-2017-008, CERN-THESIS-2017-059, <https://cds.cern.ch/record/2268298>
- [9] CMS Collaboration, <https://twiki.cern.ch/twiki/bin/view/CMSPublic/HLTplotsCHEP2016>
- [10] CMS Collaboration, <https://twiki.cern.ch/twiki/bin/view/CMSPublic/HLTplotsPixelUpgrade>
- [11] CMS Collaboration, <https://twiki.cern.ch/twiki/bin/view/CMSPublic/HLTplotsFirst2017Data>
- [12] CMS Collaboration, “High Level Trigger Performance Plots for ICHEP2016” CMS-DP-2016/056, http://cms.cern.ch/iCMS/jsp/db_notes/notestable1.jsp?CMSNoteID=DP-2016/056
- [13] CMS Collaboration, “Performance of online b tagging in proton-proton collisions at 13 TeV at the CMS experiment” CMS-DP-2017-011, http://cms.cern.ch/iCMS/jsp/db_notes/notestable1.jsp?CMSNoteID=DP-2017/011
- [14] CMS Collaboration, “Electron and photon performance in CMS with the full 2016 data sample”, CMS-DP-2017-004, http://cms.cern.ch/iCMS/jsp/db_notes/notestable1.jsp?CMSNoteID=DP-2017/004
- [15] CMS Collaboration, “Tau trigger commissioning on first 2017 13 TeV data”, CMS-DP-2017/035, http://cms.cern.ch/iCMS/jsp/db_notes/notestable1.jsp?CMSNoteID=DP-2017/035
- [16] CMS Collaboration, “Tau trigger performances on 2016 13 TeV data”, CMS-DP-2017/031, http://cms.cern.ch/iCMS/jsp/db_notes/notestable1.jsp?CMSNoteID=DP-2017/031
- [17] CMS Collaboration, <https://twiki.cern.ch/twiki/bin/view/CMSPublic/HLTplotsSummary2016>