Available on CMS information server

CMS CR -2011/183

The Compact Muon Solenoid Experiment

Conference Report

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

19 September 2011

Operation and Performance of the CMS Level-1 Trigger during 7 TeV Collisions

Pamela Renee Klabbers for the CMS Collaboration

Abstract

The Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider (LHC) has been collecting data at center-of-mass energy 7 TeV since March 2010. CMS detects the products of proton beams colliding at a rate of 40 MHz. The Level-1 trigger reduces this collision rate to an output rate of 100 kHz, which is forwarded to the High-Level trigger, a dedicated computer farm, which reduces that further to a rate of 100 Hz, suitable for storage of full event data. The Level-1 trigger uses high-speed custom electronics to combine information from electromagnetic and hadronic calorimeters and three muon detection systems and identifies potential physics objects of interest in only a few microseconds. To ensure good performance of the Level-1 trigger hardware, robust configuration and monitoring software is also required. This talk will concentrate on the performance of the Level-1 trigger in the 2010 and ongoing 2011 collision runs, as well as presenting an overall picture of the hardware and operation.

Presented at *TIPP 2011: Technology and Instrumentation in Particle Physics 2011*

Available online at www.sciencedirect.com
ScienceDirect
Physics Procedia 00 (2011) 000–000
Physics Procedia
www.elsevier.com/locate/procedia

Technology and Instrumentation in Particle Physics 2011

Operation and performance of the CMS Level-1 Trigger during 7 TeV Collisions

Pamela Klabbers for the CMS Collaboration

University of Wisconsin, Madison, WI 53706, USA

Abstract

The Compact Muon Solenoid (CMS) experiment at the Large Hadron Collider (LHC) has been collecting data at center-of-mass energy 7 TeV since March 2010. CMS detects the products of LHC proton beams colliding at a rate of 40 MHz. The Level-1 trigger reduces this collision rate to an output rate of 100 kHz, which is forwarded to the High-Level trigger, a dedicated computer farm, which reduces that further to a rate of a few hundreds of Hz, suitable for storage of full event data. The Level-1 trigger uses high-speed custom electronics to combine information from electromagnetic and hadronic calorimeters and three muon detection systems and identifies potential physics objects of interest in only a few microseconds. To ensure good performance of the Level-1 trigger hardware, robust configuration and monitoring software is also required. The talk and these proceedings concentrate on the performance of the Level-1 trigger in the 2010 and on-going 2011 collision runs, as well as presenting an overall picture of the hardware and operation.

Keywords: Trigger, LHC, CMS

1. Introduction

The general-purpose Compact Muon Solenoid (CMS) Experiment is located at the Large Hadron Collider (LHC) at CERN near Geneva, Switzerland [1]. In March 2010, the LHC had its first collisions of protons on protons with a center-of-mass energy of 7 TeV. Gradually the instantaneous luminosity has increased, starting at 10^{27} and currently greater than 10^{33} $cm^{-2}s^{-1}$. At the time of the conference, CMS had recorded a total luminosity of almost 740 pb^{-1} with 1042 bunches colliding at CMS. Since then the number of colliding bunches has increased to 1318, and instantaneous luminosity is regularly greater than 2×10^{33} $cm^{-2}s^{-1}$. There are plans to increase it further, up to nearly 5×10^{33} $cm^{-2}s^{-1}$, by utilizing other machine and beam parameters such as bunch intensity, emittance, β^*, etc. This increasing instantaneous luminosity and other experimental factors have made commissioning and operating the CMS Level-1 Trigger (L1T) challenging. The CMS Level-1 Trigger and a summary of some of the measures taken to control the rate and ensure good quality L1 triggers are described in the talk and these proceedings.

1.1. The CMS Level-1 Trigger

The CMS Level-1 Trigger (L1T) is built almost entirely of custom electronics [2]. Its inputs come from the electromagnetic and hadronic calorimeters (ECAL and HCAL), muon detectors, and beam monitoring detectors. It is built to operate at the beam crossing frequency of 40 MHz. The objective of the CMS L1T is to reduce the 40 MHz rate to less than 100 kHz, the target input rate to the High Level Triggers (HLTs) while keeping events that are physically interesting. The L1T is pipelined for high data-taking efficiency with minimal deadtime. Including the front-end detector readout, the entire latency is 3.2 microseconds for collection, decision, and propagation back to the front-ends. A block diagram of the L1T is in Figure 1.

The calorimeter trigger (on the left side of Figure 1) is built in two stages, a Regional Calorimeter Trigger (RCT) and Global Calorimeter Trigger (GCT). The RCT receives Trigger Primitives (TPs) in the form of an 8-bit rank (E_T) and quality bit from the over 8000 towers in the Hadron, Forward, and Electromagnetic Calorimeters for $|\eta| < 5$. It processes this information in parallel and sends 18x8 e/γ candidates and 36 energy sums per 4x4 tower region and bits for object classification. The GCT sorts the e/γ candidates further and finds jets (Central, Forward, and Tau) and calculates global quantities like Missing E_T. It then sends eight e/γ of two types, four each of Central, Tau, and Forward Jets, Jet counts, and several global quantities to the CMS Global Trigger (GT).

The muon trigger (on the right side of Figure 1) is also in two stages, but divided differently, covering $|\eta| < 2.4$. Each muon subsystem takes chamber hits and performs track finding at 40 MHz in a dedicated trigger subsystem. The Drift Tubes (DTs, $|\eta| < 1.2$) and Cathode Strip Chambers (CSCs, $0.9 < |\eta| < 2.4$) create track segments from these hits and forward them to a final track-finder to select the top 4 muon candidates per subsystem. The Resistive Plate Chambers (RPCs, $|\eta| < 1.6$) use a different approach, a pattern comparison for 4 each of barrel and endcap muons. Sharing of track segments (shown with arrows between muon systems in Figure 1) among muon subsystems helps close geometric gaps in the muon coverage and raise the overall muon trigger efficiency. These 16 muon candidates are forwarded to the Global Muon Trigger (GMT) for further sorting and removing of duplicates. The GMT then sends the final top 4 global muon candidates including positional information to the Global Trigger [2].

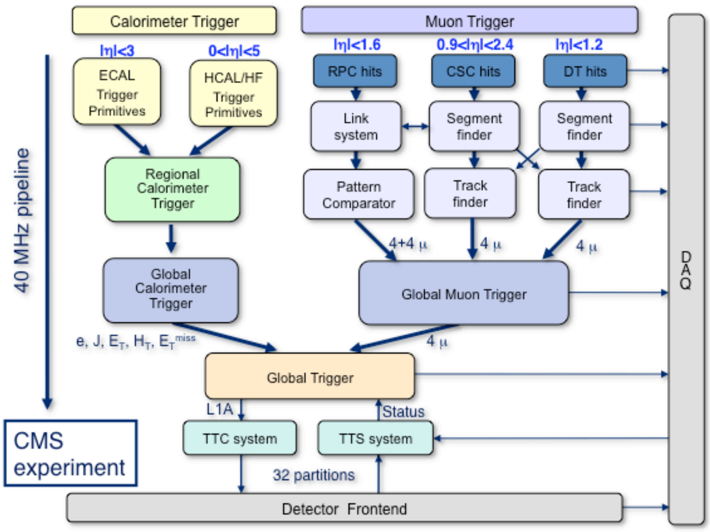

Fig. 1. CMS Level-1 Trigger block diagram: calorimeter, muon, and beam monitoring (not shown) participate in CMS L1T.

The Global Trigger is the final step before the L1 decision is made. It receives all the regional trigger objects and makes the final L1 decision. It can have up to 128 different physics algorithms, combining trigger objects, applying different thresholds, and making topological cuts, all configurable in firmware. There are also are inputs for up to 64 technical trigger signals provided by different CMS sub-detectors. Finally, it sends the Level-1 Accept to the Trigger Timing and Control (TTC) system for distribution to the detector front-end electronics [3]. The TTC interfaces to the LHC provide the orbit and clock information. The GT sends out commands via the TTC to keep the detectors and their electronics and links in sync, stop a run, start a run, etc.

1.2. Operation and monitoring

The CMS Trigger requires several custom software tools to operate and monitor. The primary interface, called the L1 Page, is accessed with a web browser (e.g. Firefox) by the trigger shifter and experts. It displays real time information about the trigger in the current run, as well as the status of the individual processes of the software controls (Trigger Supervisor, described below) for each trigger subsystem. Additionally it has drop-down menus that link to various tools that the trigger expert and shifter need for operation, configuration, and monitoring of the software and hardware of the L1 Trigger.

The Trigger Supervisor [4], linked through the L1 Page, is a software framework used to interface to all the individual trigger subsystem hardware. CMS Run Control, managed by the data acquisition, controls the global configuration of CMS during normal running and also shares state information with the Central Cell of the L1 Trigger Supervisor. This information is passed to all the subsystems, and includes a global key, the state of the run (stop, start, etc.) and information about the components included. In turn, the subsystems return information about their status to the Central Cell that then passes it on to the L1 Page. This can include errors and warnings regarding trigger subsystem components, or a failure to complete a transition of the run process due to hardware or software problems. Samples of the L1 Page, the Central Trigger Supervisor Cell and a monitoring page of the RCT are shown in Figure 2.

Another important part of the trigger monitoring is the online Data Quality Monitoring (DQM) [5]. During running, this tool uses raw data to make plots and histograms for all the CMS subsystems, including the L1 Trigger. A summary DQM web page (Figure 3) is the main interface for the Trigger Shifter, and is constantly updated as the run progresses. It displays occupancies for the trigger subsystems and can flag a plot bad if it doesn't agree with a reference histogram or is not within a defined range. Additionally, there are plots that are filled with comparisons of the real data with the L1 Trigger Emulator (L1TEMU). These provide clues as to what may be happening if there is a problem, and a summary of both the L1T and L1TEMU is sent to the main DQM shifter to monitor as well. This is shown for L1TEMU in Figure 3.

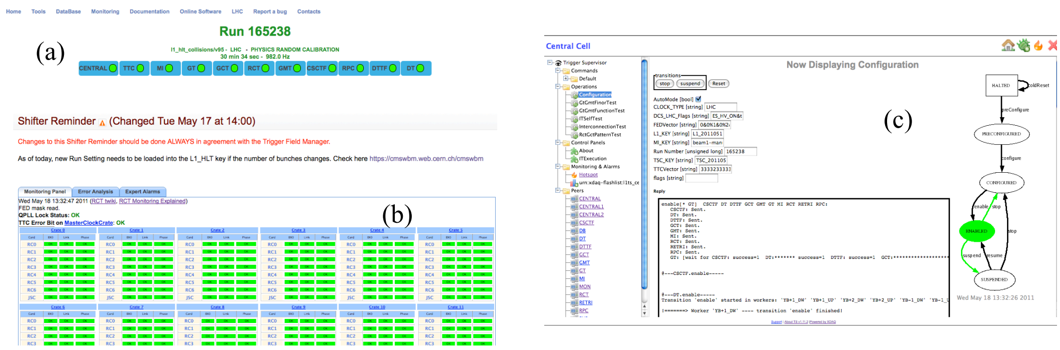

Fig. 2. (a) CMS Level-1 Page, (b) a sample monitoring page for the RCT, and (c) Central Cell of the Trigger Supervisor.

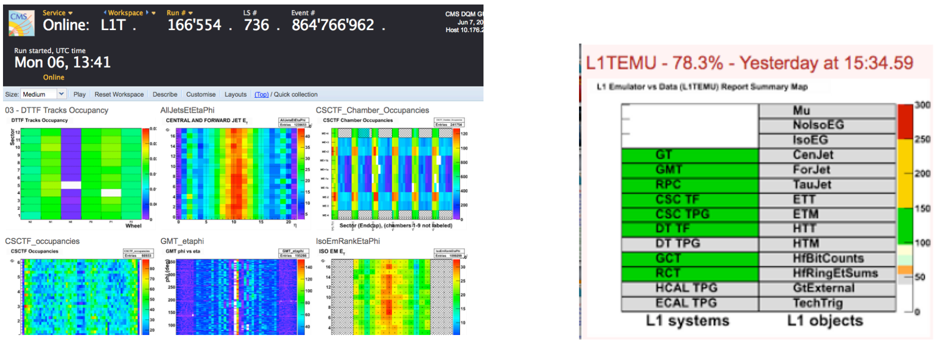

Fig. 3. CMS L1 Trigger DQM for the L1 Trigger shifter and the L1 Trigger Emulator Summary as seen by the main DQM shifter.

2. LHC First Collisions and Commissioning in 2010

On Tuesday, March 30, 2010 the LHC achieved first p-p collisions at 7 TeV. For this fill a simple Minimum Bias trigger was used to collect the data, time-in, and verify the performance of detector components. The LHC gradually increased the intensity until the last of the p-p running in October 2010.

2.1. Synchronization

In order to ensure proper readout of the many sub-detectors, an important first step was to synchronize the trigger. At CMS, collision products take a full bunch crossing longer to reach the outermost triggering detector, the CSCs, than to reach the innermost triggering detector, ECAL. Varying cable lengths must also be taken into account, even within a sub-detector. The first runs were taken with a simple Minimum Bias trigger using the Beam Scintillation Counters, and the detector-based trigger bits were recorded with the Minimum Bias triggers, but did not trigger readout. Timing scans were done, taking data with a number of delay settings for each sub-detector. The result of one of these scans is shown for the CSCs in Figure 4. To monitor for timing shift during running, crosschecks are always available in real time in DQM. A sample from a 2010 collision run is shown in Figure 4.

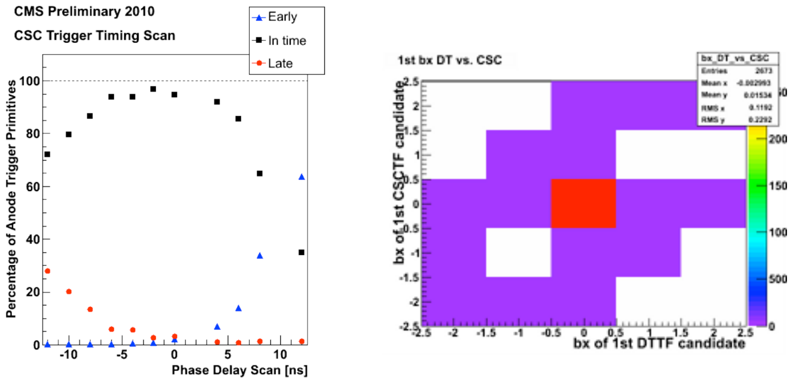

Fig. 4. Left: CMS Level-1 CSC Trigger Timing scan for synchronization of the trigger. Right: A cross-check of timing from the online DQM, DT vs. CSC timing relative to each other in terms of bunch crossing, centered at zero, showing good relative timing.

2.2. Rates

In the final month of the LHC 2010 p-p collisions, the maximum instantaneous luminosity increased to 2×10^{32} $cm^{-2}s^{-1}$. In this period we could no longer rely solely on a Minimum Bias trigger, since the rate was too high and had to be prescaled, i.e. the rate scaled down by an integer value. The L1T and HLT [6] groups worked together to come up with a set of triggers and rate limiting prescale factors to maintain a total rate to the HLT of about 65 kHz and a sustainable output rate of about 400 Hz. This combination of L1 and HLT triggers and prescales is called a trigger menu at CMS. Each L1 Trigger is used as a “seed” for a specific HLT path. Multiple “seeds” may be used in a single path.

To control the rates of the L1T, the thresholds were raised, conditions were tightened, and prescales adjusted. Several sets of prescales are tuned to allow as much data to tape as possible, and keep the thresholds as low as possible for physics. When the luminosity drops during a fill, the prescale factors can be reduced to keep the input rates about the same. This allows monitoring and lower threshold data for performance studies to be collected in sufficient quantities.

To set up these initial menus, Monte Carlo was used to predict the output rates of various triggers. Further tuning of the rates with data was necessary to take into account additional factors like cosmic ray backgrounds (muon triggers), detector response (noise), pile-up effects, and so on.

2.3. Performance in 2010

Measurements of trigger efficiency in 2010 are shown in Figure 5 for muons and e/γ triggers. The performance of subsystems was very good during this first period of proton physics. Each subsystem/trigger type used a different technique to measure the efficiency and each is briefly described in the figure caption.

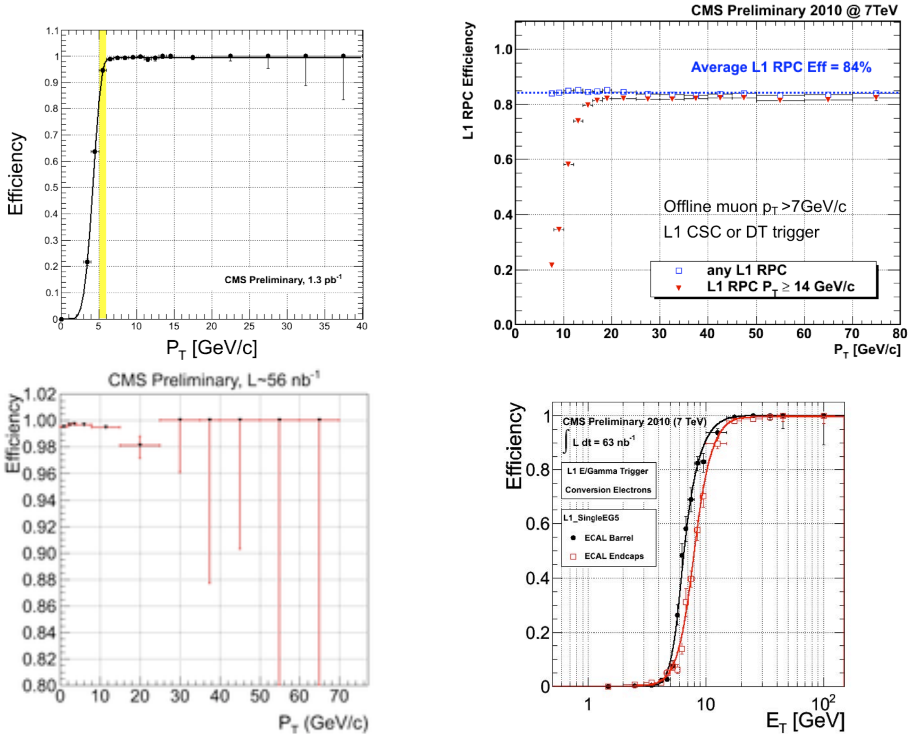

Fig. 5. From left to right, top to bottom: CMS Level-1 Trigger Efficiencies in 2010 for the DT Track Finder (DTTF) using J/psi to two muons with a tag-and-probe method, RPC using global muons (triggered muons), CSC Track Finder (CSCTF) using global muons, and the e/gamma (calorimeter) using Z boson decays to two electrons with a tag-and-probe analysis.

3. Collisions in 2011

Collisions in 2011 were even more challenging. The LHC increased the luminosity rapidly this year, the bunch spacing was reduced from 75 to 50 ns, and the trigger group had to keep rates under control and keep interesting events for the physics analysis.

Similar to 2010, trigger menus were developed to keep the L1T and HLT rates under control. Typically the total rate was kept around 70 kHz at L1 and 300-400 Hz out of the HLT was written to tape. In addition to raising the thresholds, employing prescales, and setting physics priorities, a number of detector issues were addressed to ensure good quality triggers and suppress backgrounds.

For the e/γ triggers a few different approaches were used. Tower level, η-dependent energy corrections to improve resolution were implemented using Monte Carlo electron data. This had the effect of making the triggers reach full efficiency at lower values of E_T. The ECAL implemented spike killing at the Trigger Primitive level. Energy deposited by heavily ionizing particles in ECAL barrel single-crystal avalanche photodiodes look like very energetic isolated e/γ candidates. Crystals around a suspect spike are checked for additional energy and the Trigger Primitive zeroed if it is isolated to a single crystal. At the RCT the Fine Grain (FG) veto, requiring energy deposit (from an electron bending in a magnetic field) in a tower to mostly within a narrow strip in η, and $E_{T,HCAL}/E_{T,ECAL}$ (H/E) veto, requiring most of the energy deposit in the ECAL, were applied for the e/γ candidate to the tower with maximum E_T.

For the jet triggers the resolution was improved with jet energy corrections and its response made closer to that of the e/γ triggers. The online jets were corrected to Monte Carlo derived factors and a correction matrix in η and p_T applied to jets at the GCT level, resulting in better jet trigger performance.

The muon triggers also benefited from improvements. A "ghost" muon occurs when the muon tracking algorithms find false muons due to detector geometry or timing effects, e.g. cracks between chambers, drift time, etc. The ghost muons are difficult to separate out, and vary from sub-system to sub-system, depending on the configuration, so a number of different "ghost busting" algorithms have been applied to reduce the rate. Additionally, the RPC muon trigger improved pattern recognition. They required more layers in pattern matching using 4/6 layers instead of 3/6, reducing the rate.

And finally, development of the L1-HLT menus with prescale sets has continued this year, adjusting thresholds and prescale factors to take into account the needs for physics as beam conditions changed and luminosity rose.

In spring 2011, the LHC moved from 75 to 50 ns proton bunch spacing. This had the desired effect of increasing the delivered luminosity. There was already pre-firing of the trigger, caused by the forward calorimeters, and at this time it was addressed. Since firing on the bunch crossing before vetoes a trigger in the following bunch crossing, this is a source of inefficiency (about 5%). To fix this, CMS used the BPTX_AND trigger. The BPTX_AND trigger is a AND of two precision electrostatic beam pick-ups near the beam pipe, on opposite sides of CMS. This was already available, providing the LHC colliding bunch structure for every fill.

The signal was copied, its delay reduced by 1 bunch crossing, and its signal was shortened to ~20 ns to eliminate the possibility of it vetoing the next bunch crossing, and it was used to veto the pre-triggers. The effect can be seen in the bunch structure plots in Figure 6, before and after the implementation of the veto. Unfortunately, with 50 ns bunch spacing, this had the undesired effect of vetoing the Heavy Stable Charged Particle (HSCP) trigger, which is set up to trigger in the bunch crossing after the collision.

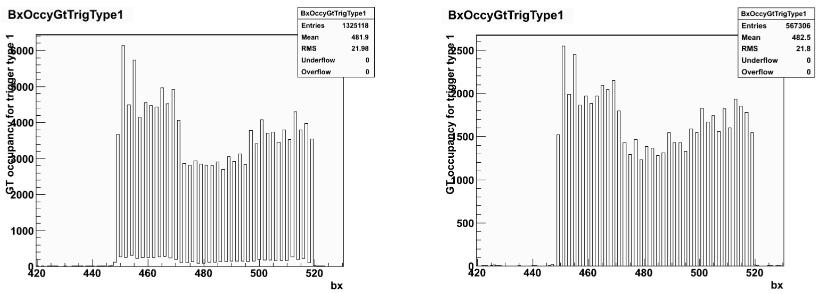

Fig. 6. Triggers as a function of bunch crossing, before (left) and after (right) the implementation of the BPTX veto.

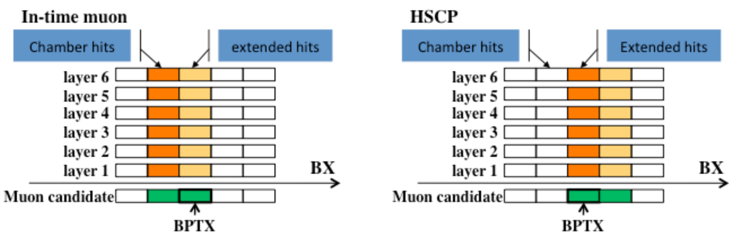

Fig. 7. A diagram showing an in-time and out-of-time muon captured in the collision bunch crossing (labeled as BPTX).

A remedy was found by using a feature of the RPC trigger, which is able to adjust its signal length, so single bunch crossing signals were extended over two bunch crossings. This, along with reducing the delay of the RPC trigger at the Global Muon Trigger by one bunch crossing, allowed for both in-time and delayed muons to trigger in the primary (not vetoed) bunch crossing. A diagram showing how it worked is in Figure 7.

3.1. Trigger Performance

Preliminary trigger performance plots for 2011 are shown for the CSC Track Finder (CSCTF) and the e/γ trigger in Figure 8. The CSCTF performs better than in 2010 for the low P_T range of the high cutoffs due to improvements in the algorithms, and the e/γ trigger is performing as expected.

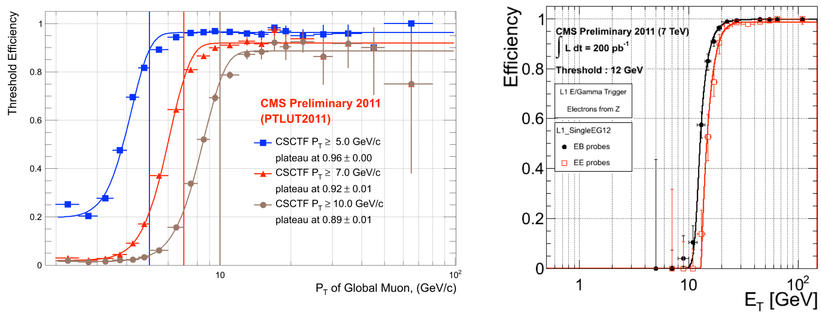

Fig. 8. CMS Level-1 Trigger efficiencies in 2011 for the CSCTF and the e/gamma (calorimeter), using the same techniques as in Figure 4.

4. Conclusions

The CMS Level-1 Trigger has risen to the challenges of operating in the new high-luminosity conditions of the LHC. By carefully monitoring the trigger, good performance has been maintained. Overall careful planning of the trigger groups has controlled rates. The challenges continue in the following months it is expected that the LHC will attain a luminosity close to $5\times10^{33}cm^{-2}s^{-1}$ in the remainder of the 2011 running.

References

[1] CMS Collaboration, *The CMS experiment at the CERN LHC*, JINST 3:S08004, 2008.

[2] CMS Collaboration, The TRIDAS Project Technical Design Report, Volume 1: The Trigger Systems, CERN/LHCC 2000-38, CMS TDR 6.1.

[3] http://ttc.web.cern.ch/TTC/intro.html

[4] M. Magrans de Abril, et al., *Homogeneous User Interface Infrastructure for Expert Control of the Level-1 Trigger*, CMS-NOTE-2010-005.

[5] S. Morovic, *CMS Online Data Quality Monitoring: Real-Time Event Processing Infrastructure*, CHEP 2010 Proceedings, Taipei, Taiwan.

[6] L. Apanasevich, *Commissioning and Performance of the CMS High Level Trigger*, these proceedings.