

The CMS trigger system



The CMS collaboration

E-mail: cms-publication-committee-chair@cern.ch

ABSTRACT: This paper describes the CMS trigger system and its performance during Run 1 of the LHC. The trigger system consists of two levels designed to select events of potential physics interest from a GHz (MHz) interaction rate of proton-proton (heavy ion) collisions. The first level of the trigger is implemented in hardware, and selects events containing detector signals consistent with an electron, photon, muon, τ lepton, jet, or missing transverse energy. A programmable menu of up to 128 object-based algorithms is used to select events for subsequent processing. The trigger thresholds are adjusted to the LHC instantaneous luminosity during data taking in order to restrict the output rate to 100 kHz, the upper limit imposed by the CMS readout electronics. The second level, implemented in software, further refines the purity of the output stream, selecting an average rate of 400 Hz for offline event storage. The objectives, strategy and performance of the trigger system during the LHC Run 1 are described.

KEYWORDS: Trigger concepts and systems (hardware and software); Trigger detectors; Data acquisition circuits

ARXIV EPRINT: [1609.02366](https://arxiv.org/abs/1609.02366)



Contents

1	Introduction	1
1.1	The CMS detector	3
2	The trigger system	3
2.1	The L1 trigger overview	3
2.2	The L1 calorimeter trigger system	5
2.2.1	The ECAL trigger primitives	5
2.2.2	HCAL trigger primitives	6
2.2.3	Regional calorimeter trigger system	6
2.2.4	Global calorimeter trigger system	7
2.3	The L1 muon trigger system	11
2.3.1	Muon local trigger segments	11
2.3.2	Drift tube track finder	12
2.3.3	Cathode strip chambers track finder	12
2.3.4	Resistive plate chambers trigger system	13
2.3.5	Global muon trigger system	14
2.4	The L1 global trigger system	15
2.5	Beam position timing trigger system	17
2.6	High-level trigger system	17
3	Object identification	18
3.1	Tracking and vertex finding	19
3.1.1	Primary vertex reconstruction	19
3.1.2	HLT tracking	21
3.2	Electron and photon triggers	22
3.2.1	L1 electron/photon identification	24
3.3	Online anomalous signals and their suppression	28
3.3.1	HLT electron and photon identification	31
3.4	Muon triggers	35
3.4.1	The L1 muon trigger performance	35
3.4.2	HLT muon identification	41
3.5	Jets and global energy sums	45
3.5.1	The L1 jet trigger	46
3.5.2	The L1 energy sums	48
3.5.3	L1 jet and energy sum rates	50
3.5.4	The HLT jet triggers	50
3.5.5	The HLT E_T^{miss} triggers	51
3.6	τ lepton triggers	54
3.6.1	The L1 τ identification	54
3.6.2	The HLT τ lepton identification	55

3.7	b-quark jet tagging	58
3.7.1	Tracking for b tagging	59
3.7.2	Performance of online b-tagging	59
3.8	Heavy ion triggers	60
4	Physics performance of the trigger	65
4.1	Higgs boson physics triggers	65
4.1.1	Triggers for Higgs boson diphoton analysis	65
4.1.2	Triggers for multi-lepton Higgs boson analyses	66
4.1.3	Triggers for the di-tau Higgs boson analysis	68
4.1.4	Triggers for ZH to 2 neutrinos + b jets analysis	68
4.2	Top quark triggers	70
4.3	Triggers for supersymmetry searches	72
4.3.1	Triggers for all-hadronic events with α_T	72
4.3.2	Triggers for inclusive search with Razor variables	74
4.3.3	Triggers for photons and missing energy	75
4.3.4	Triggers for heavy stable charged particles	77
4.4	Exotic new physics scenarios	78
4.4.1	Triggers for dijet resonance searches	78
4.4.2	Triggers for black hole search	79
4.5	B physics and quarkonia triggers	81
5	Trigger menus	84
5.1	L1 menus	85
5.1.1	Menu development	85
5.2	HLT menus	88
6	Trigger system operation and evolution	91
6.1	Trigger monitoring and operations	91
6.2	Technical performance	91
6.2.1	The L1 trigger deadtime, downtime and reliability	91
6.2.2	The HLT resources and optimization	92
6.2.3	The HLT operations	94
7	Summary	95
	The CMS collaboration	102

1 Introduction

The Compact Muon Solenoid (CMS) [1] is a multipurpose detector designed for the precision measurement of leptons, photons, and jets, among other physics objects, in proton-proton as well as heavy ion collisions at the CERN LHC [2]. The LHC is designed to collide protons at a center-of-mass energy of 14 TeV and a luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. At design luminosity, the pp interaction rate exceeds 1 GHz. Only a small fraction of these collisions contain events of interest to the CMS physics program, and only a small fraction of those can be stored for later offline analysis. It is the job of the trigger system to select the interesting events for offline storage from the bulk of the inelastic collision events.

To select events of potential physics interest [3], the CMS trigger utilizes two levels while, for comparison, ATLAS uses a three-tiered system [4]. The first level (L1) of the CMS trigger is implemented in custom hardware, and selects events containing candidate objects, e.g., ionization deposits consistent with a muon, or energy clusters consistent with an electron, photon, τ lepton, missing transverse energy (E_T^{miss}), or jet. Collisions with possibly large momentum transfer can be selected by, e.g., using the scalar sum of the jet transverse momenta (H_T).

The final event selection is based on a programmable menu where, by means of up to 128 algorithms utilizing those candidate objects, events are passed to the second level (high-level trigger, HLT). The thresholds of the first level are adjusted during data taking in response to the value of the LHC instantaneous luminosity so as to restrict the output rate to 100 kHz [3], the upper limit imposed by the CMS readout electronics. The HLT, implemented in software, further refines the purity of the physics objects, and selects an average rate of 400 Hz for offline storage. The overall output rate of the L1 trigger and HLT can be adjusted by prescaling the number of events that pass the selection criteria of specific algorithms. In addition to collecting collision data, the trigger and data acquisition systems record information for the monitoring of the detector.

After commissioning periods at 0.9 and 2.36 TeV in 2009, the first long running periods were at a center-of-mass energy of 7 TeV in 2010 and 2011, and 8 TeV in 2012. These proton-proton data, together with the first ion running periods (PbPb at 2.76 TeV, and pPb at 5.02 TeV), are referred to collectively as Run 1. During this period, the CMS trigger system selected interesting pp physics events at maximum instantaneous luminosities of $2.1 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ (2010), $4 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ (2011), and $7.7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ (2012), corresponding to 0.2, 4, and 7.7 Hz nb $^{-1}$. Figure 1 shows the pp integrated and peak luminosities as a function of time for calendar years 2010, 2011 and 2012. While the nominal bunch crossing (BX) frequency is 40 MHz, corresponding to 25 ns between individual bunch collisions, the bunch spacing during regular running was never less than 50 ns through Run 1. The highest number of collisions per BX (known as “pileup”) averaged over a data run in 2011 and 2012 was 16.15 and 34.55, respectively, while the pileup averages over the year were 9 (21) in 2011 (2012).

The trigger system is also used during heavy ion running. The conditions for PbPb collisions are significantly different from those in the pp case. The instantaneous luminosity delivered by the LHC in the 2010 (2011) PbPb running period was 3×10^{25} (5×10^{26}) $\text{cm}^{-2} \text{ s}^{-1}$, resulting in maximum interaction rates of 250 Hz (4 kHz), much lower than in pp running, with a negligible pileup probability and an inter-bunch spacing of 500 ns (200 ns). During the pPb run in 2013, an instantaneous luminosity of $10^{29} \text{ cm}^{-2} \text{ s}^{-1}$ was achieved, corresponding to an interaction rate of




Figure 1. Integrated (top) and peak (bottom) proton-proton luminosities as a function of time for calendar years 2010–2012. The 2010 integrated (instantaneous) luminosity is multiplied by a factor of 100 (10). In the lower plot, 1 Hz/nb corresponds to $10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

200 kHz, again with a very low pileup probability. Due to the large data size in these events, the readout rate of the detector is limited to 3 kHz in heavy ion collisions.

This document is organized as follows. Section 2 describes the CMS trigger system (L1 and HLT) in detail. Section 3 gives an overview of the methods, algorithms, and logic used to identify physics signatures of interest in LHC collisions, and to select events accordingly. The physics performance achieved with the CMS trigger system is outlined in section 4 based on examples of several physics analyses. In section 5, details of the L1 and HLT menus are given, together with the objectives and strategies to assemble those menus. The operation and evolution of the trigger system during the first years of the LHC running is described in section 6. A summary is given in section 7.

1.1 The CMS detector

The central feature of the CMS apparatus is a superconducting solenoid, of 6 m internal diameter, providing a magnetic field of 3.8 T. Within the superconducting solenoid volume are a silicon pixel and strip tracker, a lead tungstate crystal electromagnetic calorimeter (ECAL), and a brass/scintillator hadron calorimeter (HCAL). Muons are measured in gas-ionization detectors embedded in the steel return yoke. Extensive forward calorimetry complements the coverage provided by the barrel and endcap detectors. The missing transverse momentum vector is defined as the projection on the plane perpendicular to the beams of the negative vector sum of the momenta of all reconstructed particles in an event. Its magnitude is referred to as E_T^{miss} . The transverse momentum vector is defined as the projection on the plane perpendicular to the beams of the negative vector sum of the momenta of all reconstructed particles in an event. Its magnitude is referred to as E_T . A more detailed description of the CMS detector, together with a definition of the coordinate system used and the relevant kinematic variables, can be found in ref. [1].

2 The trigger system

The trigger system is comprised of an L1 hardware trigger and an HLT array of commercially available computers running high-level physics algorithms. In this section we describe the design of the combined L1-HLT system.

2.1 The L1 trigger overview

The L1 trigger is a hardware system with a fixed latency. Within 4 μs of a collision, the system must decide if an event should be tentatively accepted or rejected using information from the calorimeter and muon detectors.

A schematic of the L1 trigger is shown in figure 2. The trigger primitives (TP) from electromagnetic and hadron calorimeters (ECAL and HCAL) and from the muon detectors (drift tubes (DT), cathode strip chambers (CSC) and resistive-plate chambers (RPC)) are processed in several steps before the combined event information is evaluated in the global trigger (GT) and a decision is made whether to accept the event or not.

The L1 calorimeter trigger comprises two stages, a regional calorimeter trigger (RCT) and a global calorimeter trigger (GCT). The RCT receives the transverse energies and quality flags from over 8000 ECAL and HCAL towers (section 2.2.1 and 2.2.2), giving trigger coverage over $|\eta| < 5$. The RCT processes this information in parallel and sends as output e/γ candidates and regional E_T




Figure 2. Overview of the CMS L1 trigger system. Data from the forward (HF) and barrel (HCAL) hadronic calorimeters, and from the electromagnetic calorimeter (ECAL), are processed first regionally (RCT) and then globally (GCT). Energy deposits (hits) from the resistive-plate chambers (RPC), cathode strip chambers (CSC), and drift tubes (DT) are processed either via a pattern comparator or via a system of segment- and track-finders and sent onwards to a global muon trigger (GMT). The information from the GCT and GMT is combined in a global trigger (GT), which makes the final trigger decision. This decision is sent to the tracker (TRK), ECAL, HCAL or muon systems (MU) via the trigger, timing and control (TTC) system. The data acquisition system (DAQ) reads data from various subsystems for offline storage. MIP stands for minimum-ionizing particle.

sums based on 4×4 towers [5]. The GCT sorts the e/γ candidates further, finds jets (classified as central, forward, and tau) using the E_T sums, and calculates global quantities such as E_T^{miss} . It sends as output four e/γ candidates each of two types, isolated and nonisolated, four each of central, tau, and forward jets, and several global quantities.

Each of the three muon detector systems in CMS participates in the L1 muon trigger to ensure good coverage and redundancy. For the DT and CSC systems ($|\eta| < 1.2$ and $|\eta| > 0.9$, respectively), the front-end trigger electronics identifies track segments from the hit information registered in multiple detector planes of a single measurement station. These segments are collected and then transmitted via optical fibers to regional track finders in the electronics service cavern, which then applies pattern recognition algorithms that identifies muon candidates and measure their momenta from the amount they bend in the magnetic field of the flux-return yoke of the solenoid. Information is shared between the DT track finder (DTTF) and CSC track finder (CSCTF) for efficient coverage in the region of overlap between the two systems at $|\eta| \approx 1$. The hits from the RPCs ($|\eta| < 1.6$) are directly sent from the front-end electronics to pattern comparator trigger (PACT) logic boards that identify muon candidates. The three regional track finders sort the identified muon candidates and

transmit to the global muon trigger (GMT) up to 4 (CSCTF, DTTF) or 8 (RPC) candidates every bunch crossing. Each candidate is assigned a p_T and quality code as well as an (η, ϕ) position in the muon system (with a granularity of ≈ 0.05). The GMT then merges muon candidates found by more than one system to eliminate a single candidate passing multiple-muon triggers (with several options on how to select p_T between the candidates). The GMT also performs a further quality assignment so that, at the final trigger stage, candidates can be discarded if their quality is low and they are reconstructed only by one muon track finder.

The GT is the final step of the CMS L1 trigger system and implements a menu of triggers, a set of selection requirements applied to the final list of objects (i.e., electrons/photons, muons, jets, or τ leptons), required by the algorithms of the HLT algorithms to meet the physics data-taking objectives. This menu includes trigger criteria ranging from simple single-object selections with E_T above a preset threshold to selections requiring coincidences of several objects with topological conditions among them. A maximum of 128 separate selections can be implemented in a menu.

2.2 The L1 calorimeter trigger system

The following is the description of the reconstruction of ECAL and HCAL energy deposits used in the L1 trigger chain followed by describing the RCT and GCT processing steps operating on these trigger primitives.

2.2.1 The ECAL trigger primitives

The ECAL trigger primitives are computed from a barrel (EB) and two endcaps (EE), comprising 75 848 lead tungstate (PbWO_4) scintillating crystals equipped with avalanche photodiode (APD) or vacuum phototriode (VPT) light detectors in the EB and EE, respectively. A preshower detector (ES), based on silicon sensors, is placed in front of the endcap crystals to aid particle identification. The ECAL is highly segmented, is radiation tolerant and has a compact and hermetic structure, covering the pseudorapidity range of $|\eta| < 3.0$. Its target resolution is 0.5% for high-energy electrons/photons. It provides excellent identification and energy measurements of electrons and photons, which are crucial to searches for many new physics signatures. In the EB, five strips of five crystals (along the azimuthal direction) are combined into trigger towers (TTs) forming a 5×5 array of crystals. The transverse energy detected by the crystals in a single TT is summed into a TP by the front-end electronics and sent to off-detector trigger concentrator cards (TCC) via optical fibers. In the EE, trigger primitive computation is completed in the TCCs, which must perform a mapping between the collected pseudo-strips trigger data from the different supercrystals and the associated trigger towers.

Mitigation of crystal transparency changes at the trigger level. Under irradiation, the ECAL crystals lose some of their transparency, part of which is recovered when the radiation exposure stops (e.g., between LHC fills). The effect of this is that the response of the ECAL varies with time. This variation is accounted for by the use of a laser system that frequently monitors the transparency of each crystal [6] and allows for offline corrections to the measured energies to be made [7]. In 2011, the levels of radiation in ECAL were quite small, and no corrections to the response were made at L1. From 2012 onwards, where the response losses were larger, particularly in the EE,

corrections to the TT energies were calculated and applied on a weekly basis in order to maintain high trigger efficiency and low trigger thresholds.

2.2.2 HCAL trigger primitives

The HCAL TPs are computed out of the digital samples of the detector pulses by the trigger primitive generator (TPG). In the barrel, one trigger primitive corresponds to one HCAL readout, whereas raw data from the two depth-segmented detector readout elements are summed in the endcap hadron calorimeter. For the forward hadron calorimeter (HF), up to 12 readouts are summed to form one trigger primitive. One of the most important tasks of the TPG is to assign a precise bunch crossing to detector pulses, which span over several clock periods. The bunch crossing assignment uses a digital filtering technique applied to the energy samples, followed by a peak finder algorithm. The amplitude filters are realized using a sliding sum of 2 consecutive samples. A single sample is used for HF where the signals are faster. The peak finder selects those samples of the filtered pulse that are larger than the two nearest neighbors. The amplitudes of the peak and peak+1 time slices are used as an estimator of the pulse energy. The position of the peak-filtered sample in the data pipeline flow determines the timing. The transverse energy of each HCAL trigger tower is calculated on a 10-bit linear scale. In case of overflow, the E_T is set to the scale maximum. Before transmission to the RCT, the 10-bit trigger tower E_T is converted to a programmable 8-bit compressed nonlinear scale in order to minimize the trigger data flux to the regional trigger. This data compression leads to a degradation in the trigger energy resolution of less than 5%. The energy in GeV is obtained from the ADC count by converting the ADC count into fC, subtracting the pedestal and correcting for the gain of each individual channel. Finally, a correction factor is applied to compensate for the fraction of signal charge not captured in the two time-slice sum.

2.2.3 Regional calorimeter trigger system

The CMS L1 electron/photon (e/γ), τ lepton, jet, H_T (where $H_T = \sum p_T^{\text{jets}}$ is the scalar sum of the p_T of all jets with $p_T > 10$ GeV and $|\eta| < 3$), and missing E_T trigger decisions are based on input from the L1 regional calorimeter trigger (RCT) [5, 8–10]. Eighteen crates of custom RCT electronics process data for the barrel, endcap, and forward calorimeters, with a separate crate for LHC clock distribution.

Twenty-four bits comprising two 8-bit calorimeter energies, either two ECAL fine-grain (FG) bits or two HCAL minimum ionizing particle (MIP) bits, an LHC bunch crossing bit, and 5 bits of error detection code, are sent from the ECAL, HCAL, and HF calorimeter back-end electronics to the nearby RCT racks on 1.2 Gbaud copper links. This is done using one of the four 24-bit channels of the Vitesse 7216-1 serial transceiver chip on the calorimeter output and the RCT input, for 8 channels of calorimeter data per chip. The RCT V7216-1 chips are mounted on receiver mezzanine cards located on each of 7 receiver cards (RC) and the single-jet summary cards (JSC) for all 18 RCT crates.

The RCT design includes five high-speed custom GaAs application-specific integrated circuits (ASICs), which were designed and manufactured by Vitesse Semiconductor: a phase ASIC, an adder ASIC, a boundary scan ASIC, a sort ASIC, and an electron isolation ASIC [11].

The RC has eight receiver mezzanine cards for the HCAL and ECAL data, four per subsystem. On the mezzanine, the V7216-1 converts the serial data to 120 MHz TTL parallel data. Eight phase

ASICs on the RC align and synchronize the data received on four channels of parallel data from the Vitesse 7216-1, check for data transmission errors, and convert 120 MHz TTL to 160 MHz emitter-coupled logic (ECL) parallel data. Lookup tables (LUTs) convert 17 bits of input (8 bits from ECAL, HCAL and the FG bit) for two separate paths. They rescale the incoming ECAL energies, and set quality bits for the e/γ path (a tower-level logical OR of the ECAL FG bits and a limit on fractional energy in the HCAL), and rescale and sum HCAL and ECAL for the regional sums path. On the RC, the boundary scan ASIC aligns the e/γ tower energy data with data shared on cables between RCT crates adjacent in η and ϕ , and makes copies so that each of 7 electron isolation cards (EIC) receives 28 central and 32 adjacent towers via the custom 160 MHz backplane. The HCAL+ECAL summed towers are added together to form 4×4 trigger tower sums by three adder ASICs, which sum up eight 11-bit energies in 25 ns, while providing bits for overflows. The tower sums are then sent to the JSC via the backplane for further processing. A logical OR of the MIP bits over the same 4×4 trigger tower regions is sent to the JSC.

The EIC receives the 32 central tower and 28 neighboring trigger tower data from the RCs via the backplane. The electron isolation algorithm is implemented in the electron isolation ASIC, which can handle four 7-bit electromagnetic energies, a veto bit, and nearest neighbor energies every 6.25 ns. It finds up to four electron candidates in two 4×4 trigger tower regions, two isolated and two non-isolated. These candidates are then transmitted via the backplane to the JSC for further processing. In this way the e/γ algorithm is seamless across the entire calorimeter.

The JSC receives 28 e/γ candidates, 14 sums, and has a single mezzanine card to receive eight HF TPs and quality bits. The JSC rescales the HF data using a lookup table and delays the data so that it is in time with the 14 regional E_T sums when they are sent to the GCT for the jet finding and calculation of global quantities such as H_T and missing E_T . In addition, for muon isolation, a quiet bit is set for each region and forwarded with the MIP bits on the same cables as the electron candidates. The 28 electron candidates (14 isolated and non-isolated) are sorted in E_T in two stages of sort ASICs on the JSC, and the top four of each type are transmitted to the GCT for further sorting. A block diagram of this dataflow is shown in figure 3.

Finally, a master clock crate (MCC) and cards are located in one of the ten RCT racks to provide clock and control signal distribution. Input to the system is provided by the CMS trigger timing and control (TTC) system. This provides the LHC clock, bunch crossing zero (BC0), and other CMS synchronization signals via an optical fiber from a TTC VME interface board which can internally generate or receive these signals from either a local trigger controller board (LTC) or from the CMS GT.

The MCC includes a clock input card (CIC) with an LHC TTC receiver mezzanine (TTCrm) to receive the TTC clocks and signals via the fiber and set the global alignment of the signals. The CIC feeds fan-out cards, a clock fan-out card midlevel (CFCm) and a clock fan-out card to crates (CFCc) to align and distribute the signals to the individual crates via low-skew cable. Adjustable delays on these two cards allow fine-tuning of the signals to the individual crates.

2.2.4 Global calorimeter trigger system

The GCT is the last stage of the L1 calorimeter trigger chain. A detailed description of the GCT design, implementation and commissioning is provided in several conference papers [12–17] that describe the changes in design since the CMS trigger technical design report [5].




Figure 3. Block diagram of the regional calorimeter trigger (RCT) system showing the data flow through the different cards in a RCT crate. At the top is the input from the calorimeters; at the bottom is the data transmitted to the global calorimeter trigger (GCT). Data exchanged on the backplane is shown as arrows between cards. Data from neighboring towers come via the backplane, but may come over cables from adjoining crates.

The trigger objects computed by the GCT from data supplied by the RCT are listed below and described in subsequent paragraphs:

- four isolated and four non-isolated electrons/photons of highest transverse energy;
- four central, four forward, and four tau jets of highest transverse energy;
- total transverse energy (S_T),

$$S_T \equiv \sum E_T,$$

calculated as the scalar sum of the E_T of all calorimeter deposits; H_T (see section 1); and (E_T^{miss});

- missing jet transverse energy; summing of feature bits and transverse energies in the HF calorimeter.

The electron/photon sort operation must determine the four highest transverse energy objects from 72 candidates supplied by the RCT, for both isolated and non-isolated electrons/photons.

To sort the jets, the GCT must first perform jet finding and calibrate the clustered jet energies. The jets are created from the 396 regional transverse energy sums supplied by the RCT. These are the sum of contributions from both the hadron and electromagnetic calorimeters. This is a substantial extension of the GCT capability beyond that specified in ref. [5]. The jet finding and subsequent sort is challenging because of the large data volume and the need to share or duplicate data between processing regions to perform cluster finding. The latter can require data flows of

a similar magnitude to the incoming data volume depending on the clustering method used. The clusters, defined as the sum of 3×3 regions, are located using a new method [13] that requires substantially less data sharing than the previously proposed sliding window method [18]. Jets are subdivided into central, forward, and tau jets based on the RCT tau veto bits and the jet pseudorapidity.

The GCT must also calculate some additional quantities. The total transverse energy is the sum of all regional transverse energies. The total missing transverse energy E_T^{miss} is calculated by splitting the regional transverse energy values into their x and y components and summing the components in quadrature. The resulting vector, after a rotation of 180° , provides the magnitude and angle of the missing energy. The jet transverse energy H_T and missing jet transverse energy are the corresponding sums over all clustered jets found.

Finally two quantities are calculated for the forward calorimeters. The transverse energy is summed for the two rings of regions closest to the beam pipe in both positive and negative pseudorapidities. The number of regions in the same rings with the fine-grain bit is also counted.

In addition to these tasks the GCT acts as a readout device for both itself and the RCT by storing information until receipt of an L1 accept (L1A) and then sending the information to the DAQ.

The GCT input data volume and processing requirements did not allow all data to be concentrated in one processing unit. Thus, many large field programmable gate arrays (FPGA) across multiple discrete electronics cards are necessary to reduce the data volume in stages. The cards must be connected together to allow data sharing and to eventually concentrate the data into a single location for the sort algorithms.

The latency allowed is 24 bunch crossings for jets and 15 bunch crossings for electrons/photons. Using many layers of high-speed serial links to transport the large data volumes between FPGAs was not possible since these typically require several clock cycles to serialize/deserialize the data and thus they have to be used sparingly to keep the latency low. The final architecture uses high-speed optical links (1.6 Gb/sec) to transmit the data and then concentrates the data in the main processing FPGAs, followed by standard FPGA I/O to connect to downstream FPGAs.

Figure 4 shows a diagram of the GCT system data flow. The input to the GCT is 18 RCT crates. The 63 source cards retransmit the data on optical high-speed serial links (shown by dashed arrows). For each RCT crate, the electron data are transmitted on 3 fibers and the jet data on 10 fibers. There are two main trigger data paths: electron and jet.

The jet data are sent to leaf cards (configured for jet finding) mounted on the wheel cards. The leaf cards are connected in a circle to search for clustered jets in one half of the CMS calorimeter (either in the positive or the negative η). The wheel card collects the results from three leaf cards, sorts the clustered jets, and forwards the data to the concentrator card. A more detailed description of each component is given below.

- Source card. The 6 differential ECL cables per RCT crate are fed into source cards, each receiving up to two RCT cables and transmitting the data over four fiber links. This has several advantages: it allows the source cards to be electrically isolated from the main GCT system, the different data within the RCT cables to be rearranged, a large amount of information to be concentrated so that it can be delivered to the processing FPGAs on leaf cards, and data to be duplicated.




Figure 4. A schematic of the global calorimeter trigger (GCT) system, showing the data flow through the various component cards.

- **Leaf card.** The leaf card is the main processing block in the GCT design. The most difficult task in the GCT is the jet finding. This is made simpler by concentrating the data in as few FPGAs as possible. Consequently, each leaf card has two Xilinx Virtex II Pro FPGAs each with 16 multi-gigabit transceivers that are used to bring the raw data in. Three Agilent 12 channel receivers provide the opto-electronic interface. The large standard I/O capacity is used to transmit the data to the wheel card.
- **Wheel card.** There are two wheel cards, one for each half of the detector. They act as carriers for three leaf cards and further concentrate the data. They sum the energy values, sort the 54 clustered jets by transverse energy into the three types (forward, central, tau). The wheel cards then forward the information to the concentrator card via high-speed Samtec low-voltage differential signal (LVDS) cables.
- **Concentrator card.** The concentrator card performs similar actions to that of the wheel card after which it transmits the resulting trigger objects to the GT and stores the information in a pipeline until receipt of an L1A signal. The concentrator card also carries two leaf cards that process the electron data. These leaf cards record the incoming RCT data in a pipeline memory until receipt of an L1A signal and perform a fast sort on the incoming data. The interface to the GT is via a mezzanine card which transmits data over 16 fiber links running at 3 Gb/s.

The CMS L1 calorimeter trigger chain does not use information from other L1 subsystems, i.e., the L1 muon trigger, which is described in the next section. L1 calorimeter and muon information is combined to a final L1 trigger decision in the GT (section 2.4).

2.3 The L1 muon trigger system

All three CMS muon detectors contribute to the L1 trigger decision. Details on how the flow of information from the DTs, CSCs, and RPCs is processed to build full muon tracks within each system, and how tracks are combined together by the GMT to provide final muon trigger candidates, are given below.

2.3.1 Muon local trigger segments

Whereas RPC trigger tracks are built by the pattern comparator trigger (PACT) using information coming from detector hits directly, local trigger track segments (primitives) are formed within DT and CSC detectors prior to the transmission to the respective track finders.

In the case of the DTs, local trigger (DTLT) track segments are reconstructed by electronics installed on the detector. Each of the 250 DTs is equipped with a mini-crate hosting readout and trigger electronics and implemented with custom ASIC [19, 20] and programmable ASIC [21] devices. Up to two DTLT per BX in the transverse plane can be generated by one chamber; DTLT information includes the radial position, the bending angle, and information about the reconstruction quality (i.e., the number of DT layers used to build a track segment). Additionally, hits along the longitudinal direction are calculated; in this case only a position is calculated as the track is assumed to be pointing to the vertex. The DTLT electronics is capable of highly efficient (94%) BX identification [1, 22], which is a challenging task given that single hits are collected with up to ≈ 400 ns drift time. A fine grained synchronization of the DTLT clock to the LHC beams is needed to ensure proper BX identification [23, 24].

The DTLT segments are received by the trigger sector collector (TSC) system, installed on the balconies surrounding the detector and implemented using flash-based FPGAs [25]. The TSC consists of 60 modules, each receiving local trigger data from one DT sector (the four or five detectors within the same muon barrel slice, called wheel, and covering 30° in azimuthal angle): trigger segments are synchronized and transmitted over 6 Gb/s optical links per sector, to the underground counting room, where optical receiver modules perform deserialization and deliver data to the DT track finder (DTTF) system.

For the CSCs, local charged-track (LCT) segments, constructed separately from the cathode (CLCT) and anode (ALCT) hits of a detector, are correlated in the trigger motherboard (TMB) when both segments exist within a detector. A CLCT provides information on the azimuthal position of a track segment, while an ALCT provides information on the radial distance of a segment from the beam line, as well as precise timing information. A maximum of two LCTs can be sent from each detector per bunch crossing. The segments from nine detectors are collected by a muon port card (MPC) residing in the same VME crate as the TMBs. The MPC accepts up to 18 LCTs and sorts them down to the best three before transmission over an optical fiber to the CSC track finder (CSCTF). There are 60 MPCs, one in each peripheral crate.

More detailed description of the DT and CSC local trigger segment reconstruction and performance in LHC collisions is given in ref. [26].

2.3.2 Drift tube track finder

The DTTF processes the DTLT information in order to reconstruct muon track candidates measured in several concentric rings of detectors, called stations, and assigns a transverse momentum value to the track candidates [27]. First, the position and bending of each DTLT is used to compute, via a LUT, that expected position at the outer stations (in case of the fourth station layer, the extrapolation is done inward towards the third one). The position of actual DTLTs is compared to the expected one and accepted if it falls within a programmable tolerance window. These windows can be tuned to achieve the desired working point, balancing the muon identification efficiency against the accepted background. To enable triggering on cosmic muon candidates, the windows can be as large as a full DT detector in order to also accept muons that are not pointing to the interaction point. All possible station pairs are linked this way and a track candidate is built. Then, the difference in azimuthal positions of the two inner segments is translated into a transverse momentum value, again using LUTs. Also the azimuthal and longitudinal coordinates of the candidate are computed, while a quality code based on the number and positions of the stations participating in the track is generated. The hardware modules are VME 9U boards hosted in 6 crates with custom backplanes and VME access; there are 72 such track finding boards, called sector processors (SP). Each SP finds up to two tracks from one DT sector. Two separate SPs analyze DTLTs from the sectors of the central wheel, to follow tracks at positive or negative pseudorapidity. Each SP receives also a subset of the DTLT information from their neighboring SPs, through parallel electrical connections, in order to perform track finding for tracks crossing detectors in different sectors. SP from external wheels also receive track segments from the CSC trigger.

The last stage of the DTTF system consists of the muon sorter (MS) [28]. First, a module called the wedge sorter (WS) collects up to 12 track candidates from the 6 SPs of one “wedge” (5 DT sectors at the same azimuthal position) through parallel backplane connections, and selects two based on the matched magnitude of the transverse momentum and on their reconstruction quality. The resulting 24 muon candidates from 12 wedge sorters are collected via parallel LVDS cables into the final sorting module, called the barrel sorter (BS), which selects the final four muon candidates to be delivered to the GMT. Both the WS and BS perform ghost cancellation algorithms before the track sorting, in order to remove duplicate tracks, e.g., multiple track candidates originating from the same muon crossing from neighboring SPs. Two WS modules are installed in each DTTF crate, while the BS is located in a separate crate called central crate. Also readout information (DTLT track segments and DTTF track candidates in a ± 1 BX window) is provided by each DTTF module and concentrated in a readout module (provided with serial link output and TTS inputs) called a data concentrator card (DCC) and located in the central crate.

2.3.3 Cathode strip chambers track finder

The CSCTF logic consists of pairwise comparisons of track segments in different detector stations that test for the compatibility in ϕ and η of a muon emanating from the collision vertex within certain tolerance windows. These comparisons are then analyzed and built into tracks consisting of two or more stations. The track finding logic has the ability to accept segments in different assigned bunch crossings by analyzing across a sliding time window of programmable length (nominally 2 BX) every bunch crossing. Duplicate tracks found on consecutive crossings are canceled. The

reported bunch crossing of a track is given by the second arriving track segment. The reported p_T of a candidate muon is calculated with large static random-access memory (SRAM) LUTs that take information such as the track type, track η , the segment ϕ differences between up to 3 stations, and the segment bend angle in the first measurement station for two-station tracks.

In addition to identifying muons from proton collisions, the CSCTF processors simultaneously identify and trigger on beam halo muons for monitoring and veto purposes by looking for trajectories approximately parallel to the beam line. A beam halo muon is created when a proton interacts with either a gas particle in the pipe or accelerator material upstream or downstream the CMS interaction point, and the produced hadrons decay. The collection of halo muons is an interesting initial data set; the muons' trajectory is highly parallel to the beam pipe and hence also to parallel to the solenoidal magnetic field; therefore, they are minimally deflected and their unbent paths are a good tool for aligning different slices of the detector disks. Additionally, these muons are a background whose rate need to be known as they have the potential to interact with multiple detector subsystems. The halo muon trigger also allows monitoring of the stability of the proton beam.

The CSCTF system is partitioned into sectors that correspond to a 60° azimuthal region of an endcap. Therefore 12 “sector processors” are required for the entire system, where each sector processor is a 9U VME card that is housed in a single crate. Three 1.6 Gbps optical links from each of five MPCs are received by each sector processor, giving a total of 180 optical links for the entire system. There is no sharing of signals across neighbor boundaries, leading to slight inefficiencies. There are several FPGAs on each processor, but the main FPGA for the track-finding algorithms is from the Xilinx Virtex-5 family. The conversion of strip and wire positions of each track segment to η , ϕ coordinates is accomplished via a set of cascaded SRAM LUTs (each $512k \times 16$ bits). The final calculation of the muon candidate p_T is also accomplished by SRAM LUTs (each $2M \times 16$ bits). In the same VME crate there is also one sorter card that receives over a custom backplane up to 3 muons from each sector processor every beam crossing and then sorts this down to the best four muons for transmission to the GMT. The crate also contains a clock and control signal distribution card, a DAQ card with a serial link interface, and a PCI-VME bridge [5, 29].

2.3.4 Resistive plate chambers trigger system

The RPCs provide a complementary, dedicated triggering detector system with excellent time resolution ($O(1ns)$), to reinforce the measurement of the correct beam-crossing time, even at the highest LHC luminosities. The RPCs are located in both the barrel and endcap regions and can provide an independent trigger over a large portion of the pseudorapidity range ($|\eta| < 1.6$). The RPCs are double-gap chambers, operated in avalanche mode to ensure reliable operation at high rates. They are arranged in six layers in the barrel and three layers in the endcaps. Details of the RPC chamber design, geometry, gas mixtures used and operating conditions can be found in refs. [1, 30]. The RPC trigger is based on the spatial and temporal coincidence of hits in different layers. It is segmented into 25 towers in η which are each subdivided into 144 segments in ϕ . The pattern comparator trigger (PACT) [31] logic compares signals from all RPC chamber layers to predefined hit patterns in order to find muon candidates. The RPCs also assign the muon p_T , charge, η , and ϕ to the matched pattern.

Unlike the CSCs and DTs, the RPC system does not form trigger primitives, but the detector hits are used directly for muon trigger candidate recognition. Analog signals from the chambers

are discriminated and digitized by front end boards (FEB), then assigned to the proper bunch crossing, zero-suppressed, and multiplexed by a system of link boards located in the vicinity of the detector. They are then sent via optical links to 84 trigger boards in 12 trigger crates located in the underground counting room. Trigger boards contain the complex PAC logic, which fits into a large FPGA. The strip pattern templates to be compared with the particle track are arranged in segments of approximately 0.1 in $|\eta|$ and 2.5° (44 mrad) in ϕ , called logical cones. Each segment can produce only one muon candidate. The trigger algorithm imposes minimum requirements on the number and pattern of hit planes, which varies with the position of the muon. As the baseline, in the barrel region ($|\eta| \leq 1.04$), a muon candidate is created by at least a 4-hit pattern, matching a valid template. To improve efficiency, this condition is relaxed and a 3-hit pattern with at least one hit found in the third or fourth station may also create a muon candidate. In addition, low- p_T muons often do not penetrate all stations. Muon candidates can also arise when three hits are found in four layers of the first and second station. In this case, only low- p_T candidates will be reconstructed. In the endcap region ($|\eta| > 1.04$) there are only 3 measurement layers available, thus any 3-hit pattern may generate a muon candidate. A muon quality value is assigned, encoded in two bits, that reflects the number of hit layers (0 to 3, corresponding to 3 to 6 planes with hits).

Hits produced by a single muon may be visible in several logical cones which overlap in space. Thus the same muon may be reconstructed, typically with different momentum and quality, in a few segments. In order to remove the duplicated candidates a special logic, called the RPC ghost buster (GB), is applied in various steps during the reconstruction of candidates. The algorithm assumes that among the muon candidates reconstructed by the PACT there is the best one, associated to the segment penetrated by a genuine muon. Since the misreconstructed muons appear as a result of hit sharing between logical cones, these muons should appear in adjacent segments. The best muon candidate should be characterized by the highest number of hits contributing to a pattern, hence highest quality. Among candidates with the same quality, the one with highest p_T is selected. The muon candidates from all the PACTs on a trigger board are collected in a GB chip. The algorithm searches for groups of adjacent candidates from the same tower. The one with the best rank, defined by quality and p_T , is selected and other candidates in the cluster are abandoned. In the second step the selected candidate is compared with candidates from the three contiguous segments in each of the neighboring towers. In the last step, the candidates are sorted based on quality criteria, and the best ranked four are forwarded to the trigger crate sorter. After further ghost rejection and sorting, the four best muons are sent to system-wide sorters, implemented in two half-sorter boards and a final-sorter board. The resulting four best muon candidates from the barrel and 4 best muon candidates from the endcap region are sent to GMT for subtrigger merging.

The RPC data record is generated on the data concentrator card that receives data from individual trigger boards.

2.3.5 Global muon trigger system

The GMT fulfills the following functions: it synchronizes incoming regional muon candidates from DTTF, CSCF, and RPC trigger systems, merges or cancels duplicate candidates, performs p_T assignment optimization for merged candidates, sorts muon candidates according to a programmable rank, assigns quality to outgoing candidates and stores the information about the incoming and outgoing candidates in the event data. The GMT is implemented as a single 9U VME module with

a front panel spanning four VME slots to accommodate connectors for 16 input cables from regional muon trigger systems. Most of the GMT logic is implemented in a form of LUTs enabling a high level of flexibility and functional adaptability without changing the FPGA firmware, e.g., to adjust selection requirements, such as transverse momentum, pseudorapidity, and quality, of the regional muon candidates [32].

The input synchronization occurs at two levels. The phase of each input with respect to the on-board clock can be adjusted in four steps corresponding to a quarter of the 25 ns clock cycle to latch correctly the incoming data. Each input can be then delayed by up to 17 full clock cycles to compensate for latency differences in regional systems such that the internal GMT logic receives in a given clock cycle regional muon candidates from the same bunch crossing.

The muon candidates from different regional triggers are then matched geometrically, according to their pseudorapidity and azimuthal angle with programmable tolerances, to account for differences in resolutions. In addition, the input η and p_T values are converted to a common scale and a sort rank is assigned to each regional muon candidate. The assignment of the sort rank is programmable and in the actual implementation it was based on a combination of input quality and estimated transverse momentum.

The matching candidates from the DT and barrel RPC and similarly from the CSC and endcap RPC triggers are then merged. Each measured parameter (η , ϕ , p_T , charge, sort rank) is merged independently according to a programmable algorithm. The η , charge, and rank were taken from the either the DT or CSC. For p_T merging, the initial setting to take the lowest p_T measurement was optimized during the data taking to become input quality dependent in certain pseudorapidity regions. In case of a match between DT and CSC, possible in the overlap region ($0.9 < |\eta| < 1.2$), one of the candidates is canceled according to a programmable logic, dependent, for example, on an additional match with RPC.

Each of the output candidates is assigned a three-bit quality value which is maximal for a merged candidate. If the candidate is not merged, its quality depends on the input quality provided by the regional trigger system and on the pseudorapidity. The quality assignment is programmable and allows for flexibility in defining looser or tighter selection of muon candidates in GT algorithms. Typically, muon candidates in double-muon triggers were allowed to have lower quality.

The final step in the GMT logic is the sorting according to the sort rank. Sorting is first done independently in the barrel and in the endcap regions and four candidates in each region with the highest rank are passed to the final sort step. Four candidates with the highest rank are then sent to the GT.

Since the GMT module and the GT system are located in the same VME crate, the two systems share a common readout. The data recorded from GMT contains a complete record of the input regional muon candidates, the four selected muon candidates from the intermediate barrel and endcap sorting steps, as well as the complete information about the four output candidates. This information is stored in five blocks corresponding to five bunch crossings centered around the trigger clock cycle.

2.4 The L1 global trigger system

The GT is the final step of the L1 Trigger system. It consists of several VME boards mounted in a VME 9U crate together with the GMT and the central trigger control system (TCS) [33, 34].

For every LHC bunch crossing, the GT decides to reject or accept a physics event for subsequent evaluation by the HLT. This decision is based on trigger objects from the L1 muon and calorimeter systems, which contain information about transverse energy E_T or transverse momentum p_T , location (pseudorapidity and azimuthal angle), and quality. Similarly, special trigger signals delivered by various subsystems are also used to either trigger or veto the trigger decision in a standalone way (“technical triggers”) or to be combined with other trigger signals into logical expressions (“external conditions”). These technical triggers (up to 64) are also used for monitoring and calibration purposes of the various CMS sub-detectors including L1 trigger system itself.

The trigger objects received from the GCT and GMT, and the input data from the other subsystems are first synchronized to each other and to the LHC orbit clock and then sent via the crate backplane to the global trigger logic (GTL) module, where the trigger algorithm calculations are performed. For the various trigger object inputs of each type (four muons, four non-isolated and four isolated e/γ objects, four central and four forward jets, four tau jets) certain conditions are applied such as E_T or p_T being above a certain threshold, pseudorapidity and/or azimuthal angle being within a selected window, or requiring the difference in pseudorapidity and/or azimuthal angle between two particles to be within a certain range. In addition, “correlation conditions” can be calculated, i.e., the difference in pseudorapidity and azimuthal angle between two objects of different kinds. Conditions can also be applied to the trigger objects formed using energy sums such as E_T^{miss} and H_T .

Several conditions are then combined by simple combinatorial logic (AND-OR-NOT) to form up to 128 algorithms. Any condition bit can be used either as a trigger or as a veto condition. The algorithm bits for each bunch crossing are combined into a “final-OR” signal by the final decision logic (FDL) module, where each algorithm can also be prescaled or blocked. An arbitrary number of sets of prescales can be defined for the algorithms in a given logic firmware version. A set of 128 concrete algorithms form an L1 menu which together with the set of prescales completely specifies the L1 trigger selection. The algorithms and the thresholds of the utilized input objects (such as transverse momentum or spatial constraints) are defined and hard-coded in firmware and are only changed by loading another firmware version. Different prescale settings allow adjustment of the trigger rate during a run by modifying the prescale values for identical copies of algorithms differing only in input thresholds. In case of a positive “final-OR” decision and if triggers are not blocked by trigger rules or detector deadtime, the TCS sends out an L1A signal to trigger the readout of the whole CMS detector and forward all data to the HLT for further scrutiny.

Trigger rules are adjustable settings to suppress trigger requests coming too soon after one or several triggers, as in this case subsystems may not be ready to accept additional triggers [35]. Sources of deadtime can be subsystems asserting “not ready” via the trigger throttling system [3], the suppression of physics triggers for calibration cycles, or the trigger rules described above.

The GT system logs all trigger rates and deadtimes in a database to allow for the correct extraction of absolute trigger cross sections from data. The trigger cross section is defined as $\sigma = R/\mathcal{L}$, where R is the trigger rate and $\mathcal{L}$ is the instantaneous luminosity.

Over the years of CMS running, the GT system has proved to be a highly flexible tool: the trigger logic implemented in the firmware of two ALTERA FPGAs (the L1 menu) was frequently updated to adapt to changing beam conditions, increasing data rates, and modified physics requirements (details in section 5). Additional subsystems (e.g., the TOTEM detector [36]) have also been configured as a part of the L1 trigger system.

2.5 Beam position timing trigger system

The two LHC beam position monitors closest to the interaction point for each LHC experiment are reserved for timing measurements and are called the Beam Pick-up Timing eXperiment (BPTX) detectors. For CMS, they are located at a distance of approximately 175 m on either side of the interaction point (BPTX+ and BPTX-).

The trigger selects valid bunch crossings using the digitized BPTX signal by requiring a coincidence of the signals from the detectors on either side (“BPTX_AND”, logical AND of BPTX+ and BPTX-).

To suppress noise in triggers with high background, a coincidence with BPTX_AND is required. Another important application has been the suppression of pre-firing from the forward hadron calorimeter caused by particles interacting in the photomultiplier anodes, rather than the detector itself. As the LHC was mostly running with a bunch spacing of 50 ns and thus there was at least one 25 ns gap without proton collisions between two occupied bunch crossings, the trigger discarded pre-firing events by vetoing the trigger for the “empty bunch crossing” before a valid bunch crossing. This is achieved by advancing the BPTX_AND signal by one bunch crossing (25 ns time unit) and using this signal to veto the L1 trigger signal (dubbed “pre-BPTX veto”). This solution also improved the physics capabilities of the L1 trigger by enabling a search for heavy stable charged particles (section 4.3.4 for details).

2.6 High-level trigger system

The event selection at the HLT is performed in a similar way to that used in the offline processing. For each event, objects such as electrons, muons, and jets are reconstructed and identification criteria are applied in order to select only those events which are of possible interest for data analysis.

The HLT hardware consists of a single processor farm composed of commodity computers, the event filter farm (EVF), which runs Scientific Linux. The event filter farm consists of filter-builder units. In the builder units, individual event fragments from the detector are assembled to form complete events. Upon request from a filter unit, the builder unit ships an assembled event to the filter unit. The filter unit in turn unpacks the raw data into detector-specific data structures and performs the event reconstruction and trigger filtering. Associated builder and filter units are located in a single multi-core machine and communicate via shared memory. In total, the EVF executed on approximately 13,000 CPU cores at the end of 2012. More information about the hardware can be found elsewhere [37].

The filtering process uses the full precision of the data from the detector, and the selection is based on offline-quality reconstruction algorithms. With the 2011 configuration of the EVF, the CPU power available allowed L1 input rates of 100 kHz to be sustained for an average HLT processing time of up to about 90 ms per event. With increased CPU power available in 2012, we were able to accommodate a per-event time budget of 175 ms per event. Before data-taking started, the HLT was commissioned extensively using cosmic ray data [38]. The HLT design specification is described in detail in [39].

The data processing of the HLT is structured around the concept of a *HLT path*, which is a set of algorithmic processing steps run in a predefined order that both reconstructs physics objects and makes selections on these objects. Each HLT path is implemented as a sequence of steps of

increasing complexity, reconstruction refinement, and physics sophistication. Selections relying on information from the calorimeters and the muon detectors reduce the rate before the CPU-expensive tracking reconstruction is performed. The reconstruction modules and selection filters of the HLT use the software framework that is also used for offline reconstruction and analyses.

Upon completion, accepted events are sent to another software process, called the storage manager, for archival storage. The event data are stored locally on disk and eventually transferred to the CMS Tier-0 computing center for offline processing and permanent storage. Events are grouped into a set of non-exclusive streams according to the HLT decisions. Most data are processed as soon as possible; however, a special “parked” data stream collected during 2012 consisted of lower-priority data that was collected and not analyzed until after the run was over [40]. This effectively increased the amount of data CMS could store on tape, albeit with a longer latency than regular, higher-priority streams. Example physics analyses enabled by the parked data stream include generic final states created via vector boson fusion, triggered by four low-momentum jets ($E_T > 75, 55, 38, 20$ GeV, for the four jets) and parton distribution function studies via Drell-Yan events at low dimuon mass, triggered by two low- p_T muons ($p_T > 17, 8$ GeV, for the two muons.)

Globally, the output rate of the HLT is limited by the size of the events and the ability of the downstream systems (CMS Tier-0) to process the events. In addition to the primary physics stream, monitoring and calibration streams are also written. Usually these streams comprise triggers that record events with reduced content, or with large prescales in order to avoid saturating the data taking bandwidth. One example is the stream set up for calibration purposes. These streams require very large data samples but typically need information only from a small portion of the detector, such that their typical event size is around 1.5 kB, while the full event size is around 0.5 MB. Among the triggers that define the calibration stream, two select events that are used for the calibration of the ECAL. The first one collects minimum bias events and only the ECAL energy deposits are recorded. By exploiting the ϕ invariance of the energy deposition in physics events, this sample allows inter-calibration of the electromagnetic calorimeter within a ϕ ring. The second ECAL calibration trigger reconstructs π^0 and η meson candidates decaying into two photons. Only the ECAL energy deposits associated with these photons are kept. Due to the small event size, CMS was able to record up to 14 kHz of π^0/η candidates in this fashion [7]. Figure 5 shows the reconstructed masses for π^0 and η candidates obtained from these calibration triggers during the 2012 run.

3 Object identification

In this section, the L1 and HLT selection of each object is discussed as well as the related main single- and double-object triggers using those objects. The event selection at the HLT is performed in a similar manner to that used in the offline event processing. For each event, objects such as electrons, muons, or jets are reconstructed and identification criteria are applied in order to select those events which are of possible interest for data analysis.

The object reconstruction is as similar as possible to the offline one, but has more rigorous timing constraints imposed by the limited number of CPUs. Section 4 describes how these objects are used in a representative set of physics triggers.

We emphasize the track reconstruction in particular as it is used in most of the trigger paths, either for lepton isolation or for particle-flow (PF) techniques [41, 42].




Figure 5. Neutral pion (left) and η (right) invariant mass peaks reconstructed in the barrel with 2012 data. The spectra are fitted with a combination of a double (single) Gaussian for the signal and a 4th (2nd) order polynomial for the background. The entire 2012 data set is used, using special online π^0/η calibration streams. The sample size is determined by the rate of this calibration stream. Signal over background (S/B) and the fitted resolution are indicated on the plots. The fitted peak positions are not exactly at the nominal π^0/η mass values mainly due to the effects of selective readout and leakage outside the 3×3 clusters used in the mass reconstruction; however, the absolute mass values are not used in the inter-calibration.

3.1 Tracking and vertex finding

Tracking and vertex finding is very important for reconstruction at the HLT. A robust and efficient tracking algorithm can help the reconstruction of particles in many ways, such as improving the momentum resolution of muons, tracking-based isolation, and b-jet tagging. Since track reconstruction is a CPU-intensive task, many strategies have been developed to balance the need for tracks with the increase in CPU time. In this section we describe the algorithm for reconstructing the primary vertex of the collision in an efficient and fast manner using only the information from the pixel detector, as well as the algorithm for reconstructing HLT tracks. More details about the tracking algorithm used in CMS, both online and offline, can be found elsewhere [43].

It is worth emphasizing that since the tracking detector data is not included in the L1 trigger, the HLT is the first place that charged particle trajectories can be included in the trigger.

3.1.1 Primary vertex reconstruction

In many triggers, knowledge of the position of the primary vertex is required. To reconstruct the primary vertex without having to run the full (and slow) tracking algorithm, we employ a special track reconstruction pass requiring only the data from the pixel detector. With these tracks, a simple gap-clustering algorithm is used for vertex reconstruction [43]. All tracks are ordered by the z coordinate of their point of closest approach to the pp interaction point. Wherever two neighboring elements in this ordered set of z coordinates has a gap exceeding a distance requirement z_{sep} , tracks on either side are split into separate vertices. In such an algorithm, interaction vertices separated by a distance less than z_{sep} are merged. Figure 6 represents the estimated number of interactions




Figure 6. Number of vertices as a function of the number of pp interactions as measured by the forward calorimeter, for fills taken in two different periods of the 2012 pp run. A linear relation can be seen between the two quantities, demonstrating good performance of the HLT pixel vertex algorithm.

versus the number of reconstructed pixel vertices for two periods, with different pileup conditions. The number of interactions is measured using the information from the HF, which covers the pseudorapidity range $3 < |\eta| < 5$. The method used is the so-called “zero counting”, which relies on the fact that the mean number of interactions per bunch crossing (μ) has a probability density described by the Poisson distribution. The average fraction of empty HF towers is measured and then μ is calculated by inverting the Poisson zero probability. Figure 6 shows that in the 2012 data, where the number of interactions per bunch crossing reached 30, the number of reconstructed vertices depends linearly on the number of pileup events for a wide range of values, demonstrating no degradation of performance due to pileup.

With increasing number of pileup collisions, we observed that the CPU time to reconstruct pixel tracks and pixel vertices increased nonlinearly. For a few HLT paths, the CPU time usage is largely dominated by the pixel track and vertex reconstruction time and it is prohibitive to use the primary-vertex finding algorithm described above.

A second method, called fast primary vertex finding, was implemented to reduce the CPU time usage. This method initially finds a coarse primary vertex and reconstructs only pixel tracks in jets associated to this vertex. The pixel tracks are then used to find the online primary vertex using the standard method described above. The coarse vertex is found as follows: initially, jets with $p_T > 40$ GeV are considered. Pixel clusters in the ϕ wedges corresponding to the jets are selected and projected to the beam axis using the jet pseudorapidity. The projections are then clustered along the z axis. If a vertex exists, the clusters will group around the z position of the vertex.

Roughly 5% of the time, the coarse vertex is not found. In these cases, the standard vertex reconstruction is run. The coarse vertex has a resolution of 0.4 cm. By using the fast primary vertex finding, the overall CPU time needed to reconstruct the vertex is reduced by a factor 4 to 6, depending on the HLT path. The reduced CPU time requirement allowed some additional paths to use b-tagging techniques than would not have been possible with the standard algorithm. The two methods have similar performance in reconstructing the online primary vertex. The efficiency of the reconstruction relative to offline is about 92% within the vertex resolution. The pixel tracks are also used in other reconstruction steps as described in the following subsections.

3.1.2 HLT tracking

Given the variety of the reconstructed objects and the fast changes in the machine conditions, it has been impossible to adopt a unique full silicon track reconstruction for all the paths. Different objects ended up using slightly different tracking configurations, which had different timing, efficiencies, and misreconstruction rates. All configurations use a combinatorial track finder (CTF) algorithm, which consists of four steps:

1. The seed generation provides initial track candidates using a few (two or three) hits and the constraint of the pp interaction point position. A seed defines the initial estimate of the trajectory, including its parameters and their uncertainties.
2. The next step is based on a global Kalman filter [44]. It extrapolates the seed trajectories along the expected flight path of a charged particle, searching for additional hits that can be assigned to the track candidate.
3. The track fitting stage uses another Kalman filter and smoother to provide the best possible estimate of the parameters of each trajectory.
4. Finally, the track selection step sets quality flags and discards tracks that fail minimum quality requirements.

Each of these steps is configurable to reduce the time at the cost of slightly degraded performance. As an example, when building track candidates from a given seed, the offline track reconstruction retains at most the five partially reconstructed candidates for extrapolation to the next layer, while at HLT only one is kept. This ensures little time increase in the presence of large occupancy events and high pileup conditions. As another example, the algorithm stops once a specified number of hits have been assigned to a track (typically eight). As a consequence, the hits in the outermost layers of the tracker tend not to be used. The different tracking configurations can be divided into four categories:

- Pixel-only tracks, i.e., tracks consisting of only three pixel hits. As stated above, the pixel-based tracking is considerably faster than the full tracking, but pixel tracks have much worse resolution and are mostly used to build the primary vertex and are also used in parts of the b- and τ -identification stages. These tracks are also used to build the seeds for the first iteration of the iterative tracking.

- Iterative tracking, i.e., a configuration which is as similar as possible to that used offline. This is used as input to the PF reconstruction.
- Lepton isolation, i.e., a regional one-step-tracking used in paths with isolated electrons and muons. On average, higher- p_T tracks are reconstructed in comparison to the iterative tracking method and as a result this variant is somewhat more time consuming than the iterative tracking.
- b tagging, i.e., a regional one-step-tracking similar to the one used for lepton isolation.

The iterative tracking approach is designed to reconstruct tracks in decreasing order of complexity. In the early iterations, easy-to-find tracks, which have high p_T and small impact parameters, are reconstructed. After each iteration hits associated with found tracks are removed, and this reduces combinatorial complexity and allows for more effective searching for lower- p_T or highly displaced tracks. For data collected in 2012, the tracking consisted of five iterations, similar (but not identical) to those run offline. The main difference between each iteration lies in the configuration of the seed generation and final track selection steps.

The first iteration is seeded with three pixel hits. Each pixel track becomes a seed. The seeds in this iteration are not required to be consistent with the primary vertex position. For the other iterations, only seeds compatible with the primary vertex z position are used. In the first iteration, we attempt to reconstruct tracks across the entire detector. For speed reasons, later iterations are seeded regionally, i.e., only seeds in a given η - ϕ region of interest are considered. These regions are defined using the η - ϕ direction of jets from tracks reconstructed in the previous iterations. Unfortunately, due to hit inefficiency in the pixel detector and the requirement of hits in each of the three pixel layers in this step, 10–15% of isolated tracks may be lost. This leads to an efficiency loss for one-prong τ lepton decays, which is recovered by adding extra regions based on the η - ϕ direction of isolated calorimeter jets. Finally, after the five iterations, all tracks are grouped together (adding the separately reconstructed muon tracks), filtered according to quality criteria and passed to the PF reconstruction.

Figure 7 shows the offline and online track reconstruction efficiency on simulated top-antitop ($t\bar{t}$) events. Online efficiencies are above 80% for track p_T above 0.9 GeV.

Figure 8 shows the time taken by the iterative track reconstruction as a function of the average pileup. As already discussed, the time spent in tracking is too high to allow the use of the tracking on each L1-accepted event. To limit the computing time, HLT tracking was only run on a subset of events that pass a set of filters, reducing it to about 30% of the total HLT CPU time.

3.2 Electron and photon triggers

The presence of high- p_T leptons and photons is a strong indicator for interesting high- Q^2 collisions and consequently much attention has been devoted to an efficient set of triggers for these processes. Electrons and photons (EG or “electromagnetic objects”) are reconstructed primarily using the lead-tungstate electromagnetic calorimeter. Each electromagnetic object deposits its energy primarily in this detector, with little energy deposited in the hadron calorimeter. The transverse shower size is of the order of one crystal. Electrons and photons are distinguished from one another by the presence of tracks pointing to electrons and lack thereof for photons. At L1, only information from




Figure 7. Tracking efficiency as a function of the momentum of the reconstructed particle, for the HLT and offline tracking, as determined from simulated $t\bar{t}$ events. Above 0.9 GeV, the online efficiency is above 80% and plateaus at around 90%.




Figure 8. The CPU time spent in the tracking reconstruction as a function of the average pileup, as measured in pp data taken during the 2012 run. The red line shows a fit to data with a second-order polynomial. On average, about 30% of the total CPU time of the HLT was devoted to tracking during this run.




Figure 9. The L1 EG resolution, reconstructed offline E_T minus L1 E_T divided by reconstructed offline E_T , in the barrel (left) and endcap (right) regions. For both distributions, a fit to a Crystal Ball function is performed. On the right curve, the red solid line shows the result after applying the transparency corrections (as discussed in section 2.2.1). For EB, the resolution after transparency correction is unchanged.

the calorimeter is available and no distinction can be made between e and γ . At the HLT level, tracks are used to resolve this ambiguity.

3.2.1 L1 electron/photon identification

L1 electron/photon trigger performance.

The L1 electron trigger resolution. Offline reconstructed electrons are matched to L1 EG candidates by looking for the RCT region which contains the highest energy trigger tower (TT) within the electron supercluster (SC) [45, 46]. In order to extract the resolution, the supercluster transverse energy reconstructed offline is compared to the corresponding L1 candidate E_T . Figure 9 shows the distribution of the L1 EG trigger resolution, offline reconstructed E_T minus L1 E_T divided by offline reconstructed E_T , in the barrel and endcap regions. The same observable is displayed as a function of the electron offline supercluster E_T and η in figure 10. Above 60 GeV, the resolution starts to degrade as the L1 saturation is reached.¹

The resolution of L1 EG candidates (figure 9) is reasonably well described by a fit to a Crystal Ball function [47]. An electron supercluster can spread its energy over a large region of the calorimeter due to the emission of photons from bremsstrahlung. The L1 EG algorithm only aggregates energy in 2 trigger towers (section 2.2.1). For this reason, the probability to trigger is reduced for electrons propagating across a significant amount of material. This effect increases with the pseudorapidity and peaks in the transition region between the EB and the EE. Figure 10 illustrates this effect by showing the L1 EG resolution as function of η . Further effects such as the transparency change of ECAL crystals with time certainly degrades the resolution further (see section 2.2.1). The resolutions shown in figures 9 and 10 were obtained after correcting for this effect.

¹The ECAL trigger primitives saturate at 127.5 GeV and RCT EG candidates at 63.5 GeV.




Figure 10. The L1 EG resolution for all electron p_T as a function of pseudorapidity η . For each η bin, a fit to a Crystal Ball function was used to model the data distribution. The vertical bars on each point represent the sigma of each fitted function which is defined as the width of the 68% area. The red points show the improved resolution after applying transparency corrections (as discussed in section 2.2.1).

L1 electron trigger efficiency. The electron trigger efficiency was measured with electrons from $Z \rightarrow ee$ events, using a tag-and-probe method [48]. The data collected in 2011 and 2012 were used. Both the tag and the probe are required to pass tight identification requirements in order to reduce significantly the background contamination. The tag electron must also trigger the event at L1, while the probe electron is used for the efficiency studies. The invariant mass of the tag-and-probe system should be consistent with the Z boson mass ($60 < M_{ee} < 120$ GeV), resulting in a pure unbiased electron data sample. The trigger efficiency is given by the fraction of probes above a given EG threshold, as a function of the probe E_T . In order to trigger, the location of the highest energy TT within the electron supercluster must match a corresponding region of an L1 candidate in the RCT.

The trigger efficiency curves are shown in figure 11 for an EG threshold of 15 GeV. The E_T on the x axis is obtained from the fully reconstructed offline energy. In the EE this includes the pre-shower energy that is not available at L1. As a consequence, the trigger efficiency turn-on point for the EE is shifted to the right with respect to the EB. For both EB and EE, corrections for crystal transparency changes were not included at L1 in 2011, which further affects the turn-on curve (section 2.2.1). The width of the turn-on curves is partly determined by the coarse trigger granularity, since only pairs of TTs are available for the formation of L1 candidates, which leads to lower energy resolution at L1. An unbinned likelihood fit was used to derive the efficiency curves. Parameters of the turn-on curves are given in table 1. Table 2 summarizes the parameters of the turn-on curves and compares them with the actual EE turn-on curve in 2011 (figure 11).

In the EE, the material in front of the detector causes more bremsstrahlung, which together with the more complex TT geometry, causes the turn-on curve to be wider than that for the EB. Some masked or faulty regions (0.2% in EB and 1.3% in EE) result in the plateaus being slightly lower than 100% (99.95% in EB and 99.84% in EE) as shown in table 1. The effect on efficiency of the L1 spike removal [49], described in section 3.3, is negligible, but will require further optimization




Figure 11. The electron trigger efficiency at L1 as a function of offline reconstructed E_T for electrons in the EB (black dots) and EE (red dots), with an EG threshold: $E_T = 15$ GeV. The curves show unbinned likelihood fits.

Table 1. The L1 electron trigger turn-on curve parameters. This table gives the electron E_T thresholds for which an efficiency of 50%, 95% and 99% are reached for EB and EE separately. The last entry corresponds to the efficiency obtained at the plateau of each curve shown in figure 11.

EG15	EB	EE
50%	$16.06^{+0.01}_{-0.01}$ GeV	$19.11^{+0.03}_{-0.06}$ GeV
95%	$22.46^{+0.04}_{-0.05}$ GeV	$27.05^{+0.01}_{-0.01}$ GeV
99%	$28.04^{+0.07}_{-0.10}$ GeV	$34.36^{+0.01}_{-0.01}$ GeV
100 GeV	$99.95^{+0.01}_{-0.88}$ %	$99.84^{+0.06}_{-0.60}$ %




Figure 12. The EE L1 electron trigger efficiency as a function of offline reconstructed E_T before (red) and after (green) transparency corrections are applied at the ECAL TP level. The curves show unbinned likelihood fits.

Table 2. The EE L1 electron trigger turn-on curve parameters. This table gives the electron E_T thresholds for which an efficiency of 50%, 95% and 99% are reached before and after transparency corrections are applied. The last entry corresponds to the efficiency obtained at the plateau of each curve shown in figure 12.

EG15	EE	EE (corr)
50%	$19.11^{+0.03}_{-0.06}$ GeV	$17.79^{+0.03}_{-0.06}$ GeV
95%	$27.05^{+0.01}_{-0.01}$ GeV	$24.46^{+0.10}_{-0.23}$ GeV
99%	$34.36^{+0.01}_{-0.01}$ GeV	$30.78^{+0.21}_{-0.48}$ GeV
100 GeV	$99.84^{+0.06}_{-0.60}$ %	$99.89^{+0.01}_{-0.67}$ %

Table 3. Turn-on points for the EG12, EG15, EG20, and EG30 L1 trigger algorithms shown in figure 13.

EG Threshold (GeV)	12	15	20	30
EB turn-on E_T (GeV)	12	16.1	20.7	29.9
EE turn-on E_T (GeV)	13	19.1	24.6	33.7

as the number of collisions per bunch crossing increases in the future. Turn-on curves for various EG thresholds are shown in figure 13, and table 3 gives their turn-on points, i.e., the E_T value where the curve attains 50% efficiency.




Figure 13. The L1 electron triggering efficiency as a function of the reconstructed offline electron E_T for barrel (left) and endcap (right). The efficiency is shown for the EG12, EG15, EG20 and EG30 L1 trigger algorithms. The curves show unbinned likelihood fits.




Figure 14. Electron trigger efficiency at L1, as a function of offline reconstructed E_T for electrons in the EB (black dots) and EE (red squares) using the 2011 data set (EG threshold: $E_T = 20$ GeV). The curves show unbinned likelihood fits.




Figure 15. Electron trigger efficiency at L1 as a function of offline reconstructed E_T for electrons in the EB (black dots) and EE (red squares) using the 2012 data set (EG threshold: $E_T = 20$ GeV). The curves show unbinned likelihood fits.

Figures 14 and 15 show the comparison of the EG20 algorithm performance obtained in 2011 and 2012. In the latter, the turn-on curve in EE is closer to that in EB. The optimizations of the ECAL trigger primitive generation (spike killing procedure and ECAL crystal transparency corrections) and RCT calibration allowed the retention of the lowest possible unprescaled trigger to be used during physics runs.




Figure 16. Rates of the isolated and nonisolated versions of the single-EG trigger versus the transverse energy threshold rescaled to an instantaneous luminosity of $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. Isolated EG rates are computed within a pseudorapidity range of $|\eta| < 2.172$ to reflect the configuration of the L1 isolated EG algorithms used in 2012.

L1 EG trigger rates. The EG trigger rates were obtained from the analysis of a dedicated data stream, containing only L1 trigger information, that was collected at high rate on the basis of L1 decision only.

For the study, events were selected using BPTX_AND trigger coincidences. This selection provides unbiased information about the L1 EG trigger response. In this fashion, it was possible to apply requirements related to the presence of L1 EG candidates with a given E_T threshold and pseudorapidity acceptance region within the analysis.

Rates of isolated and nonisolated single-EG triggers are presented in figure 16. During the 2012 run, isolated EG trigger algorithms were restricted to $|\eta| < 2.712$ at the GT level. Rates were calculated using data collected with luminosities between 4.5 and $5.5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ (for an average luminosity of $4.94 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$), and rescaled to a target instantaneous luminosity of $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. Uncertainties stemming from this small approximation are well within the fluctuations caused by data acquisition downtime variations.

3.3 Online anomalous signals and their suppression

Anomalous signals were observed in the EB shortly after collisions began in the LHC: these were identified as being due to direct ionization within the APDs, thus producing spurious isolated signals with high apparent energy. These *spikes* can induce large trigger rates at both L1 and HLT if not removed from the trigger decision. On average, one spike with $E_T > 3 \text{ GeV}$ is observed per 370 minimum bias triggers in CMS at $\sqrt{s} = 7 \text{ TeV}$. If untreated as many 60% of trigger objects containing only ECAL energy, above a threshold of 12 GeV , would be caused by spikes. At high luminosity

these would be the dominant component of the 100 kHz CMS L1 trigger rate bandwidth [50]. Spike identification and removal strategies were developed, based on specific features of these anomalous signals. In the ECAL the energy of an electromagnetic (EM) shower is distributed over several crystals, with up to 80% of the energy in a central crystal (where the electron/photon is incident) and most of the remaining energy in the four adjacent crystals. This lateral distribution can be used to discriminate spikes from EM signals. A topological variable $s = 1 - E_4/E_1$ (E_1 : E_T of the central crystal; E_4 : summed E_T of the four adjacent crystals) named “Swiss-cross” was implemented offline to serve this purpose. A similar topological variable was also developed for the on-detector electronics, a strip fine grain veto bit (sFGVB). Every TP has an associated sFGVB that is set to 1 (signifying a true EM energy deposit) if any of its 5 constituent strips has at least two crystals with E_T above a programmable trigger sFGVB threshold, of the order of a few hundred MeV. If the sFGVB is set to zero, and the trigger tower E_T is greater than a trigger killing threshold, the energy deposition is considered spike-like. The trigger tower energy is set to zero and the tower will not contribute to the triggering of CMS for the corresponding event.

As the sFGVB threshold is a single value, the electron or photon efficiency depends upon the particle energy: the higher the threshold, the more low-energy genuine EM deposits would be flagged as spikes. However, these spurious spikes may not pass the killing threshold so they would still be accepted. With a very low sFGVB threshold, spikes may not be rejected due to neighboring crystals having noise. A detailed emulation of the full L1 chain was developed in order to optimize the two thresholds to remove as large a fraction of the anomalous signals as possible whilst maintaining excellent efficiency for real electron/photon signals. In order to determine the removal efficiency, data were taken in 2010 without the killing thresholds active. Using the Swiss-cross method, spike signals were identified offline. Those signals were then matched to L1 candidates in the corresponding RCT region and the emulator used to evaluate the fraction of L1 candidates that would have been eliminated. In a similar fashion the efficiency for triggering on genuine electrons or photons could be estimated.

Three killing thresholds were emulated ($E_T = 8, 12$, and 18 GeV), combined with six sFGVB thresholds (152, 258, 289, 350, 456, 608 MeV). Figure 17 shows the electron efficiency (fraction of electrons triggered after spike removal) versus the L1 spike rejection fraction, for all sFGVB thresholds mentioned above (one point for each threshold value) and a killing threshold of 8 GeV. The optimum configuration was chosen to be an sFGVB threshold of 258 MeV and a killing threshold of 8 GeV. This corresponds to a rejection of 96% of the spikes, whilst maintaining a trigger efficiency for electrons above 98%. With these thresholds the efficiency for higher energy electrons is even larger: 99.6% for electrons with $E_T > 20$ GeV.

Table 4 summarizes the rate reduction factors obtained for L1 EG algorithms considering the working point discussed above. This optimized configuration was tested online at the beginning of 2011. It gave a rate reduction factor of about 3 (for an EG threshold of 12 GeV), and up to a factor of 10 for E_T sum triggers (which calculate the total EM energy in the whole calorimeter system).

At the end of 2011 the average pileup had peaked at 16.15, and in 2012 the highest average pileup was 34.55. Efficient identification of EM showers at trigger level became more and more challenging. As pileup events act as noise in the calorimeter, they degraded trigger object resolution and reduced the probability of observing isolated spikes. The fraction of spike-induced EG triggers was measured as a function of the number of vertices (roughly equivalent to the number of pileup




Figure 17. Electron trigger efficiency as a function of the spike rejection at L1. Each point corresponds to a different spike removal trigger sFGVB threshold. The trigger killing threshold is set to 8 GeV. The data were taken in 2010.




Figure 18. Fraction of spike-induced EG triggers as a function of the number of reconstructed vertices. The red points represent the spike removal working point used in 2011, and the green points the optimized working point for 2012. The squares (triangles) correspond to higher (lower) pileup data.

Table 4. Rate reduction factors obtained for L1 EG algorithms (considering a 258 MeV sFGVB threshold and an 8 GeV killing threshold on the ECAL Trigger Primitives) for various EG thresholds.

EG Threshold (GeV)	12	15	20	30
Rate reduction factors	3.4	4.3	6.0	9.6

events) in figure 18. The fraction of spike-induced EG triggers reaches 10% for collisions including more than 20 pileup events (red points). Using the L1 trigger emulator, a more efficient working point (sFGVB threshold = 350 MeV, killing threshold = 12 GeV) for the spike removal algorithm reduces this fraction to 6% (green points), but still preserves the same high trigger efficiency for genuine electrons and photons.

3.3.1 HLT electron and photon identification

The HLT electron and photon identifications begin with a regional reconstruction of the energy deposited in the ECAL crystals around the L1 EM candidates. This is followed by the building of the supercluster using offline reconstruction algorithms [46].

Electron and photon candidates are initially selected based on the E_T of the supercluster and on criteria based on properties of the energy deposits in the ECAL and HCAL subdetectors. Selection requirements include a cluster shape variable $\sigma_{\text{min}\eta}$ (the root-mean-square of the width in η of the shower) [46] and an isolation requirement that limits the additional energy deposits in the ECAL in a cone around the EM candidate with outer cone size of $\Delta R \equiv \sqrt{\Delta\phi^2 + \Delta\eta^2} = 0.3$, and inner cone radius corresponding to the size of three ECAL crystals ($\Delta R = 0.05$ in the barrel region.) The energy deposits in channels that are found in a strip along ϕ centered at the ECAL position of the EM candidate with an η -width of 3 crystals are also not considered. Candidates are then required to satisfy selection criteria based on the ratio of the HCAL energy in a cone of size $\Delta R = 0.3$ centered on the SC, to the SC energy.

These requirements typically reduce the trigger rate by a factor of 3–4, reaching 10 for the tightest selection used in 2012. The thresholds are such that, after this set of calorimetric criteria, the rates of electron candidates are about 1 kHz. The previously described steps are common to electron and photon selection. In addition, photon candidate selection imposes an additional isolation requirement based on tracks reconstructed in a cone around the photon candidate. In some trigger paths extra requirements are needed to keep the rate at an acceptable level. The $R_9 \equiv E_{3\times 3}/E_{SC}$ variable, where $E_{3\times 3}$ denotes the energy deposited in a small window of 3×3 crystals around the most energetic crystal in the SC, is very effective in selecting good unconverted photons even in the presence of large pileup. Finally, to distinguish electrons from photons, a nearby track is required, as described later in this section.

An improvement deployed in the e/γ triggers in 2012 was the use of corrections for radiation-induced changes in the transparency of the crystals in the endcap ECAL [7]. A new set of corrections was deployed weekly. Figure 19 shows that the introduction of these corrections in the trigger significantly improved the performance of the electron trigger in the endcap. The turn-on curve refers to a double-electron trigger requiring a 33 GeV threshold for both legs.

Double-photon trigger efficiency. The tag-and-probe method with $Z \rightarrow ee$ events is used to measure trigger efficiencies from the data. For photon triggers, the probe electron is treated as a photon and the electron SC is required to pass photon selection requirements. Events are selected from the double-electron data set with the loosest prescaled tag-and-probe trigger path. Since this path requires only one electron passing the tight HLT selection for the leading leg of the trigger, the other electron, which is only required to pass a very loose filter on its SC transverse energy, is sufficiently unbiased such that it is suitable for our measurement. We then require at least one offline electron to match the HLT electron leg, and at least two offline photons to match the HLT electron and the HLT SC leg, respectively. The two offline photons are required to have an invariant mass compatible with the Z boson (between 70 GeV and 110 GeV), and to pass offline p_T threshold of 30 GeV and 22.5 GeV, respectively. Finally the event is required to pass offline photon and event selections, e.g., for the $H \rightarrow \gamma\gamma$ measurement.




Figure 19. Efficiency of the online E_T selection as a function of the offline electron E_T , in barrel and endcap regions, before and after the deployment of online transparency corrections. The data depicts the results of a double-electron trigger requiring $p_T > 33$ GeV for both legs, and shows that applying the corrections causes a significant improvement of the online turn-on curve.




Figure 20. Efficiencies of the leading leg for the double-photon trigger as a function of the photon transverse energy (left) and pseudorapidity (right), as described in the text. The red symbols show the efficiency of the isolation plus calorimeter identification requirement, and the blue symbols show the efficiency of the R_9 selection criteria. The black symbols show the combined efficiency.

The photon matched to the HLT electron leg is also required to match to an L1 e/γ isolated object with $E_T > 22$ GeV. This photon is considered to be the tag, while the other one is the probe. Each trigger step is measured separately and, to account for the fact that electrons and photons have different R_9 distributions, each electron pair used for the trigger efficiency measurement is weighted so that the R_9 distribution of the associated SCs matches the one of a simulated photon. The net effect is an increase of the measured efficiency due to the migration of the events towards higher R_9 values.

Figures 20 to 21 show the efficiency of the leading leg selection as a function of the photon transverse energy, pseudorapidity, and number of offline reconstructed vertices (N_{vtx}).




Figure 21. Efficiencies of the leading leg of the double-photon trigger described in the text as a function of the number of offline reconstructed vertices. The red symbols show the efficiency of the isolation plus calorimeter identification requirement, and the blue symbols show the efficiency of the R_9 selection. The black symbols show the combined efficiency.

The double-photon trigger is characterized by a steep turn-on curve. The loss of efficiency shown in figure 20 (right) for the R_9 selection follows the increase of the tracker material in the region around $|\eta| \approx 1.2$, where is more likely to find converted photons with a smaller R_9 value. The flat efficiency versus N_{vtx} curve demonstrates that the path is quite insensitive to the amount of pileup events, although some small dependence is noticeable for $N_{\text{vtx}} > 30$.

Electron selection. In order to distinguish between electron and photon candidates, the presence of a reconstructed track compatible with the SC is required. Hence, after the common selection described above, the selection of online electron candidates follows with selections involving the tracker. The first step is the so called “pixel-matching”, which uses the energy and position of the SC to propagate a hypothetical trajectories through the magnetic field under each charge hypothesis to search for compatible hits in the pixel detector. Full silicon tracks are then reconstructed from the resulting pixel seeds. Timing constraints prohibit the usage of the offline tracking algorithms and a simple Kalman filter technique is used. Nevertheless, since 2012, it is complemented by the Gaussian-Sum Filtering (GSF) algorithm, which better parametrizes the highly non-Gaussian electron energy loss. Due to the large CPU time requirements of the algorithm, it was used only in paths where it is possible to achieve a large reduction of the rate before the electron tracking (e.g., in the path selecting two high- E_T electrons, where the transverse energy requirement is of 33 GeV on each electron). The electron tracks are required to have a measured momentum compatible with the SC energy. Their direction at the last tracker layer should match the SC position in η and ϕ . These selection criteria reduce the rate of misidentified electrons by a factor of 10. Finally, isolation requirements with respect to the tracks reconstructed around the electron candidate are applied, if required for rate reasons. The lowest-threshold inclusive single isolated electron path at the end of the 2012 running (corresponding to instantaneous luminosities of $7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$) had a threshold of $E_T > 27 \text{ GeV}$, with a rate of less than 50 Hz. Figure 22 shows how the rate is gradually reduced by the filtering steps of this trigger (black histogram), along with the efficiency of electrons (red points).




Figure 22. Performance of the internal stages of the lowest- E_T unprescaled single-electron trigger. The rate is shown as the black histogram (left scale); the red symbols show the efficiency for electron selection (right scale).




Figure 23. Efficiencies of the leading leg for the double-electron trigger described in the text as a function of the offline electron momentum. The trigger uses identical selection for both legs, so the other leg just has a different threshold. Efficiencies are shown for different running periods (red May, green June, blue August, and yellow November of 2012) and separately for electron reconstructed in barrel (left) and endcap (right).




Figure 24. Efficiencies of the leading leg for the double-electron trigger described in the text as a function of the number of reconstructed vertices. The trigger uses identical selection for both legs, so the other leg just has a different threshold. Efficiencies are shown for different running periods (red May, green June, blue August, and yellow November of 2012) and separately for electron reconstructed in barrel (left) and endcap (right).

Double-electron trigger efficiency. Figures 23 and 24 show the performance of the double-electron trigger. Efficiencies were measured using a tag-and-probe technique similar to that described for the photon path measurements and are computed with respect to a standard offline selection. The results are reported for various running periods; the different results reflect the different pileup conditions. Figure 24 shows that the efficiency is only loosely dependent on the pileup conditions.

3.4 Muon triggers

3.4.1 The L1 muon trigger performance

The following sections report the performance of the L1 muon trigger system described in section 2.3. Results concerning efficiency, p_T assignment resolution, rates, and timing are presented. At GT level, different GMT quality requirements are required for single- and multi-muon algorithms. Therefore, the performance for both the single- and multi-muon objects is documented.

For most of the studies offline reconstructed muons are used as a reference to measure the response of the L1 trigger. Muon identification requirements similar to the ones used by CMS offline analysis are required. These are documented in ref. [30].

The L1 muon trigger efficiency. The efficiency of the muon trigger was calculated by use of the tag-and-probe method described in [30]. Events with two reconstructed muons having an invariant mass compatible with the one of the Z boson or of the J/ψ resonance were selected out of a sample of events collected on the basis of single muon triggers.




Figure 25. The efficiency of the single-muon trigger versus the reconstructed transverse momentum of the muon for different thresholds applied on the trigger candidate p_T for the full pseudorapidity range $|\eta| < 2.4$ (left), and limited to the range $|\eta| < 2.1$ (right). The quality requirement used in the single-muon trigger algorithms (see text) was applied. Results are computed using the tag-and-probe method applied on a Z boson enriched sample.

Reconstructed tag muons were required to meet “tight” identification requirements and to be matched to SingleMu HLT objects. This allowed the removal of trigger selection biases. Reconstructed probe muons had to be identified by either the “tight” or “loose” identification criteria. The former selection matches the one used in most of the physics analyses with single muons and was used to compute the efficiency for single L1 muon triggers (figures 25 and 26), whereas the second is the muon identification baseline for many analyses with multiple muons and it was used to compute efficiencies for L1 double-muon triggers (figure 27). The L1 muon trigger efficiency was calculated on the basis of probe muons geometrically matched with L1 muon trigger candidates.

The L1 trigger candidates were matched to probes if the distance between the two was found to be smaller than $\Delta\phi = 0.15$ and $\Delta\eta = 0.2$. If two L1 trigger candidates were matched to a single probe the closest in ϕ was chosen. Tag-and-probe muons were also required to be separated by $\Delta R > 0.5$ to exclude interference of the two in the muon chambers.

The performance for different L1 p_T requirements using a sample of dimuons satisfying a mass requirement around the Z boson mass value is presented. Figure 25 shows the efficiency for single L1 muon trigger GMT quality selections as a function of the reconstructed muon p_T for $|\eta| < 2.4$ and $|\eta| < 2.1$ acceptance regions, respectively. Figure 26 shows trigger efficiency as a function of the reconstructed muon η . In this case a L1 $p_T > 16$ GeV is applied and probe muons are required to have a reconstructed p_T larger than 24 GeV.

The number of unbiased events recorded by CMS is not sufficient for a direct and precise estimation of the overall L1 double-muon trigger efficiency. In this case efficiency is obtained using the tag-and-probe method on the J/ψ resonance. Results imposing muon quality cuts as well as L1 p_T requirements from double-muon algorithms are shown in figure 27.




Figure 26. The efficiency of the single-muon trigger as a function of η for the threshold of 16 GeV (black) for muons with reconstructed $p_T > 24$ GeV. The contribution of the muon trigger subsystems to this efficiency is also presented: the red/green/blue points show the fraction of the GMT events based on the RPC/DTTF/CSCTF candidates, respectively. Results are computed using the tag-and-probe method applied to a Z boson enriched sample.




Figure 27. The efficiency of the double-muon trigger versus the reconstructed transverse momentum of the muon for different thresholds applied on the trigger candidate p_T . Results are computed using the tag-and-probe method applied to a J/ψ enriched sample.

The ability of CMS to trigger efficiently on dimuons at low p_T allowed the CMS experiment to observe the rare $B_s^0 \rightarrow \mu^+ \mu^-$ decay at 4.3σ significance [51], where a dimuon trigger with a p_T threshold of 4 GeV on each muon was applied at the HLT. The decay was established definitively at 6.2σ significance with the combination of data from both the CMS and LHCb experiments [52].

The L1 muon trigger rates. Muon trigger rate plots were obtained from the analysis of a dedicated data stream, containing L1 trigger information alone, that was collected at high rate on the basis of L1 decision only. This stream provides unbiased information about the L1 trigger response, which is ideal for L1 trigger rate studies.




Figure 28. Left: rate of the single-muon trigger versus the transverse momentum threshold for the full pseudorapidity range $|\eta| < 2.4$ and for pseudorapidity limited to $|\eta| < 2.1$. Additionally the curves for pure endcap and barrel regions are presented. Right: the rate of the single-muon trigger GMT candidates as a function of η for the p_T threshold of 16 GeV (blue histogram). The contribution of the muon trigger subsystems to this rate is also presented: the green and blue histograms show how often the above GMT candidates built using RPC or DTTF/CSCTF candidates. On both plots the rates are rescaled to an instantaneous luminosity of $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. The quality requirement used for single-muon trigger algorithms (see text) was applied.

For this analysis, events were selected on the basis of the loosest possible L1 muon trigger algorithm. The latter implies no quality or p_T requirements on the L1 muon GMT candidates, therefore any further selection (e.g., the p_T threshold or quality requirements corresponding to single- or double-muon triggers) was applied offline.

Results on the rates of single- and double-muon triggers are presented in figures 28 and 29, respectively. The single-muon trigger rate was calculated with a data recorded at instantaneous luminosities up to $7.2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ and then rescaled to an instantaneous luminosity of $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. This extrapolation was possible as the single-muon rate per instantaneous luminosity (i.e., the trigger cross section) is not a strong function of instantaneous luminosity.² The left plot of figure 28 shows a flattening of the slope of the rate curve for single-muon triggers at high L1 p_T threshold values. The effect can be explained by studying the resolution of the p_T estimation of the L1 muon trigger computed with respect to offline reconstructed “tight muons”. The results of such a comparison are presented in figure 30 and show that the muon trigger sometimes assigns very high p_T to muons with very low momentum. These candidates with overestimated transverse momentum contribute significantly to the L1 muon trigger rate, especially at high L1 p_T thresholds.

In case of the double-muon triggers, the rate increases with luminosity. The rates were calculated using data collected with the luminosities in the range $4\text{--}6 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ (for an average luminosity of $4.9 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$), and rescaled to a target instantaneous luminosity of $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. Errors from this small approximation are well within the fluctuations caused by data acquisition deadtime variations ($O(1\%)$).

²See section 3.4.2 and figure 35 specifically. The variation is at the per-mille level.




Figure 29. The rate of the double-muon trigger versus the threshold applied to the first and second muon. The rates are rescaled to the instantaneous luminosity $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.




Figure 30. The distribution of the momentum of the L1 muon candidates versus the momentum of the corresponding reconstructed muon (“tight” identification criteria). Events with both Z boson and J/ψ resonances contribute. Offline muons in the full acceptance region ($|\eta| < 2.4$) are used.




Figure 31. The overall timing distribution of L1 muon triggers. The distribution of GMT candidates is shown as a shaded histogram. The contributions from regional muon triggers (DT, CSC, RPC) are given. In addition, the GMT and RPC distributions for heavy stable charged particle trigger configurations are labeled separately.

The L1 muon trigger timing. The muon trigger timing is a product of the timing performance of muon trigger primitive generators and muon regional track-finders (DT, CSC, RPC). The GMT algorithm is executed independently for each BX. Thus no further timing corrections on candidates generated by track finders are performed at this stage. Nevertheless, the GMT algorithm, optimized for best momentum resolution and rejection of misreconstructed double-muon candidates, can discard low quality tracks, more prone to mistiming, affecting the overall L1 muon trigger timing response as well. This may result in the GMT accepting events either in the earlier or later bunch crossing (pre- or post-firing). Such errors do not currently cause incorrect L1 decisions since triggers appearing in wrong LHC bunch crossings are suppressed at the GT level by a BPTX veto.

Ideally, the trigger timing logic assigns a muon trigger candidate to the BX in which the actual muon was produced and reconstructed. In this case the difference between trigger candidate LHC BX number and LHC BX number of an event in which muon is reconstructed is 0, meaning that the candidate arrives at (relative) BX = 0. To quantify the trigger timing performance, the fraction of triggers appearing in a given BX with respect to those with ideal timing is computed. This procedure depends on an event selection used for muon reconstruction and the underlying triggers. A typical distribution of L1 muon trigger timing is shown in figure 31.

The data of figure 31 come from a stream dedicated to the express monitoring of muon reconstruction. The event selection requires the presence of a reconstructed muon with selections similar to the ones used by the “tight” identification criteria. To ensure a correspondence between L1 muon trigger candidates and reconstructed muons their position are requested to match within $\Delta R < 0.3$ of each other. No other reconstructed muons in the proximity of the one matched with the trigger are allowed. Since the most interesting candidates are the ones that may affect the GT decision, only events with p_T , $|\eta|$, and quality requirements matching the ones used for unprescaled L1 single-muon triggers in 2012 are considered.




Figure 32. The fractions of GMT candidates in early and late bunch crossings as a function of L1 muon candidate transverse momentum.

A L1 trigger is specifically implemented for heavy, stable charged particles (HSCP) (section 4.3.4), which relies on time extension of RPC signals in the RPC trigger logic. The typical response to a prompt muon thus extends to two BXs. It is therefore important that the presence of early or late signals in the RPC and DT/CSC are not correlated. Cases where both subsystem candidates respond in $BX = -1$ ($+1$), therefore not providing a GMT candidate in $BX = 0$, are rare. It is therefore typical that GMT candidates to $BX = +1$ contribute as well to $BX = 0$.

A more detailed picture, derived from the same data set and illustrating early and late GMT decisions, is given in figure 32. Here the fraction of events in $BX = +1$ and -1 is presented as a function of GMT candidate transverse momentum. The low- p_T behavior of the pre-firing curve follows the relative contribution of DT and CSC candidates. Event selection and trigger rules affect trigger timing distributions. In particular, a trigger issued in an event suppresses possible triggers in the two consecutive BXs. The above feature does not affect $BX = -1$ because triggers issued in non-colliding BXs are vetoed, but has an impact on events triggered at $BX = +1$. Therefore, in order to extract the post-firing, only events with the first GMT candidate appearing in $BX = +1$ are used. To properly normalize the plot, only events with a non-muonic additional trigger in the event were selected.

3.4.2 HLT muon identification

The muon high-level triggers combine information from both the muon and the tracker subdetectors to identify muon candidates and determine their transverse momenta, p_T . The algorithm is composed of two main steps: level-2 (L2), which uses information from the muon system only, and level-3 (L3), which combines measurements from both tracker and muon subdetectors.

Level-2. The reconstruction of a track in the muon spectrometer starts from an initial state, called the *seed*, built from patterns of DT and CSC segments. The transverse momentum of the seed is

parametrized as $p_T = f(1/\Delta\phi)$, where $\Delta\phi$ is the azimuthal angle between the two segments and f is a first-order polynomial function whose coefficients are determined using simulated CMS data. Only seeds confirmed by the L1 decision are used.

Each seed is used to start the reconstruction of a track using measurements (hits and segments) from all the muon detectors. Tracks are built with the Kalman filter technique [44], a recursive algorithm that performs pattern recognition and track fitting. After all tracks were reconstructed, possible duplicates of the same muon candidate are removed by checking that tracks do not share any hits. The interaction point position is used to constrain the track parameters to improve the transverse momentum resolution.

If one or more L2 muons are successfully reconstructed, their number and parameters are used to filter the event. The main selection is based on the L2 muon p_T . The number of muon chambers and measurements used in the track fit can also be used to suppress misreconstructed muons.

Level-3. The L3 muon reconstruction exploits the excellent momentum and vertexing resolution of the inner silicon tracker, and the larger lever arm of the muon detector, to improve the momentum resolution at high p_T (greater than ≈ 200 GeV). The L3 muon trigger algorithm consists of three main steps: seeding of tracker reconstruction starting from L2 information, track reconstruction in the tracker, and combined fit in the tracker and muon systems.

Due to HLT timing and CPU constraints, the full tracker reconstruction is not performed. Instead, tracks are seeded by L2 muon candidates. Three different seeding algorithms are available:

1. the initial state (position, momentum) for track reconstruction is the L2-track state extrapolated to the outer surface of the tracker;
2. the initial state is the L2-track state extrapolated to the outer surface of the tracker, and updated with measurements found on the outermost layers of the silicon-strip detector; and
3. the initial state is defined by pairs of hits on adjacent layers of the silicon-pixel subdetector, in small rectangular η - ϕ regions around the L2 muon track.

All these algorithms perform differently in different parts of the detector. To optimize efficiency and timing, they are run in reverse order of CPU time required: slower algorithms are only called if the faster ones fail to reconstruct a L3 muon. Starting from the initial seeds, tracks are reconstructed in the silicon tracker using a Kalman filter. These tracks and the L2 muons are propagated to a common surface (e.g., the innermost layer of the muon system) and their compatibility is evaluated using several criteria, such as their separation, directions, or relative goodness-of-fit χ^2 . If a pair of compatible L2-tracker tracks is found, a final refit of all the tracker and muon system measurements is performed.

If one or more L3 muons are successfully reconstructed, their number and parameters are used to filter the event. The main selection is based on the muon p_T . Other track parameters, such as χ^2 and impact parameter, can be used to suppress misreconstructed muons.

Isolation. The isolation of L3 muons is evaluated combining information from the silicon tracker, ECAL, and HCAL. Tracks are reconstructed in the silicon tracker in a geometrical cone of size $\Delta R = 0.3$ around the L3 muon. In the same cone, ECAL and HCAL deposits are summed. To reduce

the dependence of the isolation variable on the pileup of pp collisions, the calorimeter deposits are corrected for the average energy density in the event ρ [53]. A relative isolation variable is defined as

$$I_{\text{rel}} = \frac{1}{p_{\text{T}}^{\mu}} \left(\sum_i p_{\text{T},\text{trk}}^i + \max \left[0, \sum_j E_{\text{T},\text{ECAL}}^j + \sum_k E_{\text{T},\text{HCAL}}^k - \pi(\Delta R)^2 \rho \right] \right).$$

The standard selection is $I_{\text{rel}} < 0.15$.

Double-muon triggers. Double-muon triggers either require the presence of two L3 muons, as described above, or one L3 muon and one “tracker-muon” [30], i.e., a track in the silicon tracker compatible with one or more segments in the muon detectors. The latter class of triggers recovers possible inefficiencies of the L2 muon reconstruction (e.g., due to the muon detector acceptance). Moreover, dropping the requirement of a fitted track in the muon system allows reduction of the effective kinematic threshold, making these triggers particularly suitable for quarkonia and B physics topologies.

The two legs of double-muon triggers are generally required to originate from the same vertex to reduce the rate of misreconstructed dimuon events. In specific quarkonia triggers, additional filtering is applied to reduce the low- p_{T} background rate. This includes, for example, mass requirements on the dimuon system and requirements on the angle between the two muon candidates (section 4.5.)

Performance of muon triggers. This section describes the performance of the single- and double-muon triggers during 2012 data taking at 8 TeV. The triggers are:

- a single-muon trigger seeded by a L1 requirement of $p_{\text{T}} > 16$ GeV, and requiring a L2 track of $p_{\text{T}} > 16$ GeV and a L3 track of $p_{\text{T}} > 40$ GeV;
- a single-muon trigger seeded by an L1 trigger of $p_{\text{T}} > 16$ GeV, and requiring a L2 track of $p_{\text{T}} > 16$ GeV and a L3 track of $p_{\text{T}} > 24$ GeV; the L3 track must also be isolated;
- a double-muon trigger by a L1 trigger requiring two muon candidates of $p_{\text{T}} > 10$ and 3.5 GeV, respectively; the L2 requirement is two tracks of $p_{\text{T}} > 10$ and 3.5 GeV, and the L3 requirement is two tracks of $p_{\text{T}} > 17$ and 8 GeV; the muons are required to originate from the same vertex; by imposing a maximum distance of 0.2 cm between the points of closest approach of the two tracks to the beam line; and
- a double-muon trigger seeded by a L1 trigger requiring two muon candidates of $p_{\text{T}} > 10$ and 3.5 GeV, respectively; the L2 requires a track of $p_{\text{T}} > 10$ GeV, and the L3 a track of $p_{\text{T}} > 17$ GeV; in addition, a tracker muon of $p_{\text{T}} > 8$ GeV is required; the muons are required to come from the same vertex, by imposing a maximum distance of 0.2 cm between the points of closest approach of the two tracks to the beam line.

Trigger efficiencies are measured with the tag-and-probe method, using Z bosons decaying to muon pairs. The tag must be identified as a “tight muon” [30] and triggered by the single-isolated-muon path. The probe is selected either as a “tight muon” or a “loose muon” [30], respectively, for single- and double-muon efficiency studies. When measuring the efficiency of isolated triggers, the probe is also required to be isolated. The efficiency is obtained by fitting simultaneously the Z resonance mass for probes passing and failing the trigger in question.




Figure 33. Efficiency of single-muon triggers without isolation (top) and with isolation (bottom) in 2012 data collected at 8 TeV, as functions of η (left) and p_T , for $|\eta| < 0.9$ (right).

Figure 33 shows the efficiencies of single-muon triggers with and without isolation, as functions of η and p_T (for $|\eta| < 0.9$), in 2012 data and in simulation. The ratio between data and simulation is also shown. An agreement of the level of 1–2% is observed.

Figure 34 shows the efficiencies for the double-muon triggers with and without the tracker muon requirement for tight muons of $p_T > 20$ GeV, as functions of η of the two muons. The total efficiency includes contributions from the efficiency of each muon leg and from the dimuon vertex constraint.

Figure 35 shows the trigger cross sections of the four main muon triggers in 2012 data taking, as functions of the LHC instantaneous luminosity. As is shown in the figure, during the 2012 run, a mild pileup-dependent inefficiency was observed for paths using L3 reconstruction. This effect caused a drop in the cross section of the isolated muon trigger at high luminosity. Figure 35 shows that this effect is not visible in nonisolated triggers (such as the single-muon path with a




Figure 34. Efficiencies of double-muon triggers without (left) and with (right) the tracker muon requirement in 2012 data collected at 8 TeV as functions of the pseudorapidity $|\eta|$ of the two muons, for loose muons with $p_T > 20$ GeV.




Figure 35. Cross sections of the four main single- and double-muon triggers used in 2012 data taking, described in the text, as a function of the LHC instantaneous luminosity. Mild pileup dependencies are visible.

$p_T > 40$ GeV requirement) as in those cases it is masked by a slight luminosity-dependent cross section increase.

3.5 Jets and global energy sums

Triggers based on jet and missing transverse energy (E_T^{miss}) triggers play an important role for search for new physics. Single-jet triggers are primarily designed to study quantum chromodynamics (QCD), but can also be used for many analyses, such as searches for new physics using initial state radiation (ISR) jets. The dijet triggers are designed primarily for jet energy scale studies. The




Figure 36. Illustration of the available tower granularity for the L1 jet finding algorithm in the central region, $|\eta| < 3$ (left). The jet trigger uses a 3×3 calorimeter region sliding window technique which spans the full (η, ϕ) coverage of the calorimeter. The active tower patterns allowed for L1 τ jet candidates are shown on the right.

E_T^{miss} triggers are designed to search for new physics with invisible particles, such as neutralinos in supersymmetric models.

3.5.1 The L1 jet trigger

The L1 jet trigger uses transverse energy sums computed using both HCAL and ECAL in the central region ($|\eta| < 3.0$) or HF in the forward region ($|\eta| > 3.0$). Each central region is composed of a 4×4 matrix of trigger towers (figure 36), each spanning a region of $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ up to $|\eta| \approx 2.0$; for higher rapidities the $\Delta\phi$ granularity is preserved, while the $\Delta\eta$ granularity becomes more coarse. In the forward region, each region consists of 4 or 6 HF trigger towers and has the same $\Delta\phi$ granularity of 0.384 as in the central region, with the $\Delta\eta$ granularity of 0.5. The jet trigger uses a “sliding window” technique [5] based on a 3×3 regions (i.e., 144 trigger towers in the central region and up to 54 trigger towers in the forward region), spanning the full (η, ϕ) coverage of the CMS calorimeter. The L1 jet candidate is found if the energy deposits in the 3×3 window meet the following conditions: the central region of the 3×3 matrix must have the E_T higher than any of the eight neighbors, and this E_T must exceed a specific threshold (used to suppress the calorimeter noise). The L1 jets are characterized by the transverse energy E_T equal to the sum of transverse energies in the 3×3 regions of the sliding window centered on the jet. The L1 jet is labeled by the (η, ϕ) of its central region.

Jets with $|\eta| > 3.0$ are classified as forward jets, whereas those with $|\eta| < 3.0$ are classified as central or τ jets, depending on the OR of the nine τ veto bits associated with the 9 regions in the 3×3 window. To improve the detection efficiency for genuine L1 τ jets, a geometrical tower pattern is utilized for L1 τ jet candidates (figure 36).

The four highest energy central, forward, or central τ jets in the calorimeter are selected. After jets are found, LUTs are used to apply a programmable η -dependent jet energy scale correction.

The performance of the L1 jets is evaluated with respect to offline jets, which are formed from the standard CaloJet reconstruction, as well as PF jet reconstruction. Jets are reconstructed using the




Figure 37. Left: the L1 jet trigger efficiency as a function of the offline CaloJet transverse momentum. Right: the L1 jet trigger efficiencies as a function of the PF jet transverse momentum. In both cases, three L1 thresholds ($E_T > 16, 36, 92$ GeV) are shown.

anti- k_T algorithm and calibrated for the nonlinearity of the calorimeter response and pileup effects using a combination of studies based on simulation and collision data, as detailed in ref. [54]. A moderate level of noise rejection is applied to the offline jets by selecting jets passing “loose” [54] identification criteria.

L1 jet trigger efficiency. The L1 jet trigger efficiency was measured with a data sample from the single-muon data set requiring an isolated muon with $p_T > 24$ GeV (HLT_IsoMu24). Events from the muon paths are unbiased with respect to the jet trigger paths.

The L1 jet efficiency is calculated relative to the offline reconstructed jets. The efficiency is defined as the fraction of leading offline jets that were matched to an L1 central, forward, or central, τ jet above a certain trigger threshold, divided by the number of offline (leading) jets that were matched to an L1 central, forward, or central τ jet above any threshold. This quantity is then plotted as a function of the offline jet p_T , η , and ϕ . The efficiency is determined by matching the L1 and reconstructed offline jets spatially in η - ϕ space. This is done by calculating the minimum separation, ΔR , between the highest- E_T reconstructed jet (with $p_T > 10$ GeV and $|\eta| < 3$) and any L1 jet above a certain E_T threshold, and requiring it to be less than 0.5. Should there be more than one jet satisfying this selection, the one closest (in ΔR) is taken as the matched jet.

We evaluated the efficiency turn-on curves for various L1 jet thresholds ($E_T > 16, 36$ and 92 GeV) as a function of the offline jet p_T . The efficiency is calculated with respect to offline PF and CaloJet transverse energies (figure 37). Each curve is fitted with a function that is the cumulative distribution function of an exponentially modified Gaussian (EMG) distribution. In this functional form, a parameter, μ , determines the point of 50% efficiency and σ represents the resolution.

Pileup dependence. To evaluate the effect on the performance of the L1 triggers in different pileup scenarios, the L1 jet efficiency is also benchmarked as a function of pileup. The measure of the pileup per event is defined by the number of ‘good’ reconstructed primary vertices in the event, with each vertex satisfying the following requirements




Figure 38. The L1 jet efficiency turn-on curves as a function of the leading offline CaloJet E_T (left) and as a function of the leading offline PF jet E_T (right), for low-, medium-, and high-pileup scenarios for three different thresholds: $E_T > 16, 36$, and 92 GeV.

- $N_{\text{dof}} > 4$;
- vertex position along the beam direction of $|z_{\text{vtx}}| < 24$ cm;
- vertex position perpendicular to the beam of $\rho < 2$ cm.

Three different pileup bins of 0–10, 10–20, and >20 vertices are defined, reflecting the low-, medium-, and high-pileup running conditions in 2012 for CaloJets and PF jets, respectively. The corresponding turn-on curves are shown in figure 38.

There is no significant change of the jet trigger efficiency observed in the presence of a high number of primary vertices. The increase in hadronic activity in high-pileup events, combined with the absence of pileup subtraction within L1 jets, results in the expected observation of a decrease in the μ value of the jet turn-on curves as a function of pileup, while the widths (σ) of the turn-on curves are found to gradually increase with increasing pileup.

3.5.2 The L1 energy sums

The GCT calculates the total scalar sum of E_T over the calorimeter regions, as well as E_T^{miss} based on individual regions. In addition, it calculates the total scalar sum of L1 jet transverse energies (H_T) and the corresponding missing transverse energy H_T^{miss} based on L1 jet candidates.

Energy sum trigger efficiencies. The performance of the various L1 energy sum trigger quantities is evaluated by comparison with the corresponding offline quantities. The latter are defined at the analysis level according to the most common physics analysis usage. The following offline quantities are defined:

- Missing transverse energy, E_T^{miss} , which is the standard (uncorrected) calorimeter-based E_T^{miss} .
- Total transverse jet energy, H_T (see section 1).




Figure 39. The L1 H_T efficiency turn-on curves as a function of the offline CaloJet (left) and PF (right) H_T , for three thresholds ($H_T > 75, 100, 150$ GeV).




Figure 40. The L1 E_T^{miss} efficiency turn-on curve as a function of the offline calorimeter E_T^{miss} , for three thresholds ($E_T^{\text{miss}} > 30, 40, 50$ GeV).

Figure 39 show the L1 H_T efficiency turn-on curve for three L1 H_T thresholds of 75, 100, and 150 GeV as a function of offline CaloJet H_T (left), and PF H_T (right). Figure 40 shows the L1 E_T^{miss} efficiency curve for three L1 E_T^{miss} thresholds of 30, 40, and 50 GeV. The turn-on points in all the efficiency curves are shown to be shifted towards larger values than the corresponding L1 trigger thresholds, which is explained by the fact that the quantities are defined in different way at the trigger and offline levels; the trigger uses standard calorimeter reconstruction based object definition, whereas offline uses the PF object definition. The same reasoning explains the slow turn-on curves observed in the performance of the energy sum triggers versus the PF quantities, with the resolution appearing to worsen when compared to the performance obtained using the standard calorimeter reconstruction. In both cases, the L1 H_T and L1 E_T^{miss} efficiencies plateau at 100%.




Figure 41. The rate of the L1 single-jet trigger as a function of the E_T threshold. The rates are rescaled to the instantaneous luminosity $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.




Figure 42. Left: rate of the L1_HTT trigger versus the L1_HTT threshold. Right: rate of the L1_ETM missing transverse energy trigger as a function of the L1_ETM threshold. On both plots, the rates are rescaled to the instantaneous luminosity $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.

3.5.3 L1 jet and energy sum rates

The L1 single jet trigger rates as a function of the L1 jet threshold were also evaluated, using similar strategy to that described in the muon identification section. We used data recorded in a special data set in which only the essential needed information about the events was stored, and further selected events without any bias based on the trigger selection (i.e., zero bias triggered events) and correspond to an instantaneous luminosity of $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. Figure 41 shows the L1 single-jet trigger rate as a function of the L1 jet threshold. Similarly, the rates of the L1 energy sum triggers (L1_HTT and L1_ETM triggers here) are shown in figure 42.

3.5.4 The HLT jet triggers

At the HLT, jets are reconstructed using the anti- k_T clustering algorithm with cone size $R = 0.5$ [53, 55]. The inputs for the jet algorithm are either calorimeter towers (resulting in so-called “CaloJet” objects), or the reconstructed particle flow objects (resulting in “PFJet” objects). In 2012, most of the jet trigger paths use PFJet as their inputs. As the PF algorithm uses significant CPU

Table 5. Single-jet triggers used for $\mathcal{L} = 7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ (pileup ≈ 32), their prescales, and trigger rates at that instantaneous luminosity.

Path name	L1 seed	L1 prescale	HLT prescale	Approx. Rate (Hz)
HLT_L1SingleJet16	L1_SingleJet16	200,000	55	0.9
HLT_L1SingleJet36	L1_SingleJet36	6,000	200	1.8
HLT_PFJet40	L1_SingleJet16	200,000	5	0.2
HLT_PFJet80	L1_SingleJet36	6,000	2	1.0
HLT_PFJet140	L1_SingleJet68	300	2	1.5
HLT_PFJet200	L1_SingleJet92	60	2	1.2
HLT_PFJet260	L1_SingleJet128	1	30	1.3
HLT_PFJet320	L1_SingleJet128	1	1	12.7
HLT_PFJet400	L1_SingleJet128	1	1	3.7
HLT_Jet370_NoJetID	L1_SingleJet128	1	1	6.7

resources, PFJet trigger paths have a pre-selection based on CaloJets. Matching between CaloJets and PFJets is then required in single PFJet paths.

Single-jet paths. The L1 thresholds for the single-jet paths were chosen such that the L1 efficiency is at least 95% at the corresponding HLT threshold. The jet energy scale corrections (JEC) were applied to the single-jet paths. The lowest threshold path was a L1 pass-through path that simply requires a L1 jet in the event with $p_T > 16 \text{ GeV}$. The single PFJet trigger paths for $\mathcal{L} = 7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ (pileup ≈ 32), along with the L1, prescales, and approximate rates are listed in table 5. The trigger turn-on curves for selected single PFJet paths as a function of transverse momentum of the offline jet is shown in figure 43. The trigger efficiency was calculated from an independent data sample collected using a single isolated muon trigger with a $p_T > 24 \text{ GeV}$ threshold. As in the L1 case (section 3.5.1), the efficiency is evaluated in comparison to offline jets, in this case, PF jets.

Dijet paths. The dijet trigger is primarily used to collect data for η -dependent energy corrections using a p_T -balance technique [54]. This correction removes any variation in the calorimeter response to a fixed jet p_T as a function of jet η .

The dijet triggers require two HLT jets with an average transverse energy greater than a given threshold. The lowest threshold path requires two HLT jets with an average transverse energy greater than 40 GeV. The DiPFJet trigger paths for $\mathcal{L} = 7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ (pileup ≈ 32), along with the L1 and HLT prescales and rates are listed in table 6. The lowest transverse energy unscaled path has a threshold of 400 GeV.

3.5.5 The HLT E_T^{miss} triggers

In this section, triggers that exclusively place requirements on missing transverse energy are described. Unscaled E_T^{miss} triggers are of particular interest for searches for new physics processes beyond the standard model. Hypothetical particles, such as the lightest supersymmetric particle




Figure 43. Left: efficiency of the L1 single-jet trigger with an E_T threshold of 128 GeV as a function of the offline jet transverse momentum. Right: the HLT efficiencies as a function of transverse momentum for a calorimeter jet trigger with a 370 GeV threshold and no jet identification requirements [56], and two PF jet triggers with 320 and 400 GeV thresholds.

Table 6. Dijet-triggers used at $\mathcal{L} = 7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ (pileup ≈ 32), their prescales, and trigger rates. The main purpose of these triggers is the η -dependent calibration of the calorimeter.

Path name	L1 seed	L1 prescale	HLT prescale	Rate (Hz)
HLT_DiPFJetAve40	L1_SingleJet16	200,000	1	0.51
HLT_DiPFJetAve80	L1_SingleJet36	6,000	1	0.71
HLT_DiPFJetAve140	L1_SingleJet68	300	1	1.51
HLT_DiPFJetAve200	L1_SingleJet92	60	1	1.36
HLT_DiPFJetAve260	L1_SingleJet128	1	15	1.41
HLT_DiPFJetAve320	L1_SingleJet128	1	5	1.19
HLT_DiPFJetAve400	L1_SingleJet128	1	1	1.44

(LSP), graviton, or dark matter, will interact weakly in the CMS detector before escaping. Their presence can be inferred by a measured imbalance in the energy or momentum of the observed particles in the event.

The E_T^{miss} algorithms. The E_T^{miss} at the HLT is calculated using the same algorithms as the offline analysis. Two algorithms were used to reconstruct the E_T^{miss} in the HLT. The first algorithm, called CaloMET, calculated the E_T^{miss} by summing all towers in the calorimeter,

$$E_T^{\text{miss}} = \sqrt{\left(\sum_{\text{towers}} E_x\right)^2 + \left(\sum_{\text{towers}} E_y\right)^2}. \quad (3.1)$$

Table 7. The E_T^{miss} triggers used for $\mathcal{L} = 7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ (pileup ≈ 32), their prescales, and rates at that luminosity. Note that the L1 $E_T^{\text{miss}} > 36 \text{ GeV}$ trigger (L1_ETM36) was highly prescaled starting at this luminosity and hence the need to use an OR with the L1 $E_T^{\text{miss}} > 40 \text{ GeV}$ trigger (L1_ETM40). The parked HLT $E_T^{\text{miss}} > 80 \text{ GeV}$ trigger (HLT_MET80_Parked) was also anticipated to be highly prescaled starting from $\mathcal{L} = 8 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. The E_T^{miss} parking triggers were available at the end of 2012. “Cleaned” refers to application of dedicated algorithms to remove noise events.

Path name	L1 seed	HLT prescale	Rate (Hz)
Prompt triggers			
HLT_MET80	L1_ETM36 OR L1_ETM40	100	0.48
HLT_MET120	L1_ETM36 OR L1_ETM40	8	0.71
HLT_MET120_HBHENoiseCleaned	L1_ETM36 OR L1_ETM40	1	3.92
HLT_MET200	L1_ETM70	1	1.46
HLT_MET200_HBHENoiseCleaned	L1_ETM70	1	0.63
HLT_MET300	L1_ETM100	1	0.47
HLT_MET300_HBHENoiseCleaned	L1_ETM100	1	0.15
HLT_MET400	L1_ETM100	1	0.19
HLT_MET400_HBHENoiseCleaned	L1_ETM100	1	0.05
HLT_PFMET150	L1_ETM36 OR L1_ETM40	1	3.05
HLT_PFMET180	L1_ETM36 OR L1_ETM40	1	1.92
Parked triggers			
HLT_MET80_Parked	L1_ETM36 OR L1_ETM40	1	47.54
HLT_MET100_HBHENoiseCleaned	L1_ETM36 OR L1_ETM40	1	9.09

Another algorithm (PFMET) uses the negative of the vector sum over transverse momenta of reconstructed anti- k_T PF jets,

$$\text{PF } E_T^{\text{miss}} = \sqrt{\left(\sum_{\text{PFJet}} P_x\right)^2 + \left(\sum_{\text{PFJet}} P_y\right)^2}. \quad (3.2)$$

No minimum threshold requirement on jet p_T was applied in this algorithm at the HLT. As with the PFJet trigger paths, a pre-selection based on the CaloMET is applied before the PFMET is calculated to reduce the required CPU time of the PF algorithm. Table 7 shows the E_T^{miss} triggers used for $\mathcal{L} = 8 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ in 2012, together with prescale factors at L1 and HLT, and rate estimated using a 2012 dedicated data sample.

Efficiency of E_T^{miss} triggers. The trigger turn-on curves as a function of E_T^{miss} are shown in figures 40 and 44. The trigger efficiency is calculated from an independent data sample collected using the lowest- p_T unscaled isolated single muon trigger path, with $p_T > 24 \text{ GeV}$.




Figure 44. The HLT efficiencies as a function of the offline $PF E_T^{\text{miss}}$ for different E_T^{miss} thresholds ($E_T^{\text{miss}} = 80\text{--}400$ GeV).

3.6 τ lepton triggers

The τ -jet triggers are important for a wide variety of physics analyses that use τ leptons decaying hadronically. In many models of new physics, third-generation particles play a special role in elucidating the mechanism for spontaneous symmetry breaking and naturalness. The τ leptons, as the charged leptons of the third generation, constitute important signatures for $h \rightarrow \tau\tau$ searches and certain new physics scenarios. The tau triggers are designed to collect events with τ leptons decaying hadronically. Hadronic decays make up more than 60% of the tau branching fractions, mostly via final states with either one or three charged hadrons in a tightly collimated jet with little additional activity around the central cone. Leptonic tau decays are automatically collected by electron and muon triggers. In what follows, we refer to taus that decay hadronically as τ_h and τ leptons that decay to electrons (muons) as τ_e (τ_μ).

3.6.1 The L1 τ lepton identification

A common approach to separate τ leptons decaying to hadrons (τ_h) from quark and gluon jets is by using isolation criteria. This is a challenging task to perform at the L1 trigger because of the given coarse granularity of the L1 calorimeter readout (figure 36). The L1 τ objects are mandatory, however, for analyses such as $h \rightarrow \tau\tau$, with both τ leptons decaying hadronically.

The L1 τ_h identification starts from previously identified L1 jet objects (section 3.5.1), which are further investigated using an isolation variable and a τ veto bit. We require that seven out of the eight noncentral trigger regions contain small energy deposits ($E_T < 2$ GeV). This acts as an isolation requirement. In addition, for each trigger region a τ veto bit is set if the energy deposit




Figure 45. Examples of trigger regions, where trigger towers with energy deposits $E_T^{\text{ECAL}} > 4 \text{ GeV}$ or $E_T^{\text{HCAL}} > 4 \text{ GeV}$, are shown as shaded squares. The L1 τ veto bit is not set if the energy is contained in a square of 2×2 trigger towers (a). Otherwise, the τ veto bit is set (b).

is spread over more than 2×2 trigger towers (figure 45). The L1 τ objects are required to have no τ veto bit set in all nine trigger regions, further constraining the energy spread within the two most energetic trigger regions. If either the isolation or the τ veto bit requirement fails, the object is regarded as an L1 central jet.

The $h \rightarrow \tau_h \tau_h$ search [57] uses an L1 seed requiring two L1 τ objects with $p_T > 44 \text{ GeV}$ and $|\eta| < 2.17$. For large τ energies, the isolation criteria introduce an inefficiency for genuine τ leptons. This is recovered by also allowing events with two L1 jets (central or τ) with $p_T > 64 \text{ GeV}$ and $|\eta| < 3.0$ to be selected. Figure 46 shows the rate of these L1 seeds as a function of the applied p_T threshold on the two objects. The measured efficiency of this L1 seed reaches a plateau of 100% at $p_T \approx 70 \text{ GeV}$, as shown in figure 47. The efficiency as function of the pseudorapidity is obtained using τ leptons with $p_T > 45 \text{ GeV}$. This requirement emulates the p_T requirement used in the $h \rightarrow \tau_h \tau_h$ search.

3.6.2 The HLT τ lepton identification

The τ -jet triggers identify and select events with hadronic decays of the τ leptons; leptonic decays are selected as prompt electrons and muons. There are three levels of the τ HLT; each is designed to reduce the rate before running the more complex subsequent step. The first step we call the level-2 (L2) τ trigger; it is built with CaloJets. The second step is referred to as level-2.5 (L2.5); this step requires isolation for matching tracks reconstructed from the pixel detector. The last step, called level-3 (L3), uses the PF algorithm to build τ lepton candidates using information from all major subdetectors. Offline τ reconstruction with CMS is described in more detail elsewhere [58]. The HLT τ paths come in two distinct varieties. The first is for τ_h candidates triggered with the L1 trigger. These τ lepton triggers have a L2 and L2.5 step to reduce the rate before running the more advanced L3 τ reconstruction. The second type of τ trigger path is triggered at L1 by a lepton or other event quantity such as E_T^{miss} . These triggers have HLT electron, muon or missing energy selections to reduce the rate before running the L3 τ algorithm.

The L2 τ -jet trigger reconstruction is entirely based on calorimeter information. The CaloJets are built with a cone of radius equal to 0.2 seeded by L1 τ jets (section 3.6.1) or L1 central jets. The only selection applied is a p_T threshold on the jet transverse energy.




Figure 46. Rate of L1 double- τ and double-jet seeds as a function of the p_T threshold on the two objects. The double- τ objects are restricted to $|\eta| < 2.17$, while the double-jet requires two seed objects (either τ or jet) within $|\eta| < 3.0$. The given rates are scaled to an instantaneous luminosity of $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$.




Figure 47. Efficiency of the double- τ_h L1 trigger with a threshold of 44 and 64 GeV on the L1 τ and jet objects, respectively. Presented is the efficiency of one τ lepton candidate as a function of transverse momentum (left) and pseudorapidity (right).




Figure 48. Efficiency of the L2 and L2.5 τ trigger with a 35 GeV threshold as a function of the offline reconstructed τ transverse momentum (left) and pseudorapidity (right).

The L2.5 step consists of a track-based isolation applied on the L2 τ candidates that are above the p_T threshold. The isolation starts by reconstructing the pixel tracks and selecting those coming from the primary vertex and matched to the L2 τ candidate. A L2 τ is considered to be isolated if there is no pixel track from the same vertex with transverse momentum greater than 1.2 GeV in an isolation annulus between $0.2 < \Delta R < 0.4$ around the τ candidate.

Finally, the L3 τ reconstruction uses the PF algorithm. The online reconstruction uses a so-called *fixed cone* τ algorithm with a signal cone of $\Delta R = 0.18$, which contains the τ decay products, and an isolation annulus of $0.18 < \Delta R < 0.45$. The trigger uses tracker-only isolation built using tracks from a vertex compatible with the primary vertex of the τ to minimize pileup dependence. There are two isolation working points: loose and tight. A loose τ is considered isolated if no tracks with $p_T > 1.5$ GeV are found in the isolation annulus. A τ candidate is considered to be “tight” if it has no tracks with $p_T > 1.0$ GeV with a signal/isolation cone boundary at 0.15.

Trigger efficiencies are measured individually in each step. For the double- τ trigger a per-leg efficiency is measured. A sample of $Z \rightarrow \tau\tau$ events selected by a single-muon trigger is used for the measurement, with one τ decaying hadronically and the other to muon and neutrinos. The $\tau_h\tau_\mu$ candidates are selected and discriminated against multijet and W boson backgrounds using muon isolation, charge requirements, and low transverse mass M_T to achieve a τ_h purity of approximately 50%.

The efficiency for the L2/L2.5 stages of the τ trigger with a transverse momentum threshold of 35 GeV is shown in figure 48. The efficiency reaches a plateau of 93.2% at 55 GeV.

For the L3 efficiency measurement, a slightly different event selection is applied: $Z \rightarrow \tau\tau_\ell$ events (with $\ell = e$ or μ) are selected with a muon-plus- E_T^{miss} or a single-electron trigger. Tight isolation on the electron/muon and $M_T < 20$ GeV, measured between the electron/muon and the missing energy, are also required. The purities after this selection are 78% and 65% for $|\eta_{\tau_h}| < 1.5$ and $1.5 < |\eta_{\tau_h}| < 2.3$, respectively. The event samples used to calculate the efficiencies in the simulation are mixed with simulated W+jets events to produce a compatible purity. The efficiency for the L3 τ trigger with a 20 GeV threshold is shown in figure 49. The efficiency reaches a plateau




Figure 49. Efficiency of the loose L3 τ algorithm from the $\tau_h\tau_\mu$ events plotted as a function of offline τ_h transverse momentum (left), pseudorapidity (center), and number of vertices (right).

of 90% very quickly at about 22 GeV. The $\tau_h\tau_h$ triggers use the tight working point. This event topology is dominated by multijet background. The tighter working point substantially reduces the rate and provides an efficiency of 80% on the plateau. In offline analyses the efficiency of the simulation is corrected as a function of the transverse momentum to match the efficiency measured in data events.

In summary, the τ HLT is used in a variety of very important physics analyses, including standard model Higgs boson searches. These analyses combine the τ_h trigger algorithms described above with other HLT objects, such as electrons, muons, and missing transverse energy. These analyses have efficiencies as high as 90% while maintaining a manageable HLT rate.

3.7 b-quark jet tagging

Many important processes studied at the LHC contain jets originating from b quarks. The precise identification of b jets is crucial to reduce the large backgrounds. In CMS, this background can be suppressed at the HLT by using b tagging algorithms, giving an acceptable trigger rate with large signal efficiency.

The b tagging algorithms exploit the fact that B hadrons typically have longer decay lifetimes than the hadrons made of light or charm quarks. As a consequence, their decay product tracks and vertices are significantly displaced from the primary vertex. Similarly, B hadrons decay more frequently to final states with leptons than their light-flavor counterparts.

The track counting (TC) and combined secondary vertex (CSV) algorithms used for offline b tagging [59] are adapted to be used at the HLT to trigger events containing jets originating from b quarks. The TC algorithm uses the impact parameter significance of the tracks in the jets to discriminate between jets originating from b quark jets from other flavors. Combined information on impact parameter significance of the tracks and properties of reconstructed secondary vertices in the jets are combined in a multivariate discriminant in the CSV algorithm.

The choice of which b tagging is used in a particular HLT path depends on timing requirements. A compromise has to be found to keep the CPU usage and trigger rates at low levels while keeping the trigger efficiency as high as possible. Therefore, online b tagging techniques were designed to be very flexible, allowing the use of not only different algorithms, but also input objects, namely




Figure 50. The efficiency to tag b quark jets versus the mistag rate, obtained from Monte Carlo simulations, for the track counting (TC) and for the combined secondary vertex (CSV) algorithms. As expected from offline studies, the CSV algorithm performs better than the TC algorithm.

primary vertex and tracks, reconstructed with different methods. The b tagging algorithms depend on the primary vertices found via the fast primary vertex algorithm described in section 3.1.1.

3.7.1 Tracking for b tagging

Three tracking reconstruction methods are available at the HLT (section 3.1) and are used for b tagging: pixel, regional, and iterative tracking.

The reconstruction of pixel tracks is very fast; however, the performance is limited. Thus, the pixel tracks are essentially only used in online b tagging using TC algorithms with jets reconstructed from energy deposits in the calorimeter at an intermediate step (L2.5) of the trigger paths. At L2.5 the b tagging discriminant thresholds are typically loose with the exclusive aim to reduce the input rates to the slower, but better performing, reconstruction of regional tracks. The regional tracks are used as input to b tagging at a later step, called L3, of event triggering. Paths using online PF jets have tracks reconstructed with the high-performance iterative tracking, which can be used by both online algorithms.

3.7.2 Performance of online b-tagging

The performance of the online b tagging at the HLT is illustrated in figures 50 and 51. Figure 50 shows the efficiency to tag b quark jets versus the mistag rate, obtained from Monte Carlo simulations, for both algorithms. As expected from studies of the performance of the algorithms used offline, the CSV algorithm performs better than the TC algorithm.




Figure 51. The efficiency of the online CSV trigger as a function of the offline CSV tagger discriminant, obtained from the data and from Monte Carlo simulations. Good agreement between the two is observed.

The efficiency of the online CSV trigger as a function of the offline CSV tagger discriminant, obtained from the data, is shown in figure 51 for a trigger path with selections on central PF jets with $E_T > 30$ GeV and $E_T^{\text{miss}} > 80$ GeV relative to an identical (prescaled) trigger path without the b tagging part. The data are a $t\bar{t}$ -enriched control region (requiring at least three jets and at least one isolated lepton). This defines the denominator of the efficiency ratio. The numerator additionally applies a requirement such that $\epsilon_{\text{CSV}} > 70\%$ for the b tagging discriminant. For the simulation studies, a sample of $t\bar{t}$ events is used with the same selection. The choice to use $-\ln(1 - \text{CSV})$ in the x -axis is because on the fact that the distribution of the CSV discriminant is limited to the range between zero and unity, and peaks at unity. This choice makes it possible to visualize the turn-on behavior. A typical requirement of $\text{CSV} > 0.9$ corresponds to 2.3 on the x axis.

3.8 Heavy ion triggers

The running conditions for PbPb collisions are significantly different from the pp case. The instantaneous luminosity delivered by the LHC in the 2010 (2011) PbPb running periods was 3×10^{25} (5×10^{26}) $\text{cm}^{-2} \text{s}^{-1}$ resulting in maximum interaction rates of 250 Hz (4 kHz), much lower than in pp running, with a negligible pileup contribution and an inter-bunch spacing of 500 ns (200 ns). During the pPb run in 2013 an instantaneous luminosity of $10^{29} \text{cm}^{-2} \text{s}^{-1}$ was achieved, corresponding to an interaction rate of 200 kHz, again with a very low pileup contribution.

In PbPb collisions, the number of produced particles depends strongly on the geometrical overlap of the Pb ions at the time of the collisions. The number of charged particles produced per unit of pseudorapidity, $dN_{\text{ch}}/d\eta$, varies from event to event from ≈ 10 for glancing collisions to ≈ 1600 for head-on collisions. The large particle multiplicity of head-on collisions leads to very high detector occupancies in the inner layers of the silicon tracker. For such high detector occupancies

Table 8. Summary of the heavy ion running conditions in various data-taking periods.

Run period	Ion species ($\sqrt{s_{NN}}$)	Max. collision rate	Zero suppression
Winter 2010	PbPb (2.76 TeV)	200 Hz	Offline
Winter 2011	PbPb (2.76 TeV)	4500 Hz	HLT
Winter 2013	pPb (5.02 TeV)	200 kHz	FED

the hardware-based zero-suppression algorithm implemented in the front-end-drivers (FED) of the tracker does not function reliably. As a consequence the tracker had to be read out without hardware zero suppression and the zero suppression was performed offline in 2010 and in the HLT in 2011. Table 8 shows a summary of the conditions in various heavy ion running periods.

A consequence of reading out the tracker without zero suppression is the limited data throughput from the detector due to the large event size. This limits the readout rate of the detector to 3 kHz in PbPb collisions. The limit has to be taken into account when setting up the trigger menu for HI collisions.

The HI object reconstruction is based on the pp HLT reconstruction algorithms described in the previous sections. The physics objects or event selection criteria used in the trigger menu are the following:

- Hadronic interactions (minimum bias);
- Jets;
- Photons;
- Muons;
- High-multiplicity events.

In the following we discuss the differences between the algorithms used in pp running to those used offline, and the performance efficiencies of these algorithms in the PbPb case.

Hadronic interactions. Since the interaction probability per bunch crossing during HI data taking is only $\approx 10^{-3}$, it is necessary to deploy a dedicated trigger to select hadronic interactions. This selection is based on coincidences between the trigger signals from the $+z$ and $-z$ sides of either beam scintillation counters (BSCs) or the HF which cover a pseudorapidity range of $2.9 < |\eta| < 5.2$. This trigger has a selection efficiency of more than 97% for hadronic inelastic collisions and is thus also referred to as a “minimum bias” trigger. The selection efficiency of this trigger was determined using a MC simulation of HI events using the HYDJET event generator [60] and was cross-checked with a control data sample selected using the BPTX signal to identify crossing beam bunches. The event sample selected this way is referred to as “zero bias” sample. From the zero-bias sample, inelastic events can be selected by requiring a charged-particle track consistent with originating from the beam crossing region. The fraction of the zero bias sample selected using the minimum bias trigger is consistent with the selection efficiency determined from simulated events.




Figure 52. Efficiency curves for the Jet50U trigger in PbPb at $\sqrt{s_{NN}} = 2.76$ TeV, as a function of the corrected (left) and uncorrected (right) leading jet transverse momentum.

Jets. The jet reconstruction algorithm used for HI data taking closely follows the corresponding pp algorithm which reconstructs calorimeter-based jets as described in sections 3.5.1 and 3.5.4, with the addition of a step subtracting the high-multiplicity underlying event using the iterative pileup subtraction technique [61]. During the 2010 and 2011 HI data-taking periods an iterative-cone type algorithm was used for jet clustering.

The efficiency of the jet triggers deployed for the 2010 PbPb run is illustrated in figure 52 by the efficiency turn-on curve of the Jet50U trigger. This trigger was discriminating events based on uncorrected jet energies. The efficiencies are given as a function of leading-jet transverse momentum for offline-corrected (left) and for uncorrected jets (right). The given efficiencies were determined based on offline jets reconstructed using the iterative cone algorithm with pileup subtraction in a data sample collected using a minimum bias trigger. The efficiency is defined as the fraction of the minimum bias sample containing a leading jet of a given p_T that is selected by the jet trigger.

During the 2011 PbPb run, the jet triggers had energy corrections applied at the HLT, leading to sharper turn-on curves, and thereby to more efficient data taking. Figure 53 illustrates the improvement by showing the efficiency of the Jet80 trigger as a function of leading-jet transverse momentum in the $|\eta| < 2$ region. The efficiencies are evaluated from a minimum bias sample, as in the 2010 case, with the jet reconstructed using the anti- k_T algorithm based on PF objects and also subtracting the underlying event using the iterative pileup subtraction technique. The efficiencies are given for various cone radii.

Photons. During the 2010 PbPb run the photon triggers employed at HLT were based on the energy deposits in the ECAL reconstructed using the island clustering algorithm [45]. This is the same algorithm as used for offline analyses based on the 2010 data, but without energy correction already applied at HLT. The trigger efficiency for the uncorrected Photon15 trigger for minimum bias events is shown in the left panel of figure 54.

For the data taking of 2011, energy corrections were already applied in the HLT. The performance of such corrected HLT photon paths is illustrated in the right panel of figure 54, which shows




Figure 53. Efficiency curves for the Jet80 trigger in PbPb at $\sqrt{s_{NN}} = 2.76$ TeV, as a function of the leading jet transverse momentum in the $|\eta| < 2$ region evaluated from minimum bias sample. The red, black, and blue points correspond to anti- k_T jets with cone size of 0.2, 0.3, and 0.4, respectively.




Figure 54. Trigger efficiency of the uncorrected Photon15 (left) and the corrected Photon40 (right) triggers as a function of corrected offline photon transverse momentum, in PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV.

the efficiency turn-on curve for the Photon40 trigger, again determined with respect to minimum bias events.

Muons. Efficient triggering on high- p_T muons is of primary importance for the HI physics program in CMS. During data-taking both single- and double-muon triggers were deployed to allow for maximal flexibility in event selection.

The per-muon trigger efficiency of the double-muon trigger (which requires two muons with $p_T > 3$ GeV) in the 2011 PbPb data determined by a tag-and-probe method is shown in figure 55. The three panels show the efficiency as a function of transverse momentum, pseudorapidity, and the overlap between the two colliding nuclei, expressed by the “number of participants.” Data are




Figure 55. Per-muon triggering efficiency of the HLT HI double-muon trigger as a function of p_T (left), η (center), and average number of participant nucleons (right). Z bosons in data (red) are compared to simulated Z bosons embedded in HI background simulated with HYDJET (blue).




Figure 56. Single-muon trigger efficiencies as functions of probe muon transverse momentum, pseudorapidity, and number of participants in the 2011 PbPb data. Red full circles are simulation and blue full squares are data. The numbers quoted in the legends of the figures are the integrated efficiencies.

shown in red and simulated Z bosons embedded in HYDJET background are shown in blue. On average, the trigger efficiency is very good, reaching 98.2% as obtained from tag-and-probe with simulated data.

The single-muon trigger efficiencies for the daughters of J/ψ mesons with $p_T > 6.5$ GeV in the 2011 PbPb data as a function of transverse momentum, pseudorapidity, and the number of participants are shown in the various panels of figure 56. The p_T and η integrated trigger efficiency is $86.0 \pm 0.2\%$ in MC, and $91.5 \pm 0.4\%$ in data. The trigger efficiency shows no significant dependence on the number of participants, as expected, in data or simulation (figure 56, right).

High-multiplicity events. In order to trigger on high-multiplicity events, several trigger paths were deployed during the HI data-taking periods. Triggers based on energy deposits in the calorimeter systems, signals in the BSC detectors, as well as triggers based on track multiplicities were employed and used in supplementary roles. The efficiency of high-multiplicity track triggers used during the 2013 pPb run is shown in the left panel of figure 57. The histograms correspond to different thresholds of the same kind as for track-based triggers. The efficiencies are shown as a function of the offline track multiplicity. The efficiencies are determined using either minimum bias events or a lower threshold high multiplicity trigger as a reference. The efficiency is defined as the fraction of events passing a given trigger threshold in the reference sample and is shown as a function




Figure 57. Left: trigger efficiency as a function of the offline track multiplicity, for the three most selective high-multiplicity triggers. Right: the spectrum of the offline tracks for minimum bias and for all the different track-based high-multiplicity triggers in the 2013 pPb data.

of number of offline reconstructed tracks. The gain in the number of high-multiplicity events is demonstrated in the right-hand side panel of figure 57.

4 Physics performance of the trigger

In the previous sections, we described the performance of the CMS trigger system for single- and multi-object triggers. However, most physics analyses published using the data taken in the first years of the LHC were performed using more complicated triggers. These triggers either take advantage of different categories of objects, such as a mixture of jets and leptons, or are topological triggers, which look at the event as a whole and calculate quantities such as the scalar sum of jet transverse energy H_T in the event or the missing transverse energy. In this section, to illustrate the performance of the trigger system, we give specific examples of some high-priority analyses that CMS carried out based on data taken in 2012, at a center-of-mass energy $\sqrt{s} = 8$ TeV.

4.1 Higgs boson physics triggers

The observation of the Higgs boson [62, 63] is the most important CMS result in the first LHC run. In this section, we discuss the CMS trigger performance for Higgs boson physics. Single-object triggers were discussed in section 3. In this section, more complex triggers are described. The strategy of combining different trigger paths to maximize the signal acceptance for the Higgs boson measurements is also presented.

4.1.1 $h \rightarrow \gamma\gamma$

As already discussed in section 3.3.1, diphoton triggers have been designed to efficiently collect $H \rightarrow \gamma\gamma$ events. To be as inclusive as possible, any photon that passes the general identification

requirements described in section 3.3.1, and either the isolation and calorimeter identification or the R_9 requirement, is accepted in the diphoton path. Asymmetric thresholds of 26 GeV on the leading photon and 18 GeV on the subleading photon have been applied together with a minimum invariant mass requirement on the diphoton system of 60 GeV. In the very late 2012 data-taking period, a similar path with more asymmetric E_T requirements was added to the HLT menu to enhance the discriminating power for the non-standard Higgs boson spin-0 and spin-2 scenarios. The performance of the trigger was shown in figures 20 to 21.

4.1.2 $H \rightarrow ZZ \rightarrow 4\ell$

The four-lepton channel provides the cleanest experimental signature for the Higgs boson search: four isolated leptons originating from a common vertex. As the number of expected events is very low, it is necessary to preserve the highest possible signal efficiency. The analysis performance therefore heavily relies on the lepton reconstruction, identification efficiency, and, due to the low branching fraction of the Higgs boson into ZZ , a robust trigger strategy to avoid any signal loss. The events are selected requiring four leptons (electrons or muons) satisfying identification, isolation, and impact parameter requirements (sections 3.3.1 and 3.4.2). The triggers described in this section were instrumental in the Higgs boson discovery and in the studies of its properties [62, 64].

In the following, we will describe the main triggers that are used to collect most of the data, as well as a set of utility triggers used to measure the online selection efficiencies. The main trigger selects $H \rightarrow ZZ \rightarrow 4\ell$ events with an efficiency larger than 95% for $m_h = 125$ GeV, at a rate less than 10 Hz at an instantaneous luminosity of $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. This trigger has loose isolation and identification requirements applied, and these are critical for proper background estimation. To improve the absolute trigger efficiency, a combination of single-electron and dielectron triggers was used. This combination achieved a 98% overall trigger efficiency.

For the $H \rightarrow ZZ \rightarrow 4\ell$ analysis, a basic set of double-lepton triggers is complemented by the triple-electron paths in the 4e channel, providing an efficiency gain of 3.3% for signal events with $m_H = 125$ GeV. The minimum momenta of the first and second lepton are 17 and 8 GeV, respectively, for the double-lepton triggers, while they are 15, 8 and 5 GeV for the triple-electron trigger. The trigger paths used in 2012 are listed in table 9, where “CaloTrk” stands for calorimeter- and tracker-based identification and isolation requirements applied with very loose criteria, while the “CaloTrkVT” name denotes triggers that make use of the same objects as discriminators, but with more stringent requirements placed on them.

Figure 58 shows the efficiency of the trigger paths described above as a function of the Higgs boson mass, for signal events with four generated leptons in the pseudorapidity acceptance and for those that have passed the analysis selection, as determined from simulation. With these trigger paths, the trigger efficiency within the acceptance of this analysis is greater than 99% for a Higgs boson signal with $m_H > 120$ GeV.

The tag-and-probe method is used to measure the per-lepton efficiency for double-lepton triggers, as described in section 3.3.1 for electrons, and in section 3.4.2 for muons. The performance in data and simulation of the per-leg efficiencies of the double-lepton triggers are shown in those sections. The position and the steepness of the turn-on curve of the trigger efficiency as a function of the lepton p_T measured on data is in good agreement with the expectations from simulation for all the triggers considered. A measurement of the trigger efficiency on the plateau reveals generally

Table 9. Triggers used in the $H \rightarrow 4\ell$ event selection (2012 data and simulation). No prescaling is applied to these triggers.

Channel	HLT path	L1 seed
4e	HLT_Ele17_CaloTrk_Ele8_CaloTrk	L1_DoubleEG_13_7
	OR HLT_Ele15_Ele8_Ele5_CaloIdL_TrkIdVL	L1_TripleEG_12_7_5
4 μ	HLT_Mu17_Mu8	L1_Mu10_MuOpen
	OR HLT_Mu17_TkMu8	L1_Mu10_MuOpen
2e2 μ	HLT_Ele17_CaloTrk_Ele8_CaloTrk	L1_DoubleEG_13_7
	OR HLT_Mu17_Mu8	L1_Mu10_MuOpen
	OR HLT_Mu17_TkMu8	L1_Mu10_MuOpen
	OR HLT_Mu8_Ele17_CaloTrk	L1_MuOpen_EG12
	OR HLT_Mu17_Ele8_CaloTrk	L1_Mu12_EG6




Figure 58. Trigger efficiency for simulated signal events with four generated leptons in the pseudorapidity acceptance (left), and for simulated signal events that have passed the full $H \rightarrow 4\ell$ analysis selection (right).

lower efficiency in data, compared to simulation, by about 1–2%. The effect of this inefficiency is mitigated, however, by the fact that multiple leptons in the event can pass the trigger requirements, and so no correction factor is applied. A systematic uncertainty of 1.5% in the expected signal yields is included to allow for this difference in trigger performance between data and simulation.

In table 10, the trigger paths used to select the tag-and-probe pairs for the efficiency measurements are listed. In case of muons, the prescaled double-muon triggers in the J/ψ mass window are used to select a low- p_T muon probe to measure the identification and isolation efficiency for muons with $p_T < 15$ GeV.

Table 10. Triggers used for tag-and-probe (T&P) efficiency measurements of four-lepton events in 2012 data and simulation: CaloTrk = CaloIdT_CaloIsoVL_TrkIdVL_TrkIsoVL, CaloTrkVT = CaloIdVT_CaloIsoVT_TrkIdT_TrkIsoVT.

Channel	Purpose	HLT path	L1 seed	prescale
4e and 2e2 μ	Z T&P	HLT_Ele17_CaloTrkVT_Ele8_Mass50	L1_DoubleEG_13_7	5
4e and 2e2 μ	Z T&P low p_T	HLT_Ele20_CaloTrkVT_SC4_Mass50_v1	L1_SingleIsoEG18er	10
4 μ and 2e2 μ	Z T&P	HLT_IsoMu24_eta2p1	L1_SingleMu16er	
4 μ and 2e2 μ	J/ψ T&P	HLT_Mu7_Track7_Jpsi		
		HLT_Mu5_Track3p5_Jpsi		
		HLT_Mu5_Track2_Jpsi		

Table 11. List of L1 and HLT used for 2012 data for the $Z(\nu\bar{\nu})H(b\bar{b})$ channel. We use PF E_T^{miss} . All triggers are combined to maximize acceptance. In all cases, an OR of the L1 $E_T^{\text{miss}} > 36$ GeV and $E_T^{\text{miss}} > 40$ GeV are used as the L1 seed.

HLT	Run Period
$E_T^{\text{miss}} > 150$ GeV	2012
$E_T^{\text{miss}} > 80$ GeV and 2 central jets with $p_T > 30$ GeV	early 2012
$E_T^{\text{miss}} > 100$ GeV and 2 central jets and $\Delta\phi$ requirement	late 2012
$E_T^{\text{miss}} > 100$ GeV and 2 central jets with $p_T > 30$ GeV and at least one b tag	late 2012

4.1.3 $H \rightarrow \tau\tau$

The triggers used for the Higgs boson $H \rightarrow \tau\tau$ analysis in the $\tau_\mu\tau_h$ and $\tau_e\tau_h$ channels require both an electron or muon and a hadronic tau object. The electron or muon is required to be isolated; the energy in the isolation cone is corrected for the effects of the pileup [53]. The tracks for the τ_h candidate and the tracks used to compute the isolation were required to come from a vertex compatible with the electron/muon origin. The efficiencies are measured using $Z \rightarrow \tau\tau$ events with a muon-plus- E_T^{miss} or a single-electron trigger. The events are selected by requiring the electron/muon to pass the tight isolation criteria, and also to have a transverse mass $M_T < 20$ GeV measured between the electron/muon and the missing transverse momentum vector. The purities after this selection are 78% and 65% for $|\eta(\tau_h)| < 1.5$ and $1.5 < |\eta(\tau_h)| < 2.3$, respectively. The event samples used to calculate the efficiencies are mixed with W+jets simulated events to produce a compatible purity. The τ -leg trigger efficiencies are discussed in detail in section 3.6.2.

4.1.4 $Z(\nu\bar{\nu})H(b\bar{b})$

The production of the Higgs boson in association with vector bosons is the most effective way to observe the Higgs boson in the $H \rightarrow b\bar{b}$ decay mode [65]. In this section, we report on the trigger performance for the 2012 data taking period. Table 11 summarizes these triggers. The main trigger requires $E_T^{\text{miss}} > 150$ GeV and was active during the entire year. This trigger, however, attains an efficiency of 95% at $E_T^{\text{miss}} \approx 190$ GeV, as shown in figure 59 (left). To accept events with lower E_T^{miss} , we introduce a trigger that requires two central PF jets with $p_T > 30$ GeV and $E_T^{\text{miss}} > 80$ GeV,




Figure 59. Trigger efficiencies for the $Z(\nu\bar{\nu})H(b\bar{b})$ analysis, as a function of offline PF E_T^{miss} for the pure $E_T^{\text{miss}} > 150$ GeV trigger (left) using late 2012 data, dijet and E_T^{miss} trigger (center) using early 2012 data, and dijet, E_T^{miss} and $\Delta\phi$ requirement trigger (right) using 2012 late data, as described in the text.

for early 2012 data. This trigger recovers events at lower E_T^{miss} . The efficiency curve, shown in figure 59 (center) reaches a plateau of 95% at $E_T^{\text{miss}} \approx 150$ GeV.

For late 2012 running, jets due to pileup caused an increase in trigger rates, and a more complicated trigger, requiring at least two central PF jets with $p_T > 60(25)$ GeV for the leading (subleading) jet, was introduced. At least one calorimeter dijet pair with $|\sum_i \vec{p}_{T,i}| > 100$ GeV is required. The minimum $\Delta\phi$ between the E_T^{miss} and the closest calorimeter jet with $p_T > 40$ GeV is required to be greater than 0.5. Finally, we require PF $E_T^{\text{miss}} > 100$ GeV. The obtained turn-on curve for this trigger is shown in figure 59 (right). The trigger achieves 90% efficiency at $E_T^{\text{miss}} \approx 170$ GeV, with roughly 80% efficiency for E_T^{miss} in the range of 130–170 GeV.

To accept events with even lower E_T^{miss} (down to 100 GeV) we exploit triggers with a b-tag online requirement (section 3.7): two jets with $p_T > 20(30)$ GeV and $E_T^{\text{miss}} > 80$ GeV for early (late) data. These triggers by themselves achieve an efficiency of roughly 50% at $E_T^{\text{miss}} = 100$ GeV and 60% efficiency for E_T^{miss} between 100 and 130 GeV when requiring at least one PF jet with a high value of the b-tagging discriminator (tight CSV > 0.898) offline. The trigger strategy for the full 2012 period used the combination of all the aforementioned triggers to collect events with $E_T^{\text{miss}} > 100$ GeV.

Rather than measuring the efficiency curves directly in data and applying them to the simulation, the efficiencies of the simulated triggers are parametrized and corrected as a function of E_T^{miss} and the CVS b tagging discriminator to match the efficiencies measured in data (described below). This approach takes into account the non-negligible correlations among the various trigger paths. It also characterizes the online b tagging efficiency and its dependence on jet p_T and η , as the geometry and trigger algorithm are simulated in a way that are as close as possible to the actual trigger environment. Studies show that the data and simulation agree to within less than 5%, except for the b tag trigger, where the agreement is approximately 10–20%.

In figure 60, we show the total trigger efficiency as a function of E_T^{miss} for signal $Z(\nu\bar{\nu})H(b\bar{b})$ events. The cumulative efficiency is 99% for $E_T^{\text{miss}} > 170$ GeV, 98% for events with $130 < E_T^{\text{miss}} < 170$ GeV, and 88% for events with $100 < E_T^{\text{miss}} < 130$ GeV.




Figure 60. Efficiency as function of E_T^{miss} for the $Z(\nu\bar{\nu})H(b\bar{b})$ signal events. An efficiency greater than 99% is obtained for $E_T^{\text{miss}} > 170$ GeV.

The total systematic uncertainty in the trigger efficiency is of the order of a few percent in the high- p_T ($E_T^{\text{miss}} > 170$ GeV), and not more than 7% in the intermediate- p_T ($130 < E_T^{\text{miss}} < 170$ GeV), and 10% in the low- p_T regions ($100 < E_T^{\text{miss}} < 130$ GeV) search regions.

4.2 Top quark triggers

Measurement of the properties of the top quark are among the most important standard model measurements in CMS. The LHC is a top factory, and the large number of top quark pairs created allows detailed studies of its properties. One of the most fundamental measurements is the top quark pair production cross section. The most accurate measurement of this cross section can be made in the so-called ‘lepton + jets’ decay mode, where one of the W bosons from the top quark decays to a lepton and a neutrino, and the other W decays hadronically, leading to a final state with a well-isolated lepton, large missing transverse energy, and four hadronic jets (two of which are b jets) [66, 67]. In Run 1, $t\bar{t}$ production studies used several trigger paths for the semileptonic top quark decay channels, to ensure that $t\bar{t}$ signal events were recorded as efficiently as possible. To maximize the acceptance of the transverse energy (momentum) requirement applied to leptons, measurements used trigger paths requiring one online reconstructed lepton (e or μ) as well as at least 3 online reconstructed jets.

Table 12 summarizes the main paths used for the triggers deployed to accommodate the high instantaneous luminosity and pileup of the 2012 run. All leptons triggers had tight or very tight lepton identification and calorimeter isolation requirements, comparable to those used offline. Jets in PF jet triggers were restricted to the central region. At L1, single electrons or muons are required with the denoted thresholds. The L1 muons are central ($|\eta| < 2.1$). Charged-hadron subtraction [68] (labeled ‘pileup subtracted’ in the table) was implemented for pileup mitigation. Additionally, the introduction of jet energy calibrations online in the second half of 2012 resulted

Table 12. Unscaled cross-triggers used for the $t\bar{t}$ (lepton plus jets) cross section measurement in 2012. All leptons use tight or very tight identification, and lepton and calorimeter isolation requirements. All jets are PF jets and restricted to the central region. At L1, single electrons or muons are required with the denoted thresholds and the L1 muons are required to be central ($|\eta| < 2.1$). When two thresholds are listed at L1, they include a lower (possibly prescaled) threshold and a higher unscaled threshold.

HLT					L1	
e/μ	Threshold (GeV)	n_{jet}	Jet corrections	Jet threshold	L1 Seed	Threshold (GeV)
e	25	3		30	EG	20, 22
	25	3	pileup subtracted	30	EG	20, 22
	25	3	pileup subtracted	30, 30, 20	EG	20, 22
	25	3	pileup subtracted	45, 35, 25	EG	20, 22
μ	20	3		30	MU	14, 16
	20	3	pileup subtracted	30	MU	16
	17	3	pileup subtracted	30	MU	14
	17	3	pileup subtracted	30, 30, 20	MU	14
	17	3	pileup subtracted	45, 35, 25	MU	14

in higher E_T thresholds in the three-jet paths; however, the data from that period were not used in the cross section measurements due to systematic uncertainties associated with the large pileup.

Simulated events are used to estimate the top quark acceptance, and were corrected for the trigger efficiency measured in data. To estimate the trigger efficiency, simulated Drell-Yan and $t\bar{t}$ samples were used to compare with data collected with single lepton triggers. The overall efficiency for the lepton plus jets paths is parametrized as a product of two independent efficiencies for the leptonic and hadronic legs of the trigger, $\epsilon_{\text{lep}} \times \epsilon_{\text{had}}$. A cleaning requirement based on the ΔR distance between the leptons and jets motivates this approach.

The leptonic leg efficiency is measured using a tag-and-probe method with Z/γ^* events, as described in sections 3.3.1 (e) and 3.4.2 (μ). While the lepton trigger was not changed during the 2012 data-taking period, the jet trigger changed as shown in table 12. Similar to the measurement for the lepton leg, the efficiency of the jet leg of the associated cross-trigger is measured in an unbiased data sample. The reference sample is required to pass a single lepton trigger, to assure a data set independent of the hadronic trigger which fulfills the lepton leg of the cross-trigger simultaneously.

As an example, figure 61 shows the efficiency turn-on curve of the hadronic leg (transverse momentum of the 4th jet) for the electron plus jets paths in 2012, and its dependence with respect to the number of reconstructed vertices, both for a selection based on the combination of the PF jets without and with charged-hadron subtraction. The offline selection of the transverse momentum requirements on the offline jets was devised to assure a plateau behavior of the scale factors, meaning no variation of the scale factor with respect to the MC sample or jet energy calibrations. From the variation of the scale factors it was concluded that a systematic uncertainty of 2% (1.5%) in electron (muon) scale factors covered the variations around their value of 0.995 (0.987).




Figure 61. Top quark triggers: efficiency of the hadronic leg for the electron plus jets paths in 2012 as a function of the p_T of the 4th jet (left) and of the number of reconstructed vertices (right).

4.3 Triggers for supersymmetry searches

Supersymmetry (SUSY) is one of the most appealing extensions to the standard model, as it solves the mass hierarchy problem, offers a path towards grand unification, and can provide candidate dark matter particles. During the years 2010–2012, many SUSY searches were performed with CMS data. Exclusion limits were set in the context of the mSUGRA model of SUSY breaking and also on the masses of the particles involved in specific cascade decays (simplified models [69]).

For the allowed parameter space, SUSY signatures [70] are characterized by the presence and decay of heavy particles. If R-parity is conserved, stable, invisible particles are expected. Most of the final states contain significant hadronic activity and E_T^{miss} . At CMS, SUSY searches were divided into leptonic, hadronic, and photonic categories, depending on the event content.

In addition, some supersymmetric models predict the existence of heavy stable charged particles, e.g., the gluino, top quarks, or τ sleptons. Their mass is expected to be of the order of a few hundred GeV, therefore their velocity would be significantly smaller than the speed of light. The signature of heavy stable charged particles would look like a non-relativistic ionizing particle, with hits in the chambers being delayed by about one bunch crossing, either in all the layers or in the outermost one(s), with respect to an ordinary “prompt” minimum ionizing particle.

In this section we discuss the CMS trigger performance collecting events for searches for supersymmetry. Most leptonic searches in CMS were performed using the same triggers as the Higgs boson leptonic searches and therefore are not documented here. For hadronic and photonic searches, we have selected three representative triggers: the α_T trigger, the “Razor” trigger, and the photon trigger. The α_T and photon analyses were performed using a data sample corresponding to an integrated luminosity of 4 fb^{-1} , while the Razor analysis used an integrated luminosity of 20 fb^{-1} , all collected at CMS during 2012 at a center-of-mass energy of 8 TeV.

4.3.1 Triggers for all-hadronic events with α_T

We present a typical example of a purely hadronic search, where events with leptons are vetoed and events with a high jet multiplicity, large E_T^{miss} , and large H_T are selected [71]. Multijet events are the most important background in this region of the phase space. To suppress these events, the




Figure 62. Efficiency turn-on curves for the α_T triggers used to collect events for four different H_T regions: $275 < H_T < 325$ GeV (upper left), $325 < H_T < 375$ GeV (upper right), $375 < H_T < 475$ GeV (lower left), and $H_T > 475$ GeV (lower right).

analysis uses a kinematical variable called α_T . For events with exactly two jets, α_T is defined as the transverse energy of the subleading jet divided by the transverse mass of the dijet system. For events with two or more jets, two pseudo-jets are created combining jet components and selecting the configuration that minimizes the energy between the two. The value of α_T is equal to 0.5 in balanced multijet events and less than 0.5 in multijet events with jet energy mismeasurement. For SUSY signal events with genuine E_T^{miss} , α_T tends to values > 0.5 , thus providing a good discrimination between signal and background. To estimate the remaining significant backgrounds (W+jets, top quark pair, single top quark, and $Z \rightarrow \nu\bar{\nu}$), data control regions are used.

A cross-trigger based on the quantities H_T and α_T is used to record the candidate event sample. A prescaled H_T trigger, labeled henceforth as H_T , is used with various thresholds to record events for the control region. The H_T thresholds of the H_T and H_T - α_T cross-triggers are chosen to match where possible, and are 250, 300, 350, 400, and 450 GeV. The α_T thresholds of the H_T - α_T trigger are tuned according to the threshold on the H_T leg in order to suppress QCD multijet events (whilst simultaneously satisfying other criteria, such as sensitivity to trigger rates).

To ensure that the H_T leg of the H_T - α_T cross-trigger and the H_T prescaled trigger are efficient for the final event selection, the lower bounds of the offline H_T bins are offset by 25 GeV with respect to the online thresholds. Figure 62 shows the turn-on curves of the H_T and α_T legs of the trigger with respect to the offline selection.

Table 13. Measured efficiencies of the H_T and H_T - α_T triggers, as a function of α_T and H_T , as measured with respect to the offline selection used in the α_T analysis.

α_T lower threshold	H_T range (GeV)	Efficiency(%)
0.55	275–325	$89.6^{+0.5}_{-0.6}$
0.55	325–375	$98.5^{+0.3}_{-0.5}$
0.55	375–475	$99.0^{+0.5}_{-0.6}$
0.55	475– ∞	$99.4^{+0.5}_{-1.2}$

Efficiencies for the H_T - α_T triggers were calculated using an orthogonal data set based on single muons, by requiring a matching to an isolated single-muon trigger. Exactly one isolated muon that is well separated from all jets is required to “tag” the event. This muon is not considered in the calculations of H_T , E_T^{miss} -like quantities, and α_T , thereby miscalculating genuine E_T^{miss} by ignoring the muon. The assumption for the H_T triggers is that their efficiency is not sensitive to whether there is genuine E_T^{miss} in the event or not. The results (efficiencies with respect to offline selection) are shown in table 13.

4.3.2 Triggers for inclusive search with Razor variables

The Razor variables R^2 and M_R were introduced in CMS to complement other variables that can be used to probe SUSY production at the LHC [72, 73]. The analyses are designed to kinematically discriminate the pair production of heavy particles from SM backgrounds, without making strong assumptions about the E_T^{miss} spectrum or details of the decay chains of these particles. The baseline selection requires two or more reconstructed objects, which can be calorimetric jets, isolated electrons or isolated muons. The Razor kinematic construction exploits the transverse momentum imbalance of SUSY events more efficiently than the traditional E_T^{miss} -based variables, retaining events with as low as $E_T^{\text{miss}} \approx 50$ GeV while reducing the background from QCD multijet events to a negligible level. Details of the definition of R^2 and M_R can be found in the above references.

The use of E_T^{miss} and H_T triggers alone would not be practical for a Razor-based search, resulting in a nontrivial dependence of the trigger efficiency on R and M_R . Instead, a set of dedicated triggers was developed, both for the fully hadronic and the leptonic final states considered in the analysis.

The Razor triggers are based on the events with two central jets with $p_T > 64$ GeV, selected at L1. The calorimetric towers in the event are clustered using the anti- k_T algorithm with a distance parameter of 0.5. The two highest p_T jets are required to have $p_T > 65$ GeV, which is fully efficient for PF jets with $p_T > 80$ GeV. If an event has more than seven jets with $p_T > 40$ GeV, it is accepted by the trigger. Otherwise, we consider all the possible ways to divide the reconstructed jets in two groups. We then form a *mega-jet* summing the four-momenta of the jets in one group. The mega-jet pair with the smallest sum of invariant masses is used to compute the values of R and M_R . A selection on R and M_R is applied to define a leptonic Razor trigger. A looser version of this selection is used for the lepton Razor triggers, in association with one isolated muon or electron with $p_T > 12$ GeV. Electrons are selected with a loose calorimeter identification requirement and a very loose isolation requirement. The kinematic selection includes cuts on both on R and M_R : $R^2 > 0.09$ and $M_R > 150$ GeV (inclusive trigger); $R^2 > 0.04$ and $M_R > 200$ GeV (leptonic




Figure 63. Turn-on curve for M_R (left) and R^2 (right) for the inclusive Razor trigger, after requiring $R^2 > 0.25$ (left) and $M_R > 400$ GeV (right). Events passing the single-electron triggers are selected to define the denominator of the efficiency, together with the dijet requirement. The requirement of satisfying the Razor trigger defines the numerator.

triggers). A “parked” version (as described in section 2.6) of the inclusive Razor trigger was also implemented, requiring $R^2 > 0.04$.

Events selected by the single-electron (single-muon) triggers are used to measure the efficiency of the inclusive and electron (muon) Razor paths. The baseline sample for the efficiency measurement is defined requiring two jets of $p_T > 80$ GeV, passing the reference trigger, and not rejected by the event cleanup requirements (designed to remove the noisy calorimeter events from the offline analysis). The numerator of the efficiency is defined from this sample, with the requirement that the relevant Razor trigger condition is satisfied. Figure 63 shows the efficiency versus M_R and R^2 for the inclusive Razor trigger, also requiring $M_R > 400$ GeV ($R^2 > 0.25$) in order for the R^2 (M_R) plot to match the selection applied in the analysis. The efficiency is found to be flat within the statistical precision, limiting the precision on the tail or R^2 after the applied M_R requirement. The analysis uses $(95 \pm 5)\%$ as an estimate of the efficiency.

4.3.3 Triggers for photons and missing energy

We present the triggers used in a search for supersymmetry in events with at least one isolated photon, jets, and E_T^{miss} . Dominant standard model background processes are direct photon production and QCD multijet events where a jet is misreconstructed as a photon. Multijet events have small intrinsic E_T^{miss} , but the finite resolution of the jet energy measurement together with the large cross section leads to a significant contribution in the tail of the E_T^{miss} . Other backgrounds arise from electroweak electron production, e.g., $W \rightarrow e\nu$, where an electron is misreconstructed as a photon. Additional contributions are expected from initial- or final-state photon radiation in various QCD and electroweak processes. Single-photon trigger thresholds are too high for the efficient selection of many SUSY benchmark points, so that for this analysis a cross-trigger based on a single photon and H_T is used. The main backgrounds are modeled using data control samples.

To trigger on the signal as well as to collect the control samples used for estimation of the QCD multijet and electroweak backgrounds, a cross-trigger is used, requiring at least one photon with $p_T > 70$ GeV and $H_T > 400$ GeV. The control region is defined by events containing at least one




Figure 64. Supersymmetry search in the $\gamma + E_T^{\text{miss}}$ channel: trigger efficiency of the H_T leg (left column), and the photon leg (right column), using as a reference the single-photon trigger with $p_T > 50$ GeV (top row) and $p_T > 75$ GeV (bottom row). The red lines indicate offline requirements.

isolated photon with $p_T > 80$ GeV and $|\eta| < 1.4$, two or more jets with $p_T > 30$ GeV and $|\eta| < 2.6$, and $H_T > 450$ GeV. The signal region includes an additional $E_T^{\text{miss}} > 100$ GeV requirement.

The trigger efficiency was measured in data for the photon and H_T legs, using a single-photon baseline trigger, which requires a single photon with $p_T > 50$ GeV and is expected to be fully efficient in the kinematic region of interest. As the statistical power of the data sample is limited by the large prescale of the baseline trigger (prescale of 900), a cross-check is performed using a less prescaled single photon trigger with $p_T > 75$ GeV (prescale of 150). In this case, it is not possible to observe the p_T turn-on of the photon leg efficiency, as the baseline selection is more restrictive than the online selection used by the analysis; however this is a valid check of the H_T leg. Figure 64 shows the turn-on curve for the H_T and photon p_T legs, both single-photon triggers. Only in the H_T leg for the single-photon trigger with the $p_T > 75$ GeV requirement, a higher threshold in the photon




Figure 65. The principle of operation of the RPC HSCP trigger for an ordinary muon (case 1), and a slow minimum ionizing particle, which produces hits across two consecutive bunch crossings (cases 2, 3) or in the next BX (case 4). Hits that would be seen in the standard PAC configuration are effectively those shown in pale orange; additionally observed hits in the HSCP configuration are those shown in dark orange. In case 1 the output of both configurations is identical, in case 2 the HSCP configuration uses the full detector information, in case 3 only the HSCP configuration can issue a trigger, and in case 4 the HSCP configuration brings back the event to the correct BX.

$p_T > 85 \text{ GeV}$ is used to avoid regions with inefficiencies due to the cross-trigger. After applying the offline analysis requirements on the photon momentum of $p_T > 80 \text{ GeV}$ and on $H_T > 450 \text{ GeV}$, indicated in the figure, the trigger is fully efficient within an uncertainty of 4%. The uncertainty is due to the low statistical power of the data set.

4.3.4 Triggers for heavy stable charged particles

The CMS experiment has a specific RPC muon trigger configuration to increase the efficiency for triggering on heavy stable charged particles (HSCP) using the excellent time resolution of detected muon candidates. Double-gap RPCs operating in avalanche mode have an intrinsic time resolution of around 2 ns. This, folded with the uncertainty coming from the time propagation along the strip, which contributes about 2 ns, and the additional jitter that comes from small channel-by-channel differences in the electronics and cable lengths, again of the order of 1–2 ns, give an overall time resolution of about 3 ns — much lower than the 25 ns timing window of the RPC data acquisition system.

If hits are not in coincidence within one BX, the RPC PAC algorithm is likely to fail because the minimum plane requirements would not be met, or if the algorithm does succeed, a lower quality value and possibly a different p_T will be assigned to the trigger particle. In addition, if the muon trigger is one BX late with respect to the trigger clock cycle, the pixel hits will not be recorded and the muon chamber calibration constant will be suboptimal, resulting in a poor offline reconstruction of late “muon-like” candidates. The functionality to extend the RPC hits to two (or more) consecutive BXs, plus the excellent intrinsic timing capabilities of the RPCs, allow the construction a dedicated physics trigger for such “late muons”. In the PAC logic the RPC hits are extended in time to 2 BXs, hence the plane requirements are met for at least one BX and triggers can be issued. On the GMT input, the RPC candidates are advanced by one BX with respect to DT and CSC candidates, so hits of a “late muon” generate a trigger in the proper BX. Ordinary “prompt” muons will produce two trigger candidates: one in the proper BX and one in the previous BX. Misreconstructed candidates can, however, be suppressed at the GT level by a veto operated on the basis of BPTX coincidences (section 2.5). Figure 65 shows the principle of operation of the RPC-based HSCP trigger. Studies with simulated data indicate that the HSCP trigger configuration significantly increases the CMS capability to detect a slow HSCP, for example, for an 800 GeV long-lived gluino, the overall improvement in trigger efficiency ranges from 0.24 to 0.32. The gain is the largest within the range $200 < p_T < 600$ GeV and for gluino velocities $0.4 < \beta < 0.7$. The HSCP trigger configuration was the main RPC operation mode during data-taking in most of the 2011 and the entire 2012 run.

4.4 Exotic new physics scenarios

Models of physics beyond the standard model that are not supersymmetric are called ‘exotic’ in CMS. In this section we describe three exotic physics scenarios and the triggers used in searches for these signals.

4.4.1 Triggers for dijet resonance searches

During the 7 TeV run, the search for heavy resonances decaying to jet pairs was performed on events triggered by the single-jet trigger. With increasing peak luminosity, the tighter threshold applied on the jet p_T became a major problem for the analysis. At the same time, the analysis was improved by introducing the so-called *wide jets* to take into account the presence of additional jets from final-state radiation. Wide jets are formed around a given set of *seed* jets, taking as input the other jets in the event. The four-momentum of each seed jet is summed with the four-momenta of other jets within $\Delta R < 1.1$ of the seed jet and with $p_T > 40$ GeV. A jet close to more than one seed jet is associated with the closest seed. With this new approach, a trigger based on $H_T = \sum_{\text{jet}} |p_T|$ is more efficient. A further improvement in the analysis was obtained by implementing a dedicated topology-based trigger, applying a looser version of the analysis reconstruction and selection requirements at the HLT:

- Wide jets were built by looking for jets with $p_T > 40$ GeV in a cone of size $\Delta R = 1.1$ around the two highest p_T jets;
- Multijet events were removed by requiring that the two wide jets fulfill $\Delta\eta < 1.5$.

During the 8 TeV run, events were kept if the wide jets built around the two highest p_T jets had an invariant mass larger than 750 GeV (*Fat750*). While this trigger alone would have performed




Figure 66. Dijet resonance search triggers. The HLT efficiency of $H_T > 650$ GeV, $H_T > 750$ GeV, and *Fat750* triggers individually, and their logical OR as a function of the offline dijet mass. The efficiency is measured with the data sample collected with a trigger path that requires $H_T > 550$ GeV. The horizontal dashed line marks the trigger efficiency $\geq 99\%$.

similarly to the H_T trigger already in use, the combination of the two triggers in a logical OR allowed us to recover the inefficiency for mass values close to the applied threshold, making the overall efficiency turn-on curve sharper. The loosest H_T -based L1 path (L1_HTT150) was used as a seed for all triggers. The trigger efficiency was measured in data, taking the events triggered by the prescaled $H_T > 550$ GeV trigger as a baseline. These events were filtered by applying the analysis selection (particularly, the $\Delta\eta$ requirement on the two wide jets) to define the denominator of the efficiency curve. The subset of these events also satisfying the analysis requirements defines the numerator of the efficiency. Figure 66 shows the trigger efficiency as a function of the offline dijet mass for individual triggers and for their logical OR. While the combination of the H_T and *Fat750* triggers already represents a sizable improvement with respect to the individual triggers, a further increase in the efficiency was obtained with the introduction of the PF-based H_T trigger. The combination of the three triggers made the analysis $\geq 99\%$ efficient for invariant masses above 890 GeV. As a result of the trigger improvements, the threshold for the dijet resonance search for the 8 TeV run was 100 GeV lower than would have been possible if the 7 TeV strategy had been used.

4.4.2 Triggers for black hole search

If the scale of quantum gravity is as low as a few TeV, it is possible for the LHC to produce microscopic black holes or their quantum precursors (“string balls”) at a significant rate [74–76].

Table 14. Black Hole trigger: unprescaled total jet activity HLT paths and their respective L1 seeds. The L1 seeds for a number of the HLT paths were revised during the data taking to account for higher instantaneous luminosity.

Path name	L1 seed	Data-taking period
HLT_HT750	L1_HTT150 OR L1_HTT175	Early
HLT_HT750	L1_HTT150 OR L1_HTT175 OR L1_HTT200	Late
HLT_PFHT650	L1_HTT150 OR L1_HTT175	Early
HLT_PFHT650	L1_HTT150 OR L1_HTT175 OR L1_HTT200	Late
HLT_PFHT700	L1_HTT150 OR L1_HTT175	Early
HLT_PFHT700	L1_HTT150 OR L1_HTT175 OR L1_HTT200	Late
HLT_PFHT750	L1_HTT150 OR L1_HTT175	Early
HLT_PFHT750	L1_HTT150 OR L1_HTT175 OR L1_HTT200	Late
HLT_PFN0PUHT650	L1_HTT150 OR L1_HTT175	
HLT_PFN0PUHT700	L1_HTT150 OR L1_HTT175	
HLT_PFN0PUHT750	L1_HTT150 OR L1_HTT175	

Black holes decay democratically, i.e., with identical couplings to all standard model degrees of freedom. Roughly 75% of the black holes decay products are jets. The average number of particles in the final state varies from roughly two (in case of quantum black holes) to half a dozen (semiclassical black holes and string balls). The microscopic black holes are massive objects, thus at least a few hundred GeV of visible energy in the detector is expected.

Since a priori we do not know the precise final state, we trigger on the total jet activity in an event. The common notation of such triggers is HLT_HT x , HLT_PFHT x , and HLT_PFN0PUHT x , where x denotes the total energy in GeV. All energies of HLT jets are fully corrected, and in the case of the HLT_PFN0PUHT x paths, pileup corrections are also applied to the HLT PF jets. The pileup subtraction is performed by first removing all of the jet’s charged hadrons not associated to the primary vertex, then calculating an energy offset based on the jet energy density distribution to remove the remaining pileup contribution. More details of the jet reconstruction at L1 and HLT are given in section 3.5.

After the jets are selected at both the L1 and the HLT, an H_T variable is calculated. In ref. [77], the jet E_T threshold at L1 is 10 GeV and the H_T thresholds are 150, 175, and 200 GeV (section 3.5.4.) These L1 triggers are used as seeds to the HLT algorithms. At the HLT, the jet E_T threshold is 40 GeV and the H_T thresholds have a range of 650–750 GeV. The unprescaled HLT paths and their L1 triggers are summarized in table 14. The L1 triggers for some of the “total jet activity” paths were updated in the middle of 2012 to account for higher instantaneous luminosity of the LHC. For simplicity, we refer to the data taking periods before (after) that change as “early” (“late”). In the previous iterations of the analysis [78, 79], the H_T thresholds at the HLT were as low as 100–350 GeV.

As the majority of the final-state objects are jets, we use jet-enriched collision data to search for black holes. These data are recorded using a logical OR of the following trigger groups, whose




Figure 67. Left: efficiency of unscaled total jet activity HLT paths as a function of S_T . Right: efficiency of HLT_PFNPUHT650 as a function of S_T in three bins of number of primary vertices, N_{PV} : (i) $N_{PV} \leq 10$, (ii) $10 < N_{PV} < 25$, and (iii) $N_{PV} \geq 25$. All efficiencies are calculated with respect to a prescaled total activity path with $H_T = 450$ GeV threshold.

triggers only differ by a threshold: i) total jet activity triggers, ii) paths that select high-mass dijet events, iii) triggers that require presence of significant E_T^{miss} and a jet with p_T above a few hundred GeV. The main offline quantities that describe the black hole are the multiplicity of the final-state objects, N , and a scalar sum of transverse momenta of all objects (jets, leptons, and photons) and the E_T^{miss} reconstructed in the event, $S_T = \sum p_T^{\text{jets}} + \sum p_T^{\text{leptons}} + \sum p_T^{\text{photons}} + E_T^{\text{miss}}$. We apply a 50 GeV requirement on all final-state objects p_T and E_T^{miss} , and select events with a multiplicity greater than one. Note that E_T^{miss} is not counted towards the multiplicity. The relative efficiency of unprescaled HLT paths that are used in the analysis as a function of the S_T is shown in figure 67 (left). The efficiencies are calculated using the same jet-enriched data set with respect to prescaled total jet activity path with the H_T threshold of 450 GeV. The paths with the H_T threshold of 650 (750) GeV are fully efficient starting from $S_T = 1000$ (1200) GeV, respectively, which is significantly below the low- S_T boundary of 1500 GeV that is used in the search. To check the pileup dependence of the trigger turn-on, we plot the efficiency of the selection path HLT_PFNPUHT650 as a function of S_T in three bins of the number of reconstructed primary vertices, N_{PV} : i) $N_{PV} \leq 10$, ii) $10 < N_{PV} < 25$, and iii) $N_{PV} \geq 25$ (figure 67 (right)). Although the trigger turn-on curves become less sharp when N_{PV} increases, this does not affect the point when the trigger becomes fully efficient. Thus, the pileup dependence of total jet activity triggers can be neglected in the black holes analysis.

4.5 B physics and quarkonia triggers

The CMS analyses in the fields of B physics and quarkonium production are mostly based on data samples collected with double-muon triggers. In the 2010 run, the LHC instantaneous luminosity was sufficiently low such that relatively loose triggers could be used. Essentially all the analyses made in the B physics group were based on one inclusive trigger, which requires two high quality muons. The resulting dimuon mass distribution covers dimuon mass values from threshold all the way to 200 GeV, displaying “needles” caused by the dimuon decays of resonances on top of a smooth underlying continuum.




Figure 68. Dimuon mass distributions collected with the inclusive double-muon trigger used during early data taking in 2011. The colored areas correspond to triggers requiring dimuons in specific mass windows, while the dark gray area represents a trigger only operated during the first 0.2 fb^{-1} of the 2011 run.

The significantly higher collision rates of the 2011 LHC run, and the ceiling of around 25–30 Hz for the total trigger bandwidth allocated for B physics, required the development of several specific HLT paths, each devoted to a more or less exclusive set of physics analyses. Figure 68 illustrates the corresponding dimuon mass distributions, stacked on each other. The high-rate “low- p_T double muon” path was in operation only during the first few weeks of the run; the others had their rates reduced through suitable selection requirements on the dimuon mass and on the single-muon and/or dimuon p_T .

The quarkonia trigger paths (J/ψ , ψ' and Υ) had explicit requirements on the p_T of the dimuon system but not of the single muons. First, because the analyses are made as a function of the dimuon p_T and second, because the single-muon p_T requirements induce a significant restriction of the covered phase space in terms of the angular decay variable $\cos\theta$, and this is crucial for the measurements of quarkonium polarization. To further reduce the rate, the two muons were required to bend away from each other because the ones bending towards each other have lower efficiencies. The dimuon was required to have a central rapidity, $|y| < 1.25$. This is particularly useful to distinguish the $\Upsilon(2S)$ and $\Upsilon(3S)$ resonances, as well as for analyses of P-wave quarkonia production, which require the measurement of the photon emitted in the radiative decays (e.g., $\chi_c \rightarrow J/\psi \gamma$). In fact, to resolve the $\chi_{c1}(1P)$ and $\chi_{c2}(1P)$ peaks (or, even more challenging, the $\chi_{b1}(1P)$ and $\chi_{b2}(1P)$ peaks), it is very important to have a high-resolution measurement of the photon energy, possible through the reconstruction of the conversions into e^+e^- pairs in the barrel section of the silicon tracker.

In addition to the quarkonia resonances, figure 68 shows a prominent “peak” labeled B_s , which represents the data collected to search for the elusive $B_s \rightarrow \mu\mu$ and $B_d \rightarrow \mu\mu$ decays. These triggers had no restrictions on the dimuon rapidity or relative curvature and kept p_T requirements much looser than those applied in the offline analysis. The total rate of the B_s trigger paths remained




Figure 69. Single-muon detection efficiencies (convolving trigger, reconstruction, and selection requirements) as a function of p_T , as obtained from the data, using the tag-and-probe method. Data points are shown for the pseudorapidity range $|\eta| < 0.2$, while the curves (depicting a parametrization of the measured efficiencies) correspond to the three ranges indicated in the legend.

relatively small, of the order of 5 Hz, even when the LHC instantaneous luminosity exceeded $7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, at the end of the 2012 run.

The other prominent trigger path illustrated in figure 68, the “low-mass displaced dimuons”, selected events with a pair of opposite-sign muons with a dimuon vertex pointing back to and displaced from the interaction point by more than three standard deviations. These events were collected to study decays of B mesons into final states containing a pair of muons plus one or more kaon and/or pion, as well as to measure the Λ_b cross section, lifetime, and polarization. This is the most challenging trigger path because of its very high rate, which cannot be reduced through the increase of muon p_T requirements without a significant loss of signal efficiency.

The main difference between the 2011 and 2012 runs, from the perspective of B physics, was the availability of the so-called “parked data” (section 2.6). The resulting increase in available HLT bandwidth meant that most trigger paths could have looser requirements in 2012 than in 2011. Additionally, several new triggers were added, including a like-sign dimuon trigger to study the “anomalous dimuon charge asymmetry” observed at the Tevatron [80].

Two special calibration triggers were developed to study the single-muon detection efficiencies in an unbiased way. One is a single muon trigger that requires the presence of an extra track such that the invariant mass of the muon-track pair is in the J/ψ mass region; the existence of a J/ψ peak in this event sample ensures that the track is likely to be a muon that can be used to provide an unbiased assessment of the muon-related efficiencies (offline reconstruction in the muon detectors, as well as L1 and L2 trigger efficiencies as described in section 3.4.2). The other is a dimuon trigger for those low-mass dimuons in which the muons are reconstructed without using any information from the silicon tracker hits, thereby allowing the study of the offline tracking and track quality selection efficiencies, as well as the L3 trigger efficiency (section 3.4.2). These efficiency measurements are made using a tag-and-probe methodology. As an illustration, figure 69 shows the single-muon detection efficiency as a function of p_T for three muon pseudorapidity ranges.

The rate of events with single muons is very large and it might happen that a muon is mistakenly identified as two close-by muons. To prevent such events from increasing the rate of dimuon triggers,




Figure 70. Dimuon trigger efficiencies in the $\Delta\phi$ versus $\Delta\eta$ plane for J/ψ events generated in the kinematic region $p_T > 50$ GeV and $|y| < 1.2$, illustrating the efficiency drop when the two muons are too close to each other.

the trigger logic at L1 and L2 discards dimuon signals if the two muon trajectories are too close to each other. The drawback is that this significantly reduces the efficiency of the dimuon trigger for signal dimuons where the two muons are close to each other, which happens quite often for low-mass dimuons of high p_T . This drop in the dimuon trigger efficiency, shown in figure 70, is induced through a muon pair correlation and, hence, is not taken into consideration through the simple product of the efficiencies of the two single muons. The corresponding correction can be evaluated by MC simulation and validated by studying distributions of measured events as a function of the distance between the two muon tracks. In the 2012 run, a new trigger was developed, in which a high- p_T single muon selected at L1 and L2 is associated with a tracker muon at L3 before a dimuon mass range is imposed. In such events, there is only a single muon required at the L1 and L2 steps, so that the event is not rejected in case that there is a second muon very close by. This trigger path is ideally suited to study charmonium production at very high p_T .

5 Trigger menus

A trigger menu is defined as the sum of all object definitions and algorithms that define a particular configuration of the CMS trigger system. The menu consists of definitions of L1 objects and the algorithms that are used to render the L1 decision, as well as the configuration of the software modules that are used in the HLT. Sets of prescale columns for different instantaneous luminosities are also included. By means of such a prescale set the data-archiving rate of the readout chain could be adjusted and maximized during a LHC fill as the instantaneous luminosity drops along with the current trigger rate.

In this section, we describe the L1 and HLT menus and how they have evolved in response to the physics goals and significant performance improvements of the LHC machine during the first run.

5.1 L1 menus

From 2010 to 2012, several L1 menus (and corresponding prescale columns) were developed to meet the experiment's physics goals and to cope with the evolution of the LHC operational conditions, i.e., the change of the center-of-mass energy between 2011 and 2012, the varying number of colliding bunches for LHC fills, and the growth of luminosity per bunch. While designing new L1 menus, improved algorithms and thresholds were utilized to continuously maintain the L1 trigger output rate within the 100 kHz bandwidth limit. When the luminosity ramp-up phase stabilized in 2011 and 2012, the strategy focused on reducing the number of L1 menus being developed to a few per year, and adapting for different machine operational conditions by using multiple prescale columns rather than different L1 menus. At the end of 2012, during a twelve-hour-long fill, the instantaneous luminosity delivered by LHC varied significantly spanning from $\approx 7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ to $\approx 2.5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. The average number of pileup events per interaction ranged from ≈ 30 at the beginning to ≈ 12 at the end of the run.

To aid the L1 menu development using data, a special reduced-content event data format (containing only GCT, GMT and GT readout payloads) was defined and used to record events in a special data set. These events were collected on the basis of BPTX and L1 trigger GT decision only. Hence, with such recorded zero bias and L1 bias data sets, it was possible to properly account for rate overlaps of the algorithms operated in parallel in the GT (section 2.4) while designing new menus. Additionally, since the event size was significantly smaller than the standard event sizes [3], it was possible to collect a much higher trigger rate of these events than the standard event-data payload, enabling frequent offline analysis and cross-checks of the L1 trigger decision.

5.1.1 Menu development

The L1 menu development for the first LHC run was to a large extent based on data. Data recorded during standard collision runs and from special LHC setups including high pileup runs. To better understand the features of the LHC machine, different magnet and collimator settings were used. In addition, some data were taken with very few proton bunches. Large number of protons per bunch lead to significantly more collisions per bunch crossing, resulting in high-pileup events. These events were used to project trigger rates at improved LHC performance. Simulated data samples were also used to evaluate the impact of the 7 TeV to 8 TeV LHC energy increase in 2012.

For the L1 menu development, as well as for the development of the L1 trigger algorithms, we followed the following principles and strategy:

- use single-object triggers as baseline algorithms and adjust thresholds to be sensitive to the electroweak physics as well as new physics, e.g., heavy particles, multi-object final states, events with large missing transverse energy;
- in case the thresholds of the single-object triggers are too high with respect to the given physics goals (or if the acceptance for a given signal can be largely increased), use multi-object triggers, e.g., two muons or one muon plus two jets;

Table 15. Machine operational conditions, target instantaneous luminosity used for rate estimation, and approximate overall L1 rate for three sample L1 menus, representative of the end of the year data-taking conditions for 2010, 2011, and 2012.

Year	$\sqrt{s}$ [TeV]	Ref. $\mathcal{L}$ [$\text{cm}^{-2} \text{s}^{-1}$]	$\langle \text{pileup} \rangle$	$\langle \text{L1 rate} \rangle$ [kHz]
2010	7	0.15×10^{33}	≈ 2.5	56.9
2011	7	3.00×10^{33}	≈ 14	80.9
2012	8	5.00×10^{33}	≈ 23	56.5




Figure 71. Rates (left) and cross sections (right) for a significant sample of L1 e/γ triggers from 2010, 2011, and 2012 sample menus.

- prefer algorithms which are insensitive to changing LHC run conditions, e.g., prefer algorithms that are less sensitive to pileup events; and
- the algorithms and thresholds in a new L1 menu developed, e.g., for a different instantaneous luminosity, should result, if possible, in a similar sharing of rates for the same type of triggers: i.e., the muon triggers, e/γ and jet/sum triggers should have the same rate at a different instantaneous luminosity compared to the existing L1 menu.

Table 15 gives an overview of typical output rates of the L1 trigger system in 2010, 2011, and 2012, and table 16 shows details for a typical 2012 menu. The examples are chosen for LHC run periods where the measured instantaneous luminosities were close to the ones the different menus were designed for. The overall L1 trigger output rate was significantly higher than 50 kHz and well below the 100 kHz limit, as intended. The differences of observed and predicted total trigger rates largely depended on how the L1 trigger was operated, i.e., if a prescale column was changed at instantaneous luminosities different from the desired operating instantaneous luminosity of a specific L1 menu it followed that the total trigger output rate changed significantly ($\mathcal{O}(10 \text{ kHz})$).

The average L1 total trigger output rate varied from year to year due to adaptations to the changing LHC conditions. Figures 71, 72, and 73 show trigger rates and cross sections of the




Figure 72. Rates (left) and cross sections (right) for a significant sample of L1 muon triggers from 2010, 2011, and 2012 sample menus.




Figure 73. Rates (top) and cross sections (bottom) for a significant sample of L1 jet triggers from 2010, 2011, and 2012 sample menus.

various lowest threshold, unscaled, single-object, and multi-object triggers defined for the first LHC run. It was found that for almost all given triggers in a specific menu, the rates and cross sections had to be similar or lower compared to an earlier used trigger menu. This was required to maintain the overall L1 trigger output rate of below the 100 kHz limit taking into account the increasing LHC performance. To achieve this goal for higher instantaneous luminosities, i.e., later in 2011 and 2012, multi-object trigger algorithms as well as higher object thresholds were used.

5.2 HLT menus

The configuration of all the HLT paths that are run online at one time is called the HLT menu of CMS. This menu was initially prepared, based on simulated data, before the first data was taken in 2010, and has continuously evolved since then. This evolution is driven mainly by the changes in the machine conditions, namely $\sqrt{s}$, luminosity, bunch spacing, and pileup conditions. Moreover, timing improvements in the software-based algorithms and analysis techniques allowed the online algorithms to be brought much closer to the ones adopted offline, leading to better performance, as well as closer correspondence between the online and the offline selections. In addition to the trigger paths designed to preselect the events to be used in the physics analyses, calibration and monitoring paths for the different CMS subdetectors are also necessary and were included in all menus.

The first menus in 2010 consisted of fewer than 60 separate trigger paths. The low instantaneous luminosity supplied by the LHC at that time allowed the use of several “pass-through” paths, in which the events accepted by the L1 trigger are accepted also by the HLT without further requirements and restrictions. In addition to the pass-through paths, single “physics object” triggers started to be developed, meant to trigger on inclusive isolated or non-isolated electrons, photons, muons, and jets. As the instantaneous luminosity increased, the strategies used to control the trigger rate consisted of: raising p_T thresholds, adding isolation and quality conditions in the identification of jets, leptons, and photons, using prescales, introducing cross-triggers (triggers which require several physics objects of different types), and defining dedicated τ -like jets. Moreover, a few other paths were included in the menu to study possible implementations for future data-taking periods at higher rate. During 2010, 12 different trigger menus were developed, covering the wide range of instantaneous luminosity scenarios provided by the LHC (from 1×10^{27} to $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$). In addition, prescale values were designed according to the LHC luminosity.

In 2011, with the LHC still operating at 7 TeV center-of-mass energy, six different HLT and L1 menus were designed, aimed at instantaneous luminosities ranging from 5×10^{32} to $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. Tighter selections were therefore needed, and the refinement of the trigger requirements was achieved by gradually introducing analysis-like selection criteria at the trigger level. Besides the usual “physics object”-oriented paths, the presence of cross-triggers and more complex trigger paths, based on algorithms similar to those applied in the offline analyses, became more and more relevant in the menu. A few refined techniques used offline could therefore be brought to the HLT, after adapting them to reduce the CPU time needed, at the expense of very little performance, and without greatly compromising their final response. Amongst those techniques, particle flow reconstruction [41] was used since the beginning to characterize the hadronically decaying τ leptons at the HLT. Towards the end of the 2011 data-taking period, strategies designed to mitigate the effect of pileup were also included in several trigger paths, with the intent of studying their performance in view of the 2012 data taking when the pileup effect was expected to become more relevant. In

Table 16. Rates from a significant set of unscaled algorithms participating to a typical L1 menu used during 2012 data-taking. Rates and cross sections (σ) are computed for a target luminosity of $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. The overall menu rate (including calibration and monitoring triggers) is 56.5 kHz. The corresponding average pileup is approximately 23 interactions per bunch crossing.

Seed name	rate @ $5 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ [kHz]	σ [μb]
L1_SingleIsoEG18er	7.69	1.55
L1_SingleEG20	10.5	2.14
L1_SingleMu12er	8.11	1.64
L1_SingleMu16	7.49	1.51
L1_SingleJet128	1.15	0.232
L1_SingleMu6_NotBptxOR	0.03	0.007
L1_SingleJetC32_NotBptxOR	0.13	0.026
L1_ETM36	4.35	0.881
L1_HTT150	1.10	0.223
L1_ETT300	0.21	0.043
L1_DoubleEG_13_7	6.58	1.33
L1_DoubleMu_10_Open	4.36	0.882
L1_DoubleMu0er_HighQ	5.77	1.16
L1_DoubleJetC56	7.59	1.53
L1_DoubleTauJet44er	1.88	0.381
L1_TripleMu0	0.81	0.165
L1_TripleEG_12_7_5	2.19	0.444
L1_TripleEG7	1.35	0.273
L1_TripleJet_64_48_28_VBF	2.28	0.462
L1_QuadJetC36	0.74	0.150
L1_Mu3p5_EG12	2.34	0.474
L1_Mu12_EG7	1.03	0.208
L1_Mu0_HTT100	0.46	0.094
L1_Mu7er_ETM20	1.19	0.241
L1_IsoEG12er_ETM30	1.54	0.311
L1_EG22_ForJet24	2.42	0.489
L1_DoubleMu5_EG5	0.54	0.109
L1_Mu5_DoubleEG6	0.96	0.194
L1_DoubleEG6_HTT100	1.32	0.266
L1_DoubleJetC36_ETM30	3.40	0.688
L1_Mu10er_JetC12_WdEtaPhi1_DoubleJetC_20_12	1.02	0.207

this respect the so-called FASTJET corrections [53], offset corrections which take into account the average energy density in the event and the area of each jet in order to correct its energy on a jet-by-jet basis, proved very successful.

The number of paths deployed in the 2011 menus rose from about 310 at the beginning of the data taking to approximately 430 towards the end of the year. A few paths were included specifically to monitor and calibrate CMS subdetector components. For example, the response of the electromagnetic calorimeter, which is fundamental for the selection and analysis of the Higgs boson decaying in two photons, is continuously monitored by some dedicated paths, in order to provide updates to the calibrations in a timely manner.

When the 8 TeV run began in 2012, because of the higher instantaneous luminosity achieved by the LHC, pileup effects became much more important and therefore improvements in the design of most of the paths included in the menu were required. Ideally, the acceptance rate of a trigger should be proportional to the instantaneous luminosity, however, due to pileup it may increase non-linearly. This effect, together with the higher LHC luminosity, would give rise to unacceptably high trigger rates. The rate increases can only be mitigated by either raising the acceptance thresholds in the path themselves (with the unwanted effect of reducing the physics reach of the events selected), or by improving the performance of the selections, with sharper turn-on curves at the thresholds and less sensitivity to pileup. The main handle used to achieve this goal, without affecting the CMS physics potential, was the extension of the implementation of particle flow reconstruction to most jet- and E_T^{miss} -based triggers. The replacement of calorimeter-based jet triggers with PF-based ones was introduced gradually during the year. An additional advantage of using the PF in the trigger selection is that the selection algorithms are mostly the same as those used offline for the final analysis; however, reconstruction algorithms based on PF methods eventually do need more time compared to more “classical” one-object-per-detector-type algorithms. One idea used in the HLT to reduce the overall CPU time consumption was to move the PF reconstruction after all other possible selections, which were based on more classical quantities, which are faster to calculate. Among the other technical improvements to the HLT algorithms that allowed the rate to be kept low and the CPU time manageable, the following were particularly relevant: the optimization of lepton identification and isolation; the use of a filter to select leptons coming from the same vertex in several dilepton paths; weekly updates to ECAL transparency corrections, which allows efficiently compensating for the transparency loss in the endcap region affecting the electromagnetic energy scale; and the dedicated τ lepton reconstruction for the double- τ and lepton+ τ triggers.

Different menus were used in 2012 for four different LHC peak luminosities, ranging from 5×10^{33} to $8 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$. The number of different HLT paths during 2012 was approximately 400 at the beginning, increasing to about 450 by the end of the year.

In addition to the proton-proton triggers, dedicated menus for the heavy ion (lead-lead) collisions in 2010 and 2011 and for the proton-lead collisions in the first months of 2013 were created. The different running conditions and physics requirements led to different menus for the ion-ion and the proton-ion runs. The final heavy ion menus in 2010 and 2011 consisted of 58 and 77 HLT paths, respectively, while the proton-ion menu of 2013 contains about 150 paths.

6 Trigger system operation and evolution

6.1 Trigger monitoring and operations

During data taking the angular distributions of objects satisfying the trigger and the trigger rates were monitored. As these two kinds of information are produced using two different software tools, they provide complementary information about the behavior of the trigger system that are useful in diagnosing problems.

We use the central CMS data quality monitoring (DQM) tools [81] to monitor the angular distributions. The DQM tools process a small subset of events selected by the HLT and produce plots of η and ϕ of the trigger objects. The distributions are monitored for regions with abnormal appearances of either too many or too few events.

The rates of each HLT path are monitored in each node of the HLT computing cluster, where the CMS data acquisition software records how many times each trigger path was successful. The path information is summed over all nodes to give the total rate of each path. The summation occurs at every fixed integrated luminosity, and the results are written into a database. The HLT group has developed customized software to extract the rates from the database and compare them to expectations. The expected rate behavior is defined by fitting the trigger rates as a function of instantaneous luminosity using previously recorded data certified as good. Uncertainties from the fit provide an envelope of expected rate variations which sets the threshold for displaying warnings in the control room. A selected set of approximately 20 different representative triggers, out of the 400 of the HLT menu, is used for regular online monitoring. The selected HLT paths have either a large rate or an important physics signature. For instance, some of the closely monitored triggers includes single-muon, single-electron, and diphoton triggers, where rate variations are identified with a 5–10% sensitivity.

6.2 Technical performance

6.2.1 The L1 trigger deadline, downtime and reliability

“Deadline during active beam” is defined as the percentage of time during normal data taking (data acquisition system in “run” mode) when collisions occur but CMS is not ready to record triggers. There are several contributions to this deadline:

- Partition controller deadline: it arises when any CMS subsystem (such as a subdetector or trigger subsystem) asserts “not ready” because of a transient problem (e.g., “out of sync”, requiring a “resync” command) or because the instantaneous trigger rate is too high.
- Trigger rules deadline: a set of trigger rules of the type “not more than m triggers within n bunch crossings” limits the instantaneous trigger rate.
- Calibration deadline: at a rate of 100 Hz, calibration triggers are sent (required mainly by the electromagnetic calorimeter) and a small part of the orbit is blocked for this purpose.

Usually, deadline was kept to approximately 1%, only a small fraction of which was due to trigger subsystems.




Figure 74. The average CPU busy fraction as a function of instantaneous luminosity for one LHC fill. Luminosity sections with data-taking downtime $>40\%$ are removed.

“Downtime” is the percentage of time when data acquisition system cannot be put into run mode during active beams because of a malfunctioning subsystem. During regular running, the downtime due to trigger subsystems was well fewer than one percent. In most cases, the trigger downtime caused by hardware or software crashes, which could be fixed by restarting the electronics subsystem or the software process, respectively. To take care of the rare cases where an electronic module is faulty, spare modules for all systems are kept in the electronics cavern. For the GT, a fully equipped spare crate is kept running and ready to take over at any time in case of a hardware fault. Empirically, we observe that the L1 trigger system contributes only a small fraction to the total experiment downtime.

6.2.2 The HLT resources and optimization

As described in section 2.6, the HLT runs on an EVF consisting of three different types of machines. Two complementary methods are used to monitor the usage of this farm by the HLT menu. The first method directly measures the time taken by the HLT selection and reconstruction steps for each event during data taking. The second method rapidly samples every CPU in the farm to determine its state, and the time per event is calculated based on the frequency of finding the CPU in a non-idle state. The two methods give consistent results. Using the second method, the total busy fraction of the EVF can also be determined.

To estimate the CPU usage of an HLT menu at a future (higher) instantaneous luminosity value, the average busy fraction over the course of an LHC fill is measured, and a fit is performed as a function of instantaneous luminosity, as shown in figure 74. An exponential function is found to give a good description of the data over a wide range of instantaneous luminosity and allows




Figure 75. The HLT processing time per event as a function of instantaneous luminosity for the three different machine types used in the filter farm.

extrapolation to higher luminosities. In addition, we also measure the time per event for each type of machine used in the filter farm as shown in figure 75. The time per event is observed to be approximately linear as a function of luminosity on the Intel Xeon E5430 CPUs. The other two types of CPUs employ Intel’s hyper-threading to run twice as many concurrent processes as there are physical cores by using parts of the CPU that would otherwise be idle. As a result, the time per event for these hyper-threaded CPUs increases faster than linearly as the CPU is saturated with increasing luminosity and input rate. Using this information, it is possible to calculate the maximum time per event of the HLT menu for a given L1 input rate, and also the instantaneous luminosity at which this limit would be reached. The figure of merit used is the time per event for an Intel Xeon E5430 CPU. The filter farm configuration used during 2012 data taking was able to sustain an average processing time per event of approximately 200 ms for an L1 input rate of 100 kHz.

In addition to the online monitoring of the HLT menus, each menu is validated in an offline environment before being used for online data taking. Each new version of the menu is compared to a previous version on a single machine to ensure that the CPU consumption does not exceed expectations. The menus are tested by running the HLT once with each menu over the same sample of previously collected events. The measurement is done using a machine with similar core architecture to the Intel Xeon E5430 CPU, and is performed using the direct timing measurement described above. New instantaneous luminosity and L1 input rate limits can then be determined by using the relative performance of the new menu and the measured performance of the older menu. An example of an offline comparison of the times per event for two different HLT menus is shown in figure 76. When testing a new menu, the time per event for each HLT path is also checked to determine which paths are the most CPU intensive. The algorithms for CPU-intensive paths are then optimized to ensure that the total processing time does not exceed the limitations of the system.




Figure 76. Comparison of the time per event measured for two different HLT menus using a validation machine outside of the event filter farm.

6.2.3 The HLT operations

Following offline validation, HLT menus are validated in an online environment using the HLT online (“HiLTON”) test stand. In order to be as close as possible to the online environment, the HiLTON is operated using the same run control interface as the CMS DAQ system. The HiLTON hardware consists of 30 Dell PE 1950 machines with dual quad-core 2.66 GHz CPUs and 16 GB of RAM. The test stand system is subdivided into three groups of ten machines. The first group is always kept identical to the PE 1950 machines used for data taking operations and thus can be used to validate menus on the current online software environment. The second group of machines may additionally be used to validate the performance of software updates, and HLT menus that depend on the updates, in an online environment. Finally, the third subdivision of the HiLTON is used to evaluate changes made to the HiLTON itself. Two machines of each group are dedicated to the building of new online software releases. A third machine is always used to collect the output of the HLT and save events to disk. The remaining seven machines in each group are able to process events via seven instances of the HLT per machine, although only four machines are used for typical menu validation.

The HLT validation is designed to maximize the performance and stability of HLT algorithms. As every event satisfying L1 trigger requirements is examined by the HLT, and several HLT decisions are based on analysis-quality physics objects reconstructed using information from all CMS subdetectors, HLT reliability is critical to the success of the experiment. On rare occasions, typically below a few events per month during data taking, one or more HLT algorithms will experience a processing error while examining a collision event. These events are stored locally for later analysis and are used to improve the reliability of the HLT software.

Roughly 0.5% of the downtime during collision data taking operations from 2009 until 2012 (including all proton-proton and heavy ion collision operations at any center-of-mass energy) was due to problems with the HLT; 95% of this downtime was due to a single incident when corrupt detector input resulted in the HLT failure for every incoming collision event. Prior and subsequent data taking using the same HLT menu resulted in no loss of data. Ignoring this incident, the HLT was responsible for a negligible loss of collision data.

7 Summary

The CMS trigger system consists of two levels: an L1 custom hardware trigger, and an HLT system with custom c++ software routines running on commodity CPUs.

The L1 trigger takes input from the calorimeters and the muon system to select the events of physics interest. To do this, it uses identified leptons, photons, and jet candidates, as well as event-level information such as missing transverse energy. Trigger primitives are generated on the front-ends of the subdetectors and then processed in several steps before a final decision is rendered in the global trigger.

The L1 calorimeter trigger comprises two stages, a regional calorimeter trigger (RCT) and a global calorimeter trigger (GCT). The RCT processes the regional information in parallel and sends as output e/γ candidates and regional E_T sums. The GCT sorts the e/γ candidates further, finds jets (classified as central, forward, and tau) using the E_T sums, and calculates global quantities such as E_T^{miss} .

Each of the three muon detector systems in CMS participates in the L1 muon trigger to ensure good coverage and redundancy. For the DT and CSC systems, the front-end trigger electronics identifies track segments from the hit information registered in multiple detector planes. Track finders apply pattern recognition algorithms which identify muon candidates, and measure their momenta from the amount of their bending in the magnetic field of the return yoke between measurement locations. The RPC hits are directly sent from the front-end electronics to pattern-comparator logic boards which identify muon candidates. The global muon trigger merges muon candidates, applies quality criteria, and sends the muon candidates to the global trigger.

The global trigger implements the menu of selection requirements applied to all objects. A maximum of 128 separate selections can be implemented simultaneously. Overall, the L1 decision is rendered within $4\,\mu\text{s}$ after the collision. At most 100 kHz of events are sent to the HLT for processing.

The HLT is implemented in software, and further refines the purity of the physics objects. Events are selected for offline storage at an average rate of 400 Hz. The HLT event selection is performed in a similar way to that used in the offline processing. For each event, objects such as leptons, photons, and jets are reconstructed and identification criteria are applied in order to select only those events which are of possible interest for data analysis. The HLT hardware consists of a processor farm using commodity PCs running Scientific Linux. The subunits are called builder and filter units. In the builder units, event fragments are assembled to complete events. Filter units then unpack the raw data and perform event reconstruction and trigger filtering. Both the L1 triggers and HLT include prescaling of events passing defined selection criteria.

The performance of the CMS trigger system has been evaluated in two stages. First, the performance of the L1 and HLT systems has been evaluated for individual trigger objects such as electrons, muons, photons, or jets, using tag-and-probe techniques. Most of the measurements considered come from the 2012 CMS data set, where data have been collected at $\sqrt{s} = 8$ TeV. Performance has been evaluated in terms of efficiency with respect to offline quantities and to the appropriate trigger rate. Both L1 and HLT performance have been studied, showing the high selection efficiency of the CMS trigger system. Second, the performance of the trigger system has been demonstrated by considering key examples across different physics analyses. In CMS, the HLT decisions often are derived from complex correlated combinations of single objects such as electrons, muons, or τ leptons. The broad range of capabilities of the trigger system has been shown through examples in Higgs boson, top-quark, and B physics, as well as in searches for new physics.

The trigger system has been instrumental in the successful collection of data for physics analyses in Run 1 of the CMS experiment at the LHC. Efficiencies were measured in data and compared to simulation, and shown to be high and well-understood. Many physics signals were collected with high efficiency and flexibility under rapidly-changing conditions, enabling a diverse and rich physics program, which has led to hundreds of publications based on the Run 1 data samples.

Acknowledgments

We congratulate our colleagues in the CERN accelerator departments for the excellent performance of the LHC and thank the technical and administrative staffs at CERN and at other CMS institutes for their contributions to the success of the CMS effort. In addition, we gratefully acknowledge the computing centers and personnel of the Worldwide LHC Computing Grid for delivering so effectively the computing infrastructure essential to our analyses. Finally, we acknowledge the enduring support for the construction and operation of the LHC and the CMS detector provided by the following funding agencies: the Austrian Federal Ministry of Science, Research and Economy and the Austrian Science Fund; the Belgian Fonds de la Recherche Scientifique, and Fonds voor Wetenschappelijk Onderzoek; the Brazilian Funding Agencies (CNPq, CAPES, FAPERJ, and FAPESP); the Bulgarian Ministry of Education and Science; CERN; the Chinese Academy of Sciences, Ministry of Science and Technology, and National Natural Science Foundation of China; the Colombian Funding Agency (COLCIENCIAS); the Croatian Ministry of Science, Education and Sport, and the Croatian Science Foundation; the Research Promotion Foundation, Cyprus; the Secretariat for Higher Education, Science, Technology and Innovation, Ecuador; the Ministry of Education and Research, Estonian Research Council via IUT23-4 and IUT23-6 and European Regional Development Fund, Estonia; the Academy of Finland, Finnish Ministry of Education and Culture, and Helsinki Institute of Physics; the Institut National de Physique Nucléaire et de Physique des Particules / CNRS, and Commissariat à l'Énergie Atomique et aux Énergies Alternatives / CEA, France; the Bundesministerium für Bildung und Forschung, Deutsche Forschungsgemeinschaft, and Helmholtz-Gemeinschaft Deutscher Forschungszentren, Germany; the General Secretariat for Research and Technology, Greece; the National Scientific Research Foundation, and National Innovation Office, Hungary; the Department of Atomic Energy and the Department of Science and Technology, India; the Institute for Studies in Theoretical Physics and Mathematics, Iran; the Science Foundation, Ireland; the Istituto Nazionale di Fisica Nucleare, Italy; the Ministry of Science, ICT and Future Planning, and National Research Foundation (NRF), Republic of Korea; the

Lithuanian Academy of Sciences; the Ministry of Education, and University of Malaya (Malaysia); the Mexican Funding Agencies (BUAP, CINVESTAV, CONACYT, LNS, SEP, and UASLP-FAI); the Ministry of Business, Innovation and Employment, New Zealand; the Pakistan Atomic Energy Commission; the Ministry of Science and Higher Education and the National Science Centre, Poland; the Fundação para a Ciência e a Tecnologia, Portugal; JINR, Dubna; the Ministry of Education and Science of the Russian Federation, the Federal Agency of Atomic Energy of the Russian Federation, Russian Academy of Sciences, and the Russian Foundation for Basic Research; the Ministry of Education, Science and Technological Development of Serbia; the Secretaría de Estado de Investigación, Desarrollo e Innovación and Programa Consolider-Ingenio 2010, Spain; the Swiss Funding Agencies (ETH Board, ETH Zurich, PSI, SNF, UniZH, Canton Zurich, and SER); the Ministry of Science and Technology, Taipei; the Thailand Center of Excellence in Physics, the Institute for the Promotion of Teaching Science and Technology of Thailand, Special Task Force for Activating Research and the National Science and Technology Development Agency of Thailand; the Scientific and Technical Research Council of Turkey, and Turkish Atomic Energy Authority; the National Academy of Sciences of Ukraine, and State Fund for Fundamental Researches, Ukraine; the Science and Technology Facilities Council, U.K.; the US Department of Energy, and the US National Science Foundation.

Individuals have received support from the Marie-Curie program and the European Research Council and EPLANET (European Union); the Leventis Foundation; the A. P. Sloan Foundation; the Alexander von Humboldt Foundation; the Belgian Federal Science Policy Office; the Fonds pour la Formation à la Recherche dans l'Industrie et dans l'Agriculture (FRIA-Belgium); the Agentschap voor Innovatie door Wetenschap en Technologie (IWT-Belgium); the Ministry of Education, Youth and Sports (MEYS) of the Czech Republic; the Council of Science and Industrial Research, India; the HOMING PLUS program of the Foundation for Polish Science, cofinanced from European Union, Regional Development Fund, the Mobility Plus program of the Ministry of Science and Higher Education, the National Science Center (Poland), contracts Harmonia 2014/14/M/ST2/00428, Opus 2013/11/B/ST2/04202, 2014/13/B/ST2/02543 and 2014/15/B/ST2/03998, Sonata-bis 2012/07/E/ST2/01406; the Thalís and Aristeia programs cofinanced by EU-ESF and the Greek NSRF; the National Priorities Research Program by Qatar National Research Fund; the Programa Clarín-COFUND del Principado de Asturias; the Rachadapisek Sompot Fund for Postdoctoral Fellowship, Chulalongkorn University and the Chulalongkorn Academic into Its 2nd Century Project Advancement Project (Thailand); and the Welch Foundation, contract C-1845.

References

- [1] CMS collaboration, *The CMS experiment at the CERN LHC*, [2008 JINST 3 S08004](#).
- [2] L. Evans and P. Bryant, *LHC machine*, [2008 JINST 3 S08001](#).
- [3] CMS collaboration, *CMS TriDAS project — technical design report, volume 2: data acquisition and high-level trigger*, [CERN-LHCC-2002-026](#), CERN, Geneva Switzerland (2002).
- [4] ATLAS collaboration, *Performance of the ATLAS trigger system in 2010*, [Eur. Phys. J. C 72 \(2012\) 1849](#) [[arXiv:1110.1530](#)].
- [5] CMS collaboration, *CMS TriDAS project — technical design report, volume 1: the trigger systems*, [CERN-LHCC-2000-038](#), CERN, Geneva Switzerland (2000).

- [6] M. Anfreville et al., *Laser monitoring system for the CMS lead tungstate crystal calorimeter*, *Nucl. Instrum. Meth. A* **594** (2008) 292.
- [7] CMS collaboration, *Energy calibration and resolution of the CMS electromagnetic calorimeter in pp collisions at $\sqrt{s} = 7$ TeV*, *2013 JINST* **8** P09009 [[arXiv:1306.2016](#)].
- [8] P. Klabbbers et al., *Integration of the CMS regional calorimeter trigger hardware into the CMS level-1 trigger*, in *Topical Workshop on Electronics for Particle Physics (TWEPP 2007)*, [CERN-2007-007](#), Prague Czech Republic September 3–7 2007.
- [9] P. Klabbbers et al., *Operation and monitoring of the CMS regional calorimeter trigger hardware*, in *Topical Workshop on Electronics for Particle Physics (TWEPP 2008)*, [CERN-2008-008](#), Naxos Greece September 15–19 2008.
- [10] P. Klabbbers et al., *Performance of the CMS regional calorimeter trigger*, in *Topical Workshop on Electronics for Particle Physics (TWEPP 2009)*, [CERN-2009-006](#), Paris France September 21–25 2009.
- [11] W.H. Smith et al., *CMS regional calorimeter trigger high speed ASICs*, in *Proceedings of the Sixth Workshop on Electronics for LHC Experiments*, [CERN-2000-010](#), Krakow Poland September 11–15 2000.
- [12] M. Stettler et al., *The CMS global calorimeter trigger hardware design*, in *Proceedings for the 12th Workshop on Electronics for LHC and Future Experiments*, [CERN-2007-001](#), Valencia Spain September 25–29 2006.
- [13] G. Iles et al., *Revised CMS global calorimeter trigger functionality & algorithms*, in *Proceedings for the 12th Workshop on Electronics for LHC and Future Experiments*, [CERN-2007-001](#), Valencia Spain September 25–29 2006.
- [14] C. Foudas et al., *First results on the performance of the CMS global calorimeter trigger*, in *Topical Workshop on Electronics for Particle Physics (TWEPP 2007)*, [CERN-2007-007](#), Prague Czech Republic September 3–7 2007.
- [15] G. Iles et al., *Performance and lessons of the CMS global calorimeter trigger*, in *Topical Workshop on Electronics for Particle Physics (TWEPP 2008)*, [CERN-2008-008](#), Naxos Greece September 15–19 2008.
- [16] A. Tapper et al., *Commissioning and performance of the CMS global calorimeter trigger*, *IEEE Nucl. Sci. Symp. Conf. Rec.* (2008) 1871.
- [17] J. Brooke et al., *Performance of the CMS global calorimeter trigger*, in *Proceedings for the 1st International Conference on Technology and Instrumentation in Particle Physics*, Tsukuba Japan (2009) [*Nucl. Instrum. Meth. A* **623** (2010) 546].
- [18] J.J. Brooke et al., *The design of a flexible global calorimeter trigger system for the compact muon solenoid experiment*, *2007 JINST* **2** P10002.
- [19] M. De Giorgi et al., *Test results of the ASIC front end trigger prototypes for the muon barrel detector of CMS at LHC*, *Nucl. Instrum. Meth. A* **438** (1999) 302.
- [20] CMS collaboration, *Design of the track correlator for the DTBX trigger*, [CMS-NOTE-1999-007](#), CERN, Geneva Switzerland (1999).
- [21] CMS collaboration, *Antifuse-FPGAs for the track-sorter-master of the CMS muon barrel drift tubes: design issues and irradiation test*, [CMS-NOTE-2001-028](#), CERN, Geneva Switzerland (2001).

- [22] P. Arce et al., *Bunched beam test of the CMS drift tubes local muon trigger*, [*Nucl. Instrum. Meth. A* **534** \(2004\) 441](#).
- [23] M. Aldaya et al., *Fine synchronization of the muon drift tubes local trigger*, [*Nucl. Instrum. Meth. A* **564** \(2006\) 169](#).
- [24] CMS collaboration, *Fine synchronization of the CMS muon drift-tube local trigger using cosmic rays*, [2010 JINST 5 T03004](#) [[arXiv:0911.4904](#)].
- [25] L. Guiducci et al., *DT sector collector electronics design and construction*, in *Topical Workshop on Electronics for Particle Physics (TWEPP 2007)*, [CERN-2007-007](#), Prague Czech Republic September 3–7 2007.
- [26] CMS collaboration, *The performance of the CMS muon detector in proton-proton collisions at $\sqrt{s} = 7$ TeV at the LHC*, [2013 JINST 8 P11002](#) [[arXiv:1306.6905](#)].
- [27] J. Ero et al., *The CMS drift tube trigger track finder*, [2008 JINST 3 P08006](#).
- [28] L. Guiducci et al., *Design and test of the off-detector electronics for the CMS barrel muon trigger*, in *Proceedings for the 12th Workshop on Electronics for LHC and Future Experiments*, [CERN-2007-001](#), Valencia Spain September 25–29 2006.
- [29] D. Acosta et al., *Development and test of a prototype regional track-finder for the level-1 trigger of the cathode strip chamber muon system of CMS*, [*Nucl. Instrum. Meth. A* **496** \(2003\) 64](#).
- [30] CMS collaboration, *Performance of CMS muon reconstruction in pp collision events at $\sqrt{s} = 7$ TeV*, [2012 JINST 7 P10002](#) [[arXiv:1206.4071](#)].
- [31] RD5 collaboration, M. Andlinger et al., *Pattern Comparator Trigger (PACT) for the muon system of the CMS experiment*, [*Nucl. Instrum. Meth. A* **370** \(1996\) 389](#).
- [32] H. Sakulin and A. Taurok, *The level-1 global muon trigger for the CMS experiment*, in *Proceedings for the 9th Workshop on Electronics for LHC Experiments*, [CERN-2003-006](#), Amsterdam The Netherlands September 29–October 3 2003.
- [33] M. Jeitler et al., *The level-1 global trigger for the CMS experiment at LHC*, [2007 JINST 2 P01006](#).
- [34] A. Taurok et al., *The central trigger control system of the CMS experiment at CERN*, [CMS-NOTE-2010-017](#), CERN, Geneva Switzerland (2010).
- [35] J. Varela, *CMS L1 trigger control system*, [CMS-NOTE-2002-033](#), CERN, Geneva Switzerland (2002).
- [36] TOTEM collaboration, G. Anelli et al., *The TOTEM experiment at the CERN Large Hadron Collider*, [2008 JINST 3 S08007](#).
- [37] G. Bauer et al., *Recent experience and future evolution of the CMS high level trigger system*, in *18th IEEE-NPSS Real Time Conference (RT)*, (2012), pg. 1.
- [38] CMS collaboration, *Commissioning of the CMS high-level trigger with cosmic rays*, [2010 JINST 5 T03005](#) [[arXiv:0911.4889](#)].
- [39] CMS TRIGGER AND DATA ACQUISITION GROUP collaboration, *The CMS high level trigger*, [*Eur. Phys. J. C* **46** \(2006\) 605](#) [[hep-ex/0512077](#)].
- [40] CMS collaboration, *Data parking and data scouting at the CMS experiment*, [CMS-DP-2012-022](#), CERN, Geneva Switzerland (2012).
- [41] CMS collaboration, *Particle-flow event reconstruction in CMS and performance for jets, taus and MET*, [CMS-PAS-PFT-09-001](#), CERN, Geneva Switzerland (2009).
- [42] CMS collaboration, *Commissioning of the particle-flow event reconstruction with the first LHC collisions recorded in the CMS detector*, [CMS-PAS-PFT-10-001](#), CERN, Geneva Switzerland (2010).

- [43] CMS collaboration, *Description and performance of track and primary-vertex reconstruction with the CMS tracker*, [2014 JINST 9 P10009](#) [[arXiv:1405.6569](#)].
- [44] R. Fruhwirth, *Application of Kalman filtering to track and vertex fitting*, [Nucl. Instrum. Meth. A 262 \(1987\) 444](#).
- [45] CMS collaboration, *Electron reconstruction and identification at $\sqrt{s} = 7$ TeV*, [CMS-PAS-EGM-10-004](#), CERN, Geneva Switzerland (2010).
- [46] CMS collaboration, *Performance of electron reconstruction and selection with the CMS detector in proton-proton collisions at $\sqrt{s} = 8$ TeV*, [2015 JINST 10 P06005](#) [[arXiv:1502.02701](#)].
- [47] M.J. Oreglia, *A study of the reactions $\psi' \rightarrow \gamma\gamma\psi$* , see appendix D, Ph.D. thesis, SLAC-R-236, Stanford University, Stanford U.S.A. (1980).
- [48] CMS collaboration, *Measurement of the inclusive W and Z production cross sections in pp collisions at $\sqrt{s} = 7$ TeV*, [JHEP 10 \(2011\) 132](#) [[arXiv:1107.4789](#)].
- [49] W. Bialas and D.A. Petyt, *Mitigation of anomalous APD signals in the CMS ECAL*, [2013 JINST 8 C03020](#).
- [50] CMS collaboration, D.A. Petyt, *Anomalous APD signals in the CMS electromagnetic calorimeter*, [Nucl. Instrum. Meth. A 695 \(2012\) 293](#).
- [51] CMS collaboration, *Measurement of the $B_s^0 \rightarrow \mu^+ \mu^-$ branching fraction and search for $B^0 \rightarrow \mu^+ \mu^-$ with the CMS experiment*, [Phys. Rev. Lett. 111 \(2013\) 101804](#) [[arXiv:1307.5025](#)].
- [52] LHCb and CMS collaborations, *Observation of the rare $B_s^0 \rightarrow \mu^+ \mu^-$ decay from the combined analysis of CMS and LHCb data*, [Nature 522 \(2015\) 68](#) [[arXiv:1411.4413](#)].
- [53] M. Cacciari, G.P. Salam and G. Soyez, *FastJet user manual*, [Eur. Phys. J. C 72 \(2012\) 1896](#) [[arXiv:1111.6097](#)].
- [54] CMS collaboration, *Determination of jet energy calibration and transverse momentum resolution in CMS*, [2011 JINST 6 P11002](#) [[arXiv:1107.4277](#)].
- [55] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, [JHEP 04 \(2008\) 063](#) [[arXiv:0802.1189](#)].
- [56] CMS collaboration, *Calorimeter jet quality criteria for the first CMS collision data*, [CMS-PAS-JME-09-008](#), CERN, Geneva Switzerland (2009).
- [57] CMS collaboration, *Evidence for the 125 GeV Higgs boson decaying to a pair of τ leptons*, [JHEP 05 \(2014\) 104](#) [[arXiv:1401.5041](#)].
- [58] CMS collaboration, *Performance of τ -lepton reconstruction and identification in CMS*, [2012 JINST 7 P01001](#) [[arXiv:1109.6034](#)].
- [59] CMS collaboration, *Identification of b-quark jets with the CMS experiment*, [2013 JINST 8 P04013](#) [[arXiv:1211.4462](#)].
- [60] I.P. Lokhtin and A.M. Snigirev, *A model of jet quenching in ultrarelativistic heavy ion collisions and high- p_T hadron spectra at RHIC*, [Eur. Phys. J. C 45 \(2006\) 211](#) [[hep-ph/0506189](#)].
- [61] O. Kodolova, I. Vardanian, A. Nikitenko and A. Oulianov, *The performance of the jet identification and reconstruction in heavy ions collisions with CMS detector*, [Eur. Phys. J. C 50 \(2007\) 117](#).
- [62] CMS collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, [Phys. Lett. B 716 \(2012\) 30](#) [[arXiv:1207.7235](#)].

- [63] CMS collaboration, *Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV*, *JHEP* **06** (2013) 081 [[arXiv:1303.4571](#)].
- [64] CMS collaboration, *Measurement of the properties of a Higgs boson in the four-lepton final state*, *Phys. Rev. D* **89** (2014) 092007 [[arXiv:1312.5353](#)].
- [65] CMS collaboration, *Search for the Standard Model Higgs boson decaying to bottom quarks in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Lett. B* **710** (2012) 284 [[arXiv:1202.4195](#)].
- [66] CMS collaboration, *Measurement of the $t\bar{t}$ production cross section in pp collisions at $\sqrt{s} = 7$ TeV with lepton + jets final states*, *Phys. Lett. B* **720** (2013) 83 [[arXiv:1212.6682](#)].
- [67] CMS collaboration, *Measurement of the $t\bar{t}$ production cross section in the dilepton channel in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **02** (2014) 024 [Erratum *ibid.* **02** (2014) 102] [[arXiv:1312.7582](#)].
- [68] CMS collaboration, *Pileup removal algorithms*, CMS-PAS-JME-14-001, CERN, Geneva Switzerland (2014).
- [69] J. Alwall, P. Schuster and N. Toro, *Simplified models for a first characterization of new physics at the LHC*, *Phys. Rev. D* **79** (2009) 075020 [[arXiv:0810.3921](#)].
- [70] G.R. Farrar and P. Fayet, *Phenomenology of the production, decay and detection of new hadronic states associated with supersymmetry*, *Phys. Lett. B* **76** (1978) 575.
- [71] CMS collaboration, *Search for supersymmetry in hadronic final states with missing transverse energy using the variables α_T and b-quark multiplicity in pp collisions at $\sqrt{s} = 8$ TeV*, *Eur. Phys. J. C* **73** (2013) 2568 [[arXiv:1303.2985](#)].
- [72] CMS collaboration, *Inclusive search for squarks and gluinos in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Rev. D* **85** (2012) 012004 [[arXiv:1107.1279](#)].
- [73] CMS collaboration, *Inclusive search for supersymmetry using the razor variables in pp collisions at $\sqrt{s} = 7$ TeV*, *Phys. Rev. Lett.* **111** (2013) 081802 [[arXiv:1212.6961](#)].
- [74] S. Dimopoulos and G.L. Landsberg, *Black holes at the LHC*, *Phys. Rev. Lett.* **87** (2001) 161602 [[hep-ph/0106295](#)].
- [75] S.B. Giddings and S.D. Thomas, *High-energy colliders as black hole factories: the end of short distance physics*, *Phys. Rev. D* **65** (2002) 056010 [[hep-ph/0106219](#)].
- [76] S. Dimopoulos and R. Emparan, *String balls at the LHC and beyond*, *Phys. Lett. B* **526** (2002) 393 [[hep-ph/0108060](#)].
- [77] CMS collaboration, *Search for microscopic black holes in pp collisions at $\sqrt{s} = 8$ TeV*, *JHEP* **07** (2013) 178 [[arXiv:1303.5338](#)].
- [78] CMS collaboration, *Search for microscopic black hole signatures at the Large Hadron Collider*, *Phys. Lett. B* **697** (2011) 434 [[arXiv:1012.3375](#)].
- [79] CMS collaboration, *Search for microscopic black holes in pp collisions at $\sqrt{s} = 7$ TeV*, *JHEP* **04** (2012) 061 [[arXiv:1202.6396](#)].
- [80] D0 collaboration, V.M. Abazov et al., *Evidence for an anomalous like-sign dimuon charge asymmetry*, *Phys. Rev. Lett.* **105** (2010) 081801 [[arXiv:1007.0395](#)].
- [81] L. Tuura, A. Meyer, I. Segoni and G. Della Ricca, *CMS data quality monitoring: systems and experiences*, in *Proceedings of the 17th International Conference on Computing in High Energy and Nuclear Physics (CHEP09)*, Prague Czech Republic March 21–27 2009 [*J. Phys. Conf. Ser.* **219** (2010) 072020].

The CMS collaboration

Yerevan Physics Institute, Yerevan, Armenia

V. Khachatryan, A.M. Sirunyan, A. Tumasyan

Institut für Hochenergiephysik der OeAW, Wien, Austria

W. Adam, E. Asilar, T. Bergauer, J. Brandstetter, E. Brondolin, M. Dragicevic, J. Erö, M. Flechl, M. Friedl, R. Frühwirth¹, V.M. Ghete, C. Hartl, N. Hörmann, J. Hrubec, M. Jeitler¹, V. Knünz, A. König, M. Krammer¹, I. Krätschmer, D. Liko, T. Matsushita, I. Mikulec, D. Rabady², B. Rahbaran, H. Rohringer, J. Schieck¹, R. Schöfbeck, J. Strauss, W. Treberer-Treberspurg, W. Waltenberger, C.-E. Wulz¹

National Centre for Particle and High Energy Physics, Minsk, Belarus

V. Mossolov, N. Shumeiko, J. Suarez Gonzalez

Universiteit Antwerpen, Antwerpen, Belgium

S. Alderweireldt, T. Cornelis, E.A. De Wolf, X. Janssen, A. Knutsson, J. Lauwers, S. Luyckx, M. Van De Klundert, H. Van Haevermaet, P. Van Mechelen, N. Van Remortel, A. Van Spilbeeck

Vrije Universiteit Brussel, Brussel, Belgium

S. Abu Zeid, F. Blekman, J. D'Hondt, N. Daci, I. De Bruyn, K. Deroover, N. Heracleous, J. Keaveney, S. Lowette, L. Moreels, A. Olbrechts, Q. Python, D. Strom, S. Tavernier, W. Van Doninck, P. Van Mulders, G.P. Van Onsem, I. Van Parijs

Université Libre de Bruxelles, Bruxelles, Belgium

P. Barria, H. Brun, C. Caillol, B. Clerbaux, G. De Lentdecker, G. Fasanella, L. Favart, A. Grebenyuk, G. Karapostoli, T. Lenzi, A. Léonard, T. Maerschalk, A. Marinov, L. Perniè, A. Randle-conde, T. Reis, T. Seva, C. Vander Velde, P. Vanlaer, R. Yonamine, F. Zenoni, F. Zhang³

Ghent University, Ghent, Belgium

K. Beernaert, L. Benucci, A. Cimmino, S. Crucy, D. Dobur, A. Fagot, G. Garcia, M. Gul, J. McCartin, A.A. Ocampo Rios, D. Poyraz, D. Ryckbosch, S. Salva, M. Sigamani, N. Strobbe, M. Tytgat, W. Van Driessche, E. Yazgan, N. Zaganidis

Université Catholique de Louvain, Louvain-la-Neuve, Belgium

S. Basegmez, C. Beluffi⁴, O. Bondu, S. Brochet, G. Bruno, A. Caudron, L. Ceard, G.G. Da Silva, C. Delaere, D. Favart, L. Forthomme, A. Giammanco⁵, J. Hollar, A. Jafari, P. Jez, M. Komm, V. Lemaitre, A. Mertens, M. Musich, C. Nuttens, L. Perrini, A. Pin, K. Piotrkowski, A. Popov⁶, L. Quertenmont, M. Selvaggi, M. Vidal Marono

Université de Mons, Mons, Belgium

N. Beliy, G.H. Hammad

Centro Brasileiro de Pesquisas Fisicas, Rio de Janeiro, Brazil

W.L. Aldá Júnior, F.L. Alves, G.A. Alves, L. Brito, M. Correa Martins Junior, M. Hamer, C. Hensel, C. Mora Herrera, A. Moraes, M.E. Pol, P. Rebello Teles

Universidade do Estado do Rio de Janeiro, Rio de Janeiro, Brazil

E. Belchior Batista Das Chagas, W. Carvalho, J. Chinellato⁷, A. Custódio, E.M. Da Costa, D. De Jesus Damiao, C. De Oliveira Martins, S. Fonseca De Souza, L.M. Huertas Guativa, H. Malbouisson, D. Matos Figueiredo, L. Mundim, H. Nogima, W.L. Prado Da Silva, A. Santoro, A. Sznajder, E.J. Tonelli Manganote⁷, A. Vilela Pereira

Universidade Estadual Paulista ^a, Universidade Federal do ABC ^b, São Paulo, Brazil

S. Ahuja^a, C.A. Bernardes^b, A. De Souza Santos^b, S. Dogra^a, T.R. Fernandez Perez Tomei^a, E.M. Gregores^b, P.G. Mercadante^b, C.S. Moon^{a,8}, S.F. Novaes^a, Sandra S. Padula^a, D. Romero Abad, J.C. Ruiz Vargas

Institute for Nuclear Research and Nuclear Energy, Sofia, Bulgaria

A. Aleksandrov, R. Hadjiiska, P. Iaydjiev, M. Rodozov, S. Stoykova, G. Sultanov, M. Vutova

University of Sofia, Sofia, Bulgaria

A. Dimitrov, I. Glushkov, L. Litov, B. Pavlov, P. Petkov

Institute of High Energy Physics, Beijing, China

M. Ahmad, J.G. Bian, G.M. Chen, H.S. Chen, M. Chen, T. Cheng, R. Du, C.H. Jiang, R. Plestina⁹, F. Romeo, S.M. Shaheen, A. Spiezia, J. Tao, C. Wang, Z. Wang, H. Zhang

State Key Laboratory of Nuclear Physics and Technology, Peking University, Beijing, China

C. Asawatrangkuldee, Y. Ban, Q. Li, S. Liu, Y. Mao, S.J. Qian, D. Wang, Z. Xu

Universidad de Los Andes, Bogota, Colombia

C. Avila, A. Cabrera, L.F. Chaparro Sierra, C. Florez, J.P. Gomez, B. Gomez Moreno, J.C. Sanabria

University of Split, Faculty of Electrical Engineering, Mechanical Engineering and Naval Architecture, Split, Croatia

N. Godinovic, D. Lelas, I. Puljak, P.M. Ribeiro Cipriano

University of Split, Faculty of Science, Split, Croatia

Z. Antunovic, M. Kovac

Institute Rudjer Boskovic, Zagreb, Croatia

V. Brigljevic, K. Kadija, J. Luetic, S. Micanovic, L. Sudic

University of Cyprus, Nicosia, Cyprus

A. Attikis, G. Mavromanolakis, J. Mousa, C. Nicolaou, F. Ptochos, P.A. Razis, H. Rykaczewski

Charles University, Prague, Czech Republic

M. Bodlak, M. Finger¹⁰, M. Finger Jr.¹⁰

Academy of Scientific Research and Technology of the Arab Republic of Egypt, Egyptian Network of High Energy Physics, Cairo, Egypt

Y. Assran¹¹, M. El Sawy^{12,13}, S. Elgammal¹³, A. Ellithi Kamel¹⁴, M.A. Mahmoud¹⁵

National Institute of Chemical Physics and Biophysics, Tallinn, Estonia

B. Calpas, M. Kadastik, M. Murumaa, M. Raidal, A. Tiko, C. Veelken

Department of Physics, University of Helsinki, Helsinki, Finland

P. Eerola, J. Pekkanen, M. Voutilainen

Helsinki Institute of Physics, Helsinki, Finland

J. Härkönen, V. Karimäki, R. Kinnunen, T. Lampén, K. Lassila-Perini, S. Lehti, T. Lindén, P. Luukka, T. Mäenpää, T. Peltola, E. Tuominen, J. Tuominiemi, E. Tuovinen, L. Wendland

Lappeenranta University of Technology, Lappeenranta, Finland

J. Talvitie, T. Tuuva

IRFU, CEA, Université Paris-Saclay, Gif-sur-Yvette, France

M. Besancon, F. Couderc, M. Dejardin, D. Denegri, B. Fabbro, J.L. Faure, C. Favaro, F. Ferri, S. Ganjour, A. Givernaud, P. Gras, G. Hamel de Monchenault, P. Jarry, E. Locci, M. Machet, J. Malcles, J. Rander, A. Rosowsky, M. Titov, A. Zghiche

Laboratoire Leprince-Ringuet, Ecole Polytechnique, IN2P3-CNRS, Palaiseau, France

I. Antropov, S. Baffioni, F. Beaudette, P. Busson, L. Cadamuro, E. Chapon, C. Charlot, T. Dahms, O. Davignon, N. Filipovic, A. Florent, R. Granier de Cassagnac, S. Lisniak, L. Mastrolorenzo, P. Miné, I.N. Naranjo, M. Nguyen, C. Ochando, G. Ortona, P. Paganini, P. Pigard, S. Regnard, R. Salerno, J.B. Sauvan, Y. Sirois, T. Strebler, Y. Yilmaz, A. Zabi

Institut Pluridisciplinaire Hubert Curien, Université de Strasbourg, Université de Haute Alsace Mulhouse, CNRS/IN2P3, Strasbourg, France

J.-L. Agram¹⁶, J. Andrea, A. Aubin, D. Bloch, J.-M. Brom, M. Buttignol, E.C. Chabert, N. Chanon, C. Collard, E. Conte¹⁶, X. Coubez, J.-C. Fontaine¹⁶, D. Gelé, U. Goerlach, C. Goetzmann, A.-C. Le Bihan, J.A. Merlin², K. Skovpen, P. Van Hove

Centre de Calcul de l'Institut National de Physique Nucleaire et de Physique des Particules, CNRS/IN2P3, Villeurbanne, France

S. Gadrat

Université de Lyon, Université Claude Bernard Lyon 1, CNRS-IN2P3, Institut de Physique Nucléaire de Lyon, Villeurbanne, France

S. Beauceron, C. Bernet, G. Boudoul, E. Bouvier, C.A. Carrillo Montoya, R. Chierici, D. Contardo, B. Courbon, P. Depasse, H. El Mamouni, J. Fan, J. Fay, S. Gascon, M. Gouzevitch, B. Ille, F. Lagarde, I.B. Laktineh, M. Lethuillier, L. Mirabito, A.L. Pequegnot, S. Perries, J.D. Ruiz Alvarez, D. Sabes, L. Sgandurra, V. Sordini, M. Vander Donckt, P. Verdier, S. Viret

Georgian Technical University, Tbilisi, Georgia

T. Toriashvili¹⁷

Tbilisi State University, Tbilisi, Georgia

Z. Tsamalaidze¹⁰

RWTH Aachen University, I. Physikalisches Institut, Aachen, Germany

C. Autermann, S. Beranek, M. Edelhoff, L. Feld, A. Heister, M.K. Kiesel, K. Klein, M. Lipinski, A. Ostapchuk, M. Preuten, F. Raupach, S. Schael, J.F. Schulte, T. Verlage, H. Weber, B. Wittmer, V. Zhukov⁶

RWTH Aachen University, III. Physikalisches Institut A, Aachen, Germany

M. Ata, M. Brodski, E. Dietz-Laursonn, D. Duchardt, M. Endres, M. Erdmann, S. Erdweg, T. Esch, R. Fischer, A. Güth, T. Hebbeker, C. Heidemann, K. Hoepfner, D. Klingebiel, S. Knutzen, P. Kreuzer, M. Merschmeyer, A. Meyer, P. Millet, M. Olschewski, K. Padeken, P. Papacz, T. Pook, M. Radziej, H. Reithler, M. Rieger, F. Scheuch, L. Sonnenschein, D. Teyssier, S. Thüer

RWTH Aachen University, III. Physikalisches Institut B, Aachen, Germany

V. Cherepanov, Y. Erdogan, G. Flügge, H. Geenen, M. Geisler, F. Hoehle, B. Kargoll, T. Kress, Y. Kuessel, A. Künsken, J. Lingemann², A. Nehr Korn, A. Nowack, I.M. Nugent, C. Pistone, O. Pooth, A. Stahl

Deutsches Elektronen-Synchrotron, Hamburg, Germany

M. Aldaya Martin, I. Asin, N. Bartosik, O. Behnke, U. Behrens, A.J. Bell, K. Borras¹⁸, A. Burgmeier, A. Campbell, S. Choudhury¹⁹, F. Costanza, C. Diez Pardos, G. Dolinska, S. Dooling, T. Dorland, G. Eckerlin, D. Eckstein, T. Eichhorn, G. Flucke, E. Gallo²⁰, J. Garay Garcia, A. Geiser, A. Gizhko, P. Gunnellini, J. Hauk, M. Hempel²¹, H. Jung, A. Kalogeropoulos, O. Karacheban²¹, M. Kasemann, P. Katsas, J. Kieseler, C. Kleinwort, I. Korol, W. Lange, J. Leonard, K. Lipka, A. Lobanov, W. Lohmann²¹, R. Mankel, I. Marfin²¹, I.-A. Melzer-Pellmann, A.B. Meyer, G. Mittag, J. Mnich, A. Mussgiller, S. Naumann-Emme, A. Nayak, E. Ntomari, H. Perrey, D. Pitzl, R. Placakyte, A. Raspereza, B. Roland, M.Ö. Sahin, P. Saxena, T. Schoerner-Sadenius, M. Schröder, C. Seitz, S. Spannagel, K.D. Trippkewitz, R. Walsh, C. Wissing

University of Hamburg, Hamburg, Germany

V. Blobel, M. Centis Vignali, A.R. Draeger, J. Erfle, E. Garutti, K. Goebel, D. Gonzalez, M. Görner, J. Haller, M. Hoffmann, R.S. Höing, A. Junkes, R. Klanner, R. Kogler, N. Kovalchuk, T. Lapsien, T. Lenz, I. Marchesini, D. Marconi, M. Meyer, D. Nowatschin, J. Ott, F. Pantaleo², T. Peiffer, A. Perieanu, N. Pietsch, J. Poehlsen, D. Rathjens, C. Sander, C. Scharf, H. Schettler, P. Schleper, E. Schlieckau, A. Schmidt, J. Schwandt, V. Sola, H. Stadie, G. Steinbrück, H. Tholen, D. Troendle, E. Usai, L. Vanelderden, A. Vanhoefer, B. Vormwald

Institut für Experimentelle Kernphysik, Karlsruhe, Germany

M. Akbiyik, C. Barth, C. Baus, J. Berger, C. Böser, E. Butz, T. Chwalek, F. Colombo, W. De Boer, A. Descroix, A. Dierlamm, S. Fink, F. Frensch, R. Friese, M. Giffels, A. Gilbert, D. Haitz, F. Hartmann², S.M. Heindl, U. Husemann, I. Katkov⁶, A. Kornmayer², P. Lobelle Pardo, B. Maier, H. Mildner, M.U. Mozer, T. Müller, Th. Müller, M. Plagge, G. Quast, K. Rabbertz, S. Röcker, F. Roscher, G. Sieber, H.J. Simonis, F.M. Stober, R. Ulrich, J. Wagner-Kuhr, S. Wayand, M. Weber, T. Weiler, C. Wöhrmann, R. Wolf

Institute of Nuclear and Particle Physics (INPP), NCSR Demokritos, Aghia Paraskevi, Greece

G. Anagnostou, G. Daskalakis, T. Gerasis, V.A. Giakoumopoulou, A. Kyriakis, D. Loukas, A. Psallidas, I. Topsis-Giotis

National and Kapodistrian University of Athens, Athens, Greece

A. Agapitos, S. Kesisoglou, A. Panagiotou, N. Saoulidou, E. Tziaferi

University of Ioánnina, Ioánnina, Greece

I. Evangelou, G. Flouris, C. Foudas, P. Kokkas, N. Loukas, N. Manthos, I. Papadopoulos, E. Paradas, J. Strologas

Wigner Research Centre for Physics, Budapest, Hungary

G. Bencze, C. Hajdu, A. Hazi, P. Hidas, D. Horvath²², F. Sikler, V. Veszpremi, G. Vesztergombi²³, A.J. Zsigmond

Institute of Nuclear Research ATOMKI, Debrecen, Hungary

N. Beni, S. Czellar, J. Karancsi²⁴, J. Molnar, Z. Szillasi

University of Debrecen, Debrecen, Hungary

M. Bartók²⁵, A. Makovec, P. Raics, Z.L. Trocsanyi, B. Ujvari

National Institute of Science Education and Research, Bhubaneswar, India

P. Mal, K. Mandal, D.K. Sahoo, N. Sahoo, S.K. Swain

Panjab University, Chandigarh, India

S. Bansal, S.B. Beri, V. Bhatnagar, R. Chawla, R. Gupta, U. Bhawandeep, A.K. Kalsi, A. Kaur, M. Kaur, R. Kumar, A. Mehta, M. Mittal, J.B. Singh, G. Walia

University of Delhi, Delhi, India

Ashok Kumar, A. Bhardwaj, B.C. Choudhary, R.B. Garg, A. Kumar, S. Malhotra, M. Naimuddin, N. Nishu, K. Ranjan, R. Sharma, V. Sharma

Saha Institute of Nuclear Physics, Kolkata, India

S. Bhattacharya, K. Chatterjee, S. Dey, S. Dutta, Sa. Jain, N. Majumdar, A. Modak, K. Mondal, S. Mukherjee, S. Mukhopadhyay, A. Roy, D. Roy, S. Roy Chowdhury, S. Sarkar, M. Sharan

Bhabha Atomic Research Centre, Mumbai, India

A. Abdulsalam, R. Chudasama, D. Dutta, V. Jha, V. Kumar, A.K. Mohanty², L.M. Pant, P. Shukla, A. Topkar

Tata Institute of Fundamental Research, Mumbai, India

T. Aziz, S. Banerjee, S. Bhowmik²⁶, R.M. Chatterjee, R.K. Dewanjee, S. Dugad, S. Ganguly, S. Ghosh, M. Guchait, A. Gurtu²⁷, G. Kole, S. Kumar, B. Mahakud, M. Maity²⁶, G. Majumder, K. Mazumdar, S. Mitra, G.B. Mohanty, B. Parida, T. Sarkar²⁶, N. Sur, B. Sutar, N. Wickramage²⁸

Indian Institute of Science Education and Research (IISER), Pune, India

S. Chauhan, S. Dube, K. Kothekar, S. Sharma

Institute for Research in Fundamental Sciences (IPM), Tehran, Iran

H. Bakhshiansohi, H. Behnamian, S.M. Etesami²⁹, A. Fahim³⁰, R. Goldouzian, M. Khakzad, M. Mohammadi Najafabadi, M. Naseri, S. Paktinat Mehdiabadi, F. Rezaei Hosseinabadi, B. Safarzadeh³¹, M. Zeinali

University College Dublin, Dublin, Ireland

M. Felcini, M. Grunewald

INFN Sezione di Bari ^a, Università di Bari ^b, Politecnico di Bari ^c, Bari, Italy

M. Abbrescia^{a,b}, C. Calabria^{a,b}, C. Caputo^{a,b}, A. Colaleo^a, D. Creanza^{a,c}, L. Cristella^{a,b}, N. De Filippis^{a,c}, M. De Palma^{a,b}, L. Fiore^a, G. Iaselli^{a,c}, G. Maggi^{a,c}, M. Maggi^a, G. Miniello^{a,b}, S. My^{a,c}, S. Nuzzo^{a,b}, A. Pompili^{a,b}, G. Pugliese^{a,c}, R. Radogna^{a,b}, A. Ranieri^a, G. Selvaggi^{a,b}, L. Silvestris^{a,2}, R. Venditti^{a,b}, P. Verwilligen^a

INFN Sezione di Bologna ^a, Università di Bologna ^b, Bologna, Italy

G. Abbiendi^a, C. Battilana², A.C. Benvenuti^a, D. Bonacorsi^{a,b}, S. Braibant-Giacomelli^{a,b}, L. Brigliadori^{a,b}, R. Campanini^{a,b}, P. Capiluppi^{a,b}, A. Castro^{a,b}, F.R. Cavallo^a, S.S. Chhibra^{a,b}, G. Codispoti^{a,b}, M. Cuffiani^{a,b}, G.M. Dallavalle^a, F. Fabbri^a, A. Fanfani^{a,b}, D. Fasanella^{a,b}, P. Giacomelli^a, C. Grandi^a, L. Guiducci^{a,b}, S. Marcellini^a, G. Masetti^a, A. Montanari^a, F.L. Navarria^{a,b}, A. Perrotta^a, A.M. Rossi^{a,b}, T. Rovelli^{a,b}, G.P. Siroli^{a,b}, N. Tosi^{a,b}, R. Travaglini^{a,b}

INFN Sezione di Catania ^a, Università di Catania ^b, Catania, Italy

G. Cappello^a, M. Chiorboli^{a,b}, S. Costa^{a,b}, A. Di Mattia^a, F. Giordano^{a,b}, R. Potenza^{a,b}, A. Tricomi^{a,b}, C. Tuve^{a,b}

INFN Sezione di Firenze ^a, Università di Firenze ^b, Firenze, Italy

G. Barbagli^a, V. Ciulli^{a,b}, C. Civinini^a, R. D'Alessandro^{a,b}, E. Focardi^{a,b}, S. Gonzi^{a,b}, V. Gori^{a,b}, P. Lenzi^{a,b}, M. Meschini^a, S. Paoletti^a, G. Sguazzoni^a, A. Tropiano^{a,b}, L. Viliani^{a,b,2}

INFN Laboratori Nazionali di Frascati, Frascati, Italy

L. Benussi, S. Bianco, F. Fabbri, D. Piccolo, F. Primavera

INFN Sezione di Genova ^a, Università di Genova ^b, Genova, Italy

V. Calvelli^{a,b}, F. Ferro^a, M. Lo Vetere^{a,b}, M.R. Monge^{a,b}, E. Robutti^a, S. Tosi^{a,b}

INFN Sezione di Milano-Bicocca ^a, Università di Milano-Bicocca ^b, Milano, Italy

L. Brianza, M.E. Dinardo^{a,b}, S. Fiorendi^{a,b}, S. Gennai^a, R. Gerosa^{a,b}, A. Ghezzi^{a,b}, P. Govoni^{a,b}, S. Malvezzi^a, R.A. Manzoni^{a,b}, B. Marzocchi^{a,b,2}, D. Menasce^a, L. Moroni^a, M. Paganoni^{a,b}, D. Pedrini^a, S. Ragazzi^{a,b}, N. Redaelli^a, T. Tabarelli de Fatis^{a,b}

INFN Sezione di Napoli ^a, Università di Napoli 'Federico II' ^b, Napoli, Italy, Università della Basilicata ^c, Potenza, Italy, Università G. Marconi ^d, Roma, Italy

S. Buontempo^a, N. Cavallo^{a,c}, S. Di Guida^{a,d,2}, M. Esposito^{a,b}, F. Fabozzi^{a,c}, A.O.M. Iorio^{a,b}, G. Lanza^a, L. Lista^a, S. Meola^{a,d,2}, M. Merola^a, P. Paolucci^{a,2}, C. Sciacca^{a,b}, F. Thyssen

INFN Sezione di Padova ^a, Università di Padova ^b, Padova, Italy, Università di Trento ^c, Trento, Italy

N. Bacchetta^a, M. Bellato^a, L. Benato^{a,b}, D. Bisello^{a,b}, A. Boletti^{a,b}, R. Carlin^{a,b}, P. Checchia^a, M. Dall'Osso^{a,b,2}, U. Dosselli^a, F. Gasparini^{a,b}, U. Gasparini^{a,b}, A. Gozzelino^a, S. Lacaprara^a, M. Margoni^{a,b}, A.T. Meneguzzo^{a,b}, F. Montecassiano^a, M. Passaseo^a, J. Pazzini^{a,b}, M. Pegoraro^a, N. Pozzobon^{a,b}, F. Simonetto^{a,b}, E. Torassa^a, M. Tosi^{a,b}, S. Vanini^{a,b}, S. Ventura^a, M. Zanetti, P. Zotto^{a,b}, A. Zucchetta^{a,b,2}, G. Zumerle^{a,b}

INFN Sezione di Pavia ^a, Università di Pavia ^b, Pavia, Italy

A. Braghieri^a, A. Magnani^a, P. Montagna^{a,b}, S.P. Ratti^{a,b}, V. Re^a, C. Riccardi^{a,b}, P. Salvini^a, I. Vai^a, P. Vitulo^{a,b}

INFN Sezione di Perugia ^a, Università di Perugia ^b, Perugia, Italy

L. Alunni Solestizi^{a,b}, M. Biasini^{a,b}, G.M. Bilei^a, D. Ciangottini^{a,b,2}, L. Fanò^{a,b}, P. Lariccia^{a,b}, G. Mantovani^{a,b}, M. Menichelli^a, A. Saha^a, A. Santocchia^{a,b}

INFN Sezione di Pisa ^a, Università di Pisa ^b, Scuola Normale Superiore di Pisa ^c, Pisa, Italy

K. Androsov^{a,32}, P. Azzurri^a, G. Bagliesi^a, J. Bernardini^a, T. Boccali^a, R. Castaldi^a, M.A. Ciocci^{a,32}, R. Dell’Orso^a, S. Donato^{a,c,2}, G. Fedi, L. Foà^{a,c†}, A. Giassi^a, M.T. Grippo^{a,32}, F. Ligabue^{a,c}, T. Lomtadze^a, L. Martini^{a,b}, A. Messineo^{a,b}, F. Palla^a, A. Rizzi^{a,b}, A. Savoy-Navarro^{a,33}, A.T. Serban^a, P. Spagnolo^a, R. Tenchini^a, G. Tonelli^{a,b}, A. Venturi^a, P.G. Verдини^a

INFN Sezione di Roma ^a, Università di Roma ^b, Roma, Italy

L. Barone^{a,b}, F. Cavallari^a, G. D’imperio^{a,b,2}, D. Del Re^{a,b}, M. Diemoz^a, S. Gelli^{a,b}, C. Jorda^a, E. Longo^{a,b}, F. Margaroli^{a,b}, P. Meridiani^a, G. Organtini^{a,b}, R. Paramatti^a, F. Preiato^{a,b}, S. Rahatlou^{a,b}, C. Rovelli^a, F. Santanastasio^{a,b}, P. Traczyk^{a,b,2}

INFN Sezione di Torino ^a, Università di Torino ^b, Torino, Italy, Università del Piemonte Orientale ^c, Novara, Italy

N. Amapane^{a,b}, R. Arcidiacono^{a,c,2}, S. Argiro^{a,b}, M. Arneodo^{a,c}, R. Bellan^{a,b}, C. Biino^a, N. Cartiglia^a, M. Costa^{a,b}, R. Covarelli^{a,b}, A. Degano^{a,b}, N. Demaria^a, L. Finco^{a,b,2}, B. Kiani^{a,b}, C. Mariotti^a, S. Maselli^a, E. Migliore^{a,b}, V. Monaco^{a,b}, E. Monteil^{a,b}, M.M. Obertino^{a,b}, L. Pacher^{a,b}, N. Pastrone^a, M. Pelliccioni^a, G.L. Pinna Angioni^{a,b}, F. Ravera^{a,b}, A. Romero^{a,b}, M. Ruspa^{a,c}, R. Sacchi^{a,b}, A. Solano^{a,b}, A. Staiano^a, U. Tamponi^a

INFN Sezione di Trieste ^a, Università di Trieste ^b, Trieste, Italy

S. Belforte^a, V. Candelise^{a,b,2}, M. Casarsa^a, F. Cossutti^a, G. Della Ricca^{a,b}, B. Gobbo^a, C. La Licata^{a,b}, M. Marone^{a,b}, A. Schizzi^{a,b}, A. Zanetti^a

Kangwon National University, Chunchon, Korea

A. Kropivnitskaya, S.K. Nam

Kyungpook National University, Daegu, Korea

D.H. Kim, G.N. Kim, M.S. Kim, D.J. Kong, S. Lee, Y.D. Oh, A. Sakharov, D.C. Son

Chonbuk National University, Jeonju, Korea

J.A. Brochero Cifuentes, H. Kim, T.J. Kim³⁴

Chonnam National University, Institute for Universe and Elementary Particles, Kwangju, Korea

S. Song

Korea University, Seoul, Korea

S. Choi, Y. Go, D. Gyun, B. Hong, M. Jo, H. Kim, Y. Kim, B. Lee, K. Lee, K.S. Lee, S. Lee, S.K. Park, Y. Roh

Seoul National University, Seoul, Korea

H.D. Yoo

University of Seoul, Seoul, Korea

M. Choi, H. Kim, J.H. Kim, J.S.H. Lee, I.C. Park, G. Ryu, M.S. Ryu

Sungkyunkwan University, Suwon, Korea

Y. Choi, J. Goh, D. Kim, E. Kwon, J. Lee, I. Yu

Vilnius University, Vilnius, Lithuania

V. Dudenas, A. Juodagalvis, J. Vaitkus

National Centre for Particle Physics, Universiti Malaya, Kuala Lumpur, MalaysiaI. Ahmed, Z.A. Ibrahim, J.R. Komaragiri, M.A.B. Md Ali³⁵, F. Mohamad Idris³⁶, W.A.T. Wan Abdullah, M.N. Yusli**Centro de Investigacion y de Estudios Avanzados del IPN, Mexico City, Mexico**E. Casimiro Linares, H. Castilla-Valdez, E. De La Cruz-Burelo, I. Heredia-De La Cruz³⁷, A. Hernandez-Almada, R. Lopez-Fernandez, A. Sanchez-Hernandez**Universidad Iberoamericana, Mexico City, Mexico**

S. Carrillo Moreno, F. Vazquez Valencia

Benemerita Universidad Autonoma de Puebla, Puebla, Mexico

I. Pedraza, H.A. Salazar Ibarguen

Universidad Autónoma de San Luis Potosí, San Luis Potosí, Mexico

A. Morelos Pineda

University of Auckland, Auckland, New Zealand

D. Krofcheck

University of Canterbury, Christchurch, New Zealand

P.H. Butler

National Centre for Physics, Quaid-I-Azam University, Islamabad, Pakistan

A. Ahmad, M. Ahmad, Q. Hassan, H.R. Hoorani, W.A. Khan, T. Khurshid, M. Shoaib

National Centre for Nuclear Research, Swierk, Poland

H. Bialkowska, M. Bluj, B. Boimska, T. Frueboes, M. Górski, M. Kazana, K. Nawrocki, K. Romanowska-Rybinska, M. Szleper, P. Zalewski

Institute of Experimental Physics, Faculty of Physics, University of Warsaw, Warsaw, PolandG. Brona, K. Bunkowski, A. Byszuk³⁸, K. Doroba, A. Kalinowski, K. Kierzkowski, M. Konecki, J. Krolikowski, M. Misiura, W. Oklinski, M. Olszewski, K. Pozniak³⁸, M. Walczak, W. Zabolotny**Laboratório de Instrumentação e Física Experimental de Partículas, Lisboa, Portugal**

P. Bargassa, C. Beirão Da Cruz E Silva, A. Di Francesco, P. Faccioli, P.G. Ferreira Parracho, M. Gallinaro, N. Leonardo, L. Lloret Iglesias, F. Nguyen, J. Rodrigues Antunes, J. Seixas, O. Toldaiev, D. Vadrucio, J. Varela, P. Vischia

Joint Institute for Nuclear Research, Dubna, Russia

S. Afanasiev, P. Bunin, M. Gavrilenko, I. Golutvin, I. Gorbunov, A. Kamenev, V. Karjavin, V. Konoplyanikov, A. Lanev, A. Malakhov, V. Matveev^{39,40}, P. Moisezenz, V. Palichik, V. Pereygin, S. Shmatov, S. Shulha, N. Skatchkov, V. Smirnov, A. Zarubin

Petersburg Nuclear Physics Institute, Gatchina (St. Petersburg), Russia

V. Golovtsov, Y. Ivanov, V. Kim⁴¹, E. Kuznetsova, P. Levchenko, V. Murzin, V. Oreshkin, I. Smirnov, V. Sulimov, L. Uvarov, S. Vavilov, A. Vorobyev

Institute for Nuclear Research, Moscow, Russia

Yu. Andreev, A. Dermenev, S. Gninenko, N. Golubev, A. Karneyeu, M. Kirsanov, N. Krasnikov, A. Pashenkov, D. Tlisov, A. Toropin

Institute for Theoretical and Experimental Physics, Moscow, Russia

V. Epshteyn, V. Gavrilov, N. Lychkovskaya, V. Popov, I. Pozdnyakov, G. Safronov, A. Spiridonov, E. Vlasov, A. Zhokin

National Research Nuclear University 'Moscow Engineering Physics Institute' (MEPhI), Moscow, Russia

A. Bylinkin

P.N. Lebedev Physical Institute, Moscow, Russia

V. Andreev, M. Azarkin⁴⁰, I. Dremin⁴⁰, M. Kirakosyan, A. Leonidov⁴⁰, G. Mesyats, S.V. Rusakov

Skobeltsyn Institute of Nuclear Physics, Lomonosov Moscow State University, Moscow, Russia

A. Baskakov, A. Belyaev, E. Boos, M. Dubinin⁴², L. Dudko, A. Ershov, A. Gribushin, A. Kaminskiy⁴³, V. Klyukhin, O. Kodolova, I. Lokhtin, I. Myagkov, S. Obraztsov, S. Petrushanko, V. Savrin

State Research Center of Russian Federation, Institute for High Energy Physics, Protvino, Russia

I. Azhgirey, I. Bayshev, S. Bitioukov, V. Kachanov, A. Kalinin, D. Konstantinov, V. Krychkine, V. Petrov, R. Ryutin, A. Sobol, L. Tourtchanovitch, S. Troshin, N. Tyurin, A. Uzunian, A. Volkov

University of Belgrade, Faculty of Physics and Vinca Institute of Nuclear Sciences, Belgrade, Serbia

P. Adzic⁴⁴, J. Milosevic, V. Rekovic

Centro de Investigaciones Energéticas Medioambientales y Tecnológicas (CIEMAT), Madrid, Spain

J. Alcaraz Maestre, E. Calvo, M. Cerrada, M. Chamizo Llatas, N. Colino, B. De La Cruz, A. Delgado Peris, D. Domínguez Vázquez, A. Escalante Del Valle, C. Fernandez Bedoya, J.P. Fernández Ramos, J. Flix, M.C. Fouz, P. Garcia-Abia, O. Gonzalez Lopez, S. Goy Lopez, J.M. Hernandez, M.I. Josa, E. Navarro De Martino, A. Pérez-Calero Yzquierdo, J. Puerta Pelayo, A. Quintario Olmeda, I. Redondo, L. Romero, J. Santaolalla, M.S. Soares

Universidad Autónoma de Madrid, Madrid, Spain

C. Albajar, J.F. de Trocóniz, M. Missiroli, D. Moran

Universidad de Oviedo, Oviedo, Spain

J. Cuevas, J. Fernandez Menendez, S. Folgueras, I. Gonzalez Caballero, E. Palencia Cortezon, J.M. Vizan Garcia

Instituto de Física de Cantabria (IFCA), CSIC-Universidad de Cantabria, Santander, Spain

I.J. Cabrillo, A. Calderon, J.R. Castiñeiras De Saa, P. De Castro Manzano, J. Duarte Campderros, M. Fernandez, J. Garcia-Ferrero, G. Gomez, A. Lopez Virto, J. Marco, R. Marco, C. Martinez Rivero, F. Matorras, F.J. Munoz Sanchez, J. Piedra Gomez, T. Rodrigo, A.Y. Rodríguez-Marrero, A. Ruiz-Jimeno, L. Scodellaro, N. Trevisani, I. Vila, R. Vilar Cortabitarte

CERN, European Organization for Nuclear Research, Geneva, Switzerland

D. Abbaneo, E. Auffray, G. Auzinger, M. Bachtis, P. Baillon, A.H. Ball, D. Barney, A. Benaglia, J. Bendavid, L. Benhabib, J.F. Benitez, G.M. Berruti, P. Bloch, A. Bocci, A. Bonato, C. Botta, H. Breuker, T. Camporesi, R. Castello, G. Cerminara, M. D'Alfonso, D. d'Enterria, A. Dabrowski, V. Daponte, A. David, M. De Gruttola, F. De Guio, A. De Roeck, S. De Visscher, E. Di Marco, M. Dobson, M. Dordevic, B. Dorney, T. du Pree, M. Dünser, N. Dupont, A. Elliott-Peisert, G. Franzoni, W. Funk, D. Gigi, K. Gill, D. Giordano, M. Girone, F. Glege, R. Guida, S. Gundacker, M. Guthoff, J. Hammer, P. Harris, J. Hegeman, V. Innocente, P. Janot, H. Kirschenmann, M.J. Kortelainen, K. Kousouris, K. Krajczar, P. Lecoq, C. Lourenço, M.T. Lucchini, N. Magini, L. Malgeri, M. Mannelli, A. Martelli, L. Masetti, F. Meijers, S. Mersi, E. Meschi, F. Moortgat, S. Morovic, M. Mulders, M.V. Nemallapudi, H. Neugebauer, S. Orfanelli⁴⁵, L. Orsini, L. Pape, E. Perez, M. Peruzzi, A. Petrilli, G. Petrucciani, A. Pfeiffer, D. Piparo, A. Racz, G. Rolandi⁴⁶, M. Rovere, M. Ruan, H. Sakulin, C. Schäfer, C. Schwick, M. Seidel, A. Sharma, P. Silva, M. Simon, P. Sphicas⁴⁷, J. Steggemann, B. Stieger, M. Stoye, Y. Takahashi, D. Treille, A. Triossi, A. Tsiros, G.I. Veres²³, N. Wardle, H.K. Wöhri, A. Zagozdinska³⁸, W.D. Zeuner

Paul Scherrer Institut, Villigen, Switzerland

W. Bertl, K. Deiters, W. Erdmann, R. Horisberger, Q. Ingram, H.C. Kaestli, D. Kotlinski, U. Langenegger, D. Renker, T. Rohe

Institute for Particle Physics, ETH Zurich, Zurich, Switzerland

F. Bachmair, L. Bäni, L. Bianchini, B. Casal, G. Dissertori, M. Dittmar, M. Donegà, P. Eller, C. Grab, C. Heidegger, D. Hits, J. Hoss, G. Kasieczka, W. Lustermann, B. Mangano, M. Marionneau, P. Martinez Ruiz del Arbol, M. Masciovecchio, D. Meister, F. Micheli, P. Musella, F. Nessi-Tedaldi, F. Pandolfi, J. Pata, F. Pauss, L. Perrozzi, M. Quittnat, M. Rossini, A. Starodumov⁴⁸, M. Takahashi, V.R. Tavolaro, K. Theofilatos, R. Wallny

Universität Zürich, Zurich, Switzerland

T.K. Aarrestad, C. Amsler⁴⁹, L. Caminada, M.F. Canelli, V. Chiochia, A. De Cosa, C. Galloni, A. Hinzmann, T. Hreus, B. Kilminster, C. Lange, J. Ngadiuba, D. Pinna, P. Robmann, F.J. Ronga, D. Salerno, Y. Yang

National Central University, Chung-Li, Taiwan

M. Cardaci, K.H. Chen, T.H. Doan, Sh. Jain, R. Khurana, M. Konyushikhin, C.M. Kuo, W. Lin, Y.J. Lu, S.S. Yu

National Taiwan University (NTU), Taipei, Taiwan

Arun Kumar, R. Bartek, P. Chang, Y.H. Chang, Y.W. Chang, Y. Chao, K.F. Chen, P.H. Chen, C. Dietz, F. Fiori, U. Grundler, W.-S. Hou, Y. Hsiung, Y.F. Liu, R.-S. Lu, M. Miñano Moya, E. Petrakou, J.f. Tsai, Y.M. Tzeng

Chulalongkorn University, Faculty of Science, Department of Physics, Bangkok, Thailand

B. Asavapibhop, K. Kovitanggoon, G. Singh, N. Srimanobhas, N. Suwonjandee

Cukurova University, Adana, Turkey

A. Adiguzel, M.N. Bakirci⁵⁰, Z.S. Demiroglu, C. Dozen, E. Eskut, S. Girgis, G. Gokbulut, Y. Guler, E. Gurpinar, I. Hos, E.E. Kangal⁵¹, G. Onengut⁵², K. Ozdemir⁵³, A. Polatoz, D. Sunar Cerci⁵⁴, B. Tali⁵⁴, H. Topakli⁵⁰, M. Vergili, C. Zorbilmez

Middle East Technical University, Physics Department, Ankara, Turkey

I.V. Akin, B. Bilin, S. Bilmis, B. Isildak⁵⁵, G. Karapinar⁵⁶, M. Yalvac, M. Zeyrek

Bogazici University, Istanbul, Turkey

E. Gülmez, M. Kaya⁵⁷, O. Kaya⁵⁸, E.A. Yetkin⁵⁹, T. Yetkin⁶⁰

Istanbul Technical University, Istanbul, Turkey

A. Cakir, K. Cankocak, S. Sen⁶¹, F.I. Vardarli

Institute for Scintillation Materials of National Academy of Science of Ukraine, Kharkov, Ukraine

B. Grynyov

National Scientific Center, Kharkov Institute of Physics and Technology, Kharkov, Ukraine

L. Levchuk, P. Sorokin

University of Bristol, Bristol, United Kingdom

R. Aggleton, F. Ball, L. Beck, J.J. Brooke, E. Clement, D. Cussans, H. Flacher, J. Goldstein, M. Grimes, G.P. Heath, H.F. Heath, J. Jacob, L. Kreczko, C. Lucas, Z. Meng, D.M. Newbold⁶², S. Paramesvaran, A. Poll, T. Sakuma, S. Seif El Nasr-storey, S. Senkin, D. Smith, V.J. Smith

Rutherford Appleton Laboratory, Didcot, United Kingdom

K.W. Bell, A. Belyaev⁶³, C. Brew, R.M. Brown, L. Calligaris, D. Cieri, D.J.A. Cockerill, J.A. Coughlan, K. Harder, S. Harper, E. Olaiya, D. Petyt, C.H. Shepherd-Themistocleous, A. Thea, I.R. Tomalin, T. Williams, W.J. Womersley, S.D. Worm

Imperial College, London, United Kingdom

M. Baber, R. Bainbridge, O. Buchmuller, A. Bundock, D. Burton, S. Casasso, M. Citron, D. Colling, L. Corpe, N. Cripps, P. Dauncey, G. Davies, A. De Wit, M. Della Negra, P. Dunne, A. Elwood, W. Ferguson, J. Fulcher, D. Futyan, G. Hall, G. Iles, M. Kenzie, R. Lane, R. Lucas⁶², L. Lyons, A.-M. Magnan, S. Malik, J. Nash, A. Nikitenko⁴⁸, J. Pela, M. Pesaresi, K. Petridis, D.M. Raymond, A. Richards, A. Rose, C. Seez, A. Tapper, K. Uchida, M. Vazquez Acosta⁶⁴, T. Virdee, S.C. Zenz

Brunel University, Uxbridge, United Kingdom

J.E. Cole, P.R. Hobson, A. Khan, P. Kyberd, D. Leggat, D. Leslie, I.D. Reid, P. Symonds, L. Teodorescu, M. Turner

Baylor University, Waco, U.S.A.

A. Borzou, K. Call, J. Dittmann, K. Hatakeyama, H. Liu, N. Pastika

The University of Alabama, Tuscaloosa, U.S.A.

O. Charaf, S.I. Cooper, C. Henderson, P. Rumerio

Boston University, Boston, U.S.A.

D. Arcaro, A. Avetisyan, T. Bose, C. Fantasia, D. Gastler, P. Lawson, D. Rankin, C. Richardson, J. Rohlf, J. St. John, L. Sulak, D. Zou

Brown University, Providence, U.S.A.

J. Alimena, E. Berry, S. Bhattacharya, D. Cutts, N. Dhirga, A. Ferapontov, A. Garabedian, J. Hakala, U. Heintz, E. Laird, G. Landsberg, Z. Mao, M. Narain, S. Piperov, S. Sagir, R. Syarif

University of California, Davis, Davis, U.S.A.

R. Breedon, G. Breto, M. Calderon De La Barca Sanchez, S. Chauhan, M. Chertok, J. Conway, R. Conway, P.T. Cox, R. Erbacher, M. Gardner, W. Ko, R. Lander, M. Mulhearn, D. Pellett, J. Pilot, F. Ricci-Tam, S. Shalhout, J. Smith, M. Squires, D. Stolp, M. Tripathi, S. Wilbur, R. Yohay

University of California, Los Angeles, U.S.A.

R. Cousins, P. Everaerts, C. Farrell, J. Hauser, M. Ignatenko, D. Saltzberg, E. Takasugi, V. Valuev, M. Weber

University of California, Riverside, Riverside, U.S.A.

K. Burt, R. Clare, J. Ellison, J.W. Gary, G. Hanson, J. Heilman, M. Iova, P. Jandir, E. Kennedy, F. Lacroix, O.R. Long, A. Luthra, M. Malberti, M. Olmedo Negrete, A. Shrinivas, H. Wei, S. Wimpenny, B. R. Yates

University of California, San Diego, La Jolla, U.S.A.

J.G. Branson, G.B. Cerati, S. Cittolin, R.T. D'Agnolo, M. Derdzinski, A. Holzner, R. Kelley, D. Klein, J. Letts, I. Macneill, D. Olivito, S. Padhi, M. Pieri, M. Sani, V. Sharma, S. Simon, M. Tadel, A. Vartak, S. Wasserbaech⁶⁵, C. Welke, F. Würthwein, A. Yagil, G. Zevi Della Porta

University of California, Santa Barbara - Department of Physics, Santa Barbara, U.S.A.

J. Bradmiller-Feld, C. Campagnari, A. Dishaw, V. Dutta, K. Flowers, M. Franco Sevilla, P. Geffert, C. George, F. Golf, L. Gouskos, J. Gran, J. Incandela, N. Mccoll, S.D. Mullin, J. Richman, D. Stuart, I. Suarez, C. West, J. Yoo

California Institute of Technology, Pasadena, U.S.A.

D. Anderson, A. Apresyan, A. Bornheim, J. Bunn, Y. Chen, J. Duarte, A. Mott, H.B. Newman, C. Pena, M. Pierini, M. Spiropulu, J.R. Vlimant, S. Xie, R.Y. Zhu

Carnegie Mellon University, Pittsburgh, U.S.A.

M.B. Andrews, V. Azzolini, A. Calamba, B. Carlson, T. Ferguson, M. Paulini, J. Russ, M. Sun, H. Vogel, I. Vorobiev

University of Colorado Boulder, Boulder, U.S.A.

J.P. Cumalat, W.T. Ford, A. Gaz, F. Jensen, A. Johnson, M. Krohn, T. Mulholland, U. Nauenberg, K. Stenson, S.R. Wagner

Cornell University, Ithaca, U.S.A.

J. Alexander, A. Chatterjee, J. Chaves, J. Chu, S. Dittmer, N. Eggert, N. Mirman, G. Nicolas Kaufman, J.R. Patterson, A. Rinkevicius, A. Ryd, L. Skinnari, L. Soffi, W. Sun, S.M. Tan, W.D. Teo, J. Thom, J. Thompson, J. Tucker, Y. Weng, P. Wittich

Fermi National Accelerator Laboratory, Batavia, U.S.A.

S. Abdullin, M. Albrow, J. Anderson, G. Apollinari, S. Banerjee, L.A.T. Bauerdick, A. Beretvas, J. Berryhill, P.C. Bhat, G. Bolla, K. Burkett, J.N. Butler, H.W.K. Cheung, F. Chlebana, S. Cihangir, V.D. Elvira, I. Fisk, J. Freeman, E. Gottschalk, L. Gray, D. Green, S. Grünendahl, O. Gutsche, J. Hanlon, D. Hare, R.M. Harris, S. Hasegawa, J. Hirschauer, Z. Hu, B. Jayatilaka, S. Jindariani, M. Johnson, U. Joshi, A.W. Jung, B. Klima, B. Kreis, S. Kwan[†], S. Lammel, J. Linacre, D. Lincoln, R. Lipton, T. Liu, R. Lopes De Sá, J. Lykken, K. Maeshima, J.M. Marraffino, V.I. Martinez Outschoorn, S. Maruyama, D. Mason, P. McBride, P. Merkel, K. Mishra, S. Mrenna, S. Nahn, C. Newman-Holmes, V. O'Dell, K. Pedro, O. Prokofyev, G. Rakness, E. Sexton-Kennedy, A. Soha, W.J. Spalding, L. Spiegel, L. Taylor, S. Tkaczyk, N.V. Tran, L. Uplegger, E.W. Vaandering, C. Vernieri, M. Verzocchi, R. Vidal, H.A. Weber, A. Whitbeck, F. Yang

University of Florida, Gainesville, U.S.A.

D. Acosta, P. Avery, P. Bortignon, D. Bourilkov, A. Carnes, M. Carver, D. Curry, S. Das, G.P. Di Giovanni, R.D. Field, I.K. Furic, S.V. Gleyzer, J. Hugon, J. Konigsberg, A. Korytov, J.F. Low, P. Ma, K. Matchev, H. Mei, P. Milenovic⁶⁶, G. Mitselmakher, D. Rank, R. Rossin, L. Shchutska, M. Snowball, D. Sperka, N. Terentyev, L. Thomas, J. Wang, S. Wang, J. Yelton

Florida International University, Miami, U.S.A.

S. Hewamanage, S. Linn, P. Markowitz, G. Martinez, J.L. Rodriguez

Florida State University, Tallahassee, U.S.A.

A. Ackert, J.R. Adams, T. Adams, A. Askew, J. Bochenek, B. Diamond, J. Haas, S. Hagopian, V. Hagopian, K.F. Johnson, A. Khatiwada, H. Prosper, M. Weinberg

Florida Institute of Technology, Melbourne, U.S.A.

M.M. Baarmand, V. Bhopatkar, S. Colafranceschi⁶⁷, M. Hohlmann, H. Kalakhety, D. Noonan, T. Roy, F. Yumiceva

University of Illinois at Chicago (UIC), Chicago, U.S.A.

M.R. Adams, L. Apanasevich, D. Berry, R.R. Betts, I. Bucinskaite, R. Cavanaugh, O. Evdokimov, L. Gauthier, C.E. Gerber, D.J. Hofman, P. Kurt, C. O'Brien, I.D. Sandoval Gonzalez, C. Silkworth, P. Turner, N. Varelas, Z. Wu, M. Zakaria

The University of Iowa, Iowa City, U.S.A.

B. Bilki⁶⁸, W. Clarida, K. Dilsiz, S. Durgut, R.P. Gandrajula, M. Haytmyradov, V. Khristenko, J.-P. Merlo, H. Mermerkaya⁶⁹, A. Mestvirishvili, A. Moeller, J. Nachtman, H. Ogul, Y. Onel, F. Ozok⁷⁰, A. Penzo, C. Snyder, E. Tiras, J. Wetzel, K. Yi

Johns Hopkins University, Baltimore, U.S.A.

I. Anderson, B.A. Barnett, B. Blumenfeld, N. Eminizer, D. Fehling, L. Feng, A.V. Gritsan, P. Maksimovic, C. Martin, M. Osherson, J. Roskes, A. Sady, U. Sarica, M. Swartz, M. Xiao, Y. Xin, C. You

The University of Kansas, Lawrence, U.S.A.

P. Baringer, A. Bean, G. Benelli, C. Bruner, R.P. Kenny III, D. Majumder, M. Malek, M. Murray, S. Sanders, R. Stringer, Q. Wang

Kansas State University, Manhattan, U.S.A.

A. Ivanov, K. Kaadze, S. Khalil, M. Makouski, Y. Maravin, A. Mohammadi, L.K. Saini, N. Skhirtladze, S. Toda

Lawrence Livermore National Laboratory, Livermore, U.S.A.

D. Lange, F. Rebassoo, D. Wright

University of Maryland, College Park, U.S.A.

C. Anelli, A. Baden, O. Baron, A. Belloni, B. Calvert, S.C. Eno, C. Ferraioli, J.A. Gomez, N.J. Hadley, S. Jabeen, R.G. Kellogg, T. Kolberg, J. Kunkle, Y. Lu, A.C. Mignerey, Y.H. Shin, A. Skuja, M.B. Tonjes, S.C. Tonwar

Massachusetts Institute of Technology, Cambridge, U.S.A.

A. Apyan, R. Barbieri, A. Baty, K. Bierwagen, S. Brandt, W. Busza, I.A. Cali, Z. Demiragli, L. Di Matteo, G. Gomez Ceballos, M. Goncharov, D. Gulhan, Y. Iiyama, G.M. Innocenti, M. Klute, D. Kovalskyi, Y.S. Lai, Y.-J. Lee, A. Levin, P.D. Luckey, A.C. Marini, C. McGinn, C. Mironov, S. Narayanan, X. Niu, C. Paus, D. Ralph, C. Roland, G. Roland, J. Salfeld-Nebgen, G.S.F. Stephans, K. Sumorok, M. Varma, D. Velicanu, J. Veverka, J. Wang, T.W. Wang, B. Wyslouch, M. Yang, V. Zhukova

University of Minnesota, Minneapolis, U.S.A.

B. Dahmes, A. Evans, A. Finkel, A. Gude, P. Hansen, S. Kalafut, S.C. Kao, K. Klapoetke, Y. Kubota, Z. Lesko, J. Mans, S. Nourbakhsh, N. Ruckstuhl, R. Rusack, N. Tambe, J. Turkewitz

University of Mississippi, Oxford, U.S.A.

J.G. Acosta, S. Oliveros

University of Nebraska-Lincoln, Lincoln, U.S.A.

E. Avdeeva, K. Bloom, S. Bose, D.R. Claes, A. Dominguez, C. Fangmeier, R. Gonzalez Suarez, R. Kamalieddin, J. Keller, D. Knowlton, I. Kravchenko, F. Meier, J. Monroy, F. Ratnikov, J.E. Siado, G.R. Snow

State University of New York at Buffalo, Buffalo, U.S.A.

M. Alyari, J. Dolen, J. George, A. Godshalk, C. Harrington, I. Iashvili, J. Kaisen, A. Kharchilava, A. Kumar, S. Rappoccio, B. Roozbahani

Northeastern University, Boston, U.S.A.

G. Alverson, E. Barberis, D. Baumgartel, M. Chasco, A. Hortiangtham, A. Massironi, D.M. Morse, D. Nash, T. Orimoto, R. Teixeira De Lima, D. Trocino, R.-J. Wang, D. Wood, J. Zhang

Northwestern University, Evanston, U.S.A.

K.A. Hahn, A. Kubik, N. Mucia, N. Odell, B. Pollack, A. Pozdnyakov, M. Schmitt, S. Stoynev, K. Sung, M. Trovato, M. Velasco

University of Notre Dame, Notre Dame, U.S.A.

A. Brinkerhoff, N. Dev, M. Hildreth, C. Jessop, D.J. Karmgard, N. Kellams, K. Lannon, S. Lynch, N. Marinelli, F. Meng, C. Mueller, Y. Musienko³⁹, T. Pearson, M. Planer, A. Reinsvold, R. Ruchti, G. Smith, S. Taroni, N. Valls, M. Wayne, M. Wolf, A. Woodard

The Ohio State University, Columbus, U.S.A.

L. Antonelli, J. Brinson, B. Bylsma, L.S. Durkin, S. Flowers, A. Hart, C. Hill, R. Hughes, W. Ji, K. Kotov, T.Y. Ling, B. Liu, W. Luo, D. Puigh, M. Rodenburg, B.L. Winer, H.W. Wulsin

Princeton University, Princeton, U.S.A.

O. Driga, P. Elmer, J. Hardenbrook, P. Hebda, S.A. Koay, P. Lujan, D. Marlow, T. Medvedeva, M. Mooney, J. Olsen, C. Palmer, P. Piroué, H. Saka, D. Stickland, C. Tully, A. Zuranski

University of Puerto Rico, Mayaguez, U.S.A.

S. Malik

Purdue University, West Lafayette, U.S.A.

V.E. Barnes, D. Benedetti, D. Bortoletto, L. Gutay, M.K. Jha, M. Jones, K. Jung, D.H. Miller, N. Neumeister, B.C. Radburn-Smith, X. Shi, I. Shipsey, D. Silvers, J. Sun, A. Svyatkovskiy, F. Wang, W. Xie, L. Xu

Purdue University Calumet, Hammond, U.S.A.

N. Parashar, J. Stupak

Rice University, Houston, U.S.A.

A. Adair, B. Akgun, Z. Chen, K.M. Ecklund, F.J.M. Geurts, M. Guilbaud, W. Li, B. Michlin, M. Northup, B.P. Padley, R. Redjimi, J. Roberts, J. Rorie, Z. Tu, J. Zabel

University of Rochester, Rochester, U.S.A.

B. Betchart, A. Bodek, P. de Barbaro, R. Demina, Y. Eshaq, T. Ferbel, M. Galanti, A. Garcia-Bellido, J. Han, A. Harel, O. Hindrichs, A. Khukhunaishvili, G. Petrillo, P. Tan, M. Verzetti

Rutgers, The State University of New Jersey, Piscataway, U.S.A.

S. Arora, A. Barker, J.P. Chou, C. Contreras-Campana, E. Contreras-Campana, D. Duggan, D. Ferencek, Y. Gershtein, R. Gray, E. Halkiadakis, D. Hidas, E. Hughes, S. Kaplan, R. Kunawalkam Elayavalli, A. Lath, K. Nash, S. Panwalkar, M. Park, S. Salur, S. Schnetzer, D. Sheffield, S. Somalwar, R. Stone, S. Thomas, P. Thomassen, M. Walker

University of Tennessee, Knoxville, U.S.A.

M. Foerster, G. Riley, K. Rose, S. Spanier, A. York

Texas A&M University, College Station, U.S.A.

O. Bouhali⁷¹, A. Castaneda Hernandez⁷¹, M. Dalchenko, M. De Mattia, A. Delgado, S. Dildick, R. Eusebi, J. Gilmore, T. Kamon⁷², V. Krutelyov, R. Mueller, I. Osipenko, Y. Pakhotin, R. Patel, A. Perloff, A. Rose, A. Safonov, A. Tatarinov, K.A. Ulmer²

Texas Tech University, Lubbock, U.S.A.

N. Akchurin, C. Cowden, J. Damgov, C. Dragoiu, P.R. Duder, J. Faulkner, S. Kunori, K. Lamichhane, S.W. Lee, T. Libeiro, S. Undleeb, I. Volobouev

Vanderbilt University, Nashville, U.S.A.

E. Appelt, A.G. Delannoy, S. Greene, A. Gurrola, R. Janjam, W. Johns, C. Maguire, Y. Mao, A. Melo, H. Ni, P. Sheldon, B. Snook, S. Tuo, J. Velkovska, Q. Xu

University of Virginia, Charlottesville, U.S.A.

M.W. Arenton, B. Cox, B. Francis, J. Goodell, R. Hirosky, A. Ledovskoy, H. Li, C. Lin, C. Neu, T. Sinthuprasith, X. Sun, Y. Wang, E. Wolfe, J. Wood, F. Xia

Wayne State University, Detroit, U.S.A.

C. Clarke, R. Harr, P.E. Karchin, C. Kottachchi Kankanamge Don, P. Lamichhane, J. Sturdy

University of Wisconsin - Madison, Madison, WI, U.S.A.

D.A. Belknap, D. Carlsmith, M. Cepeda, S. Dasu, L. Dodd, S. Duric, B. Gomber, M. Grothe, R. Hall-Wilton, M. Herndon, A. Hervé, P. Klabbers, A. Lanaro, A. Levine, K. Long, R. Loveless, A. Mohapatra, I. Ojalvo, T. Perry, G.A. Pierro, G. Polese, T. Ruggles, T. Sarangi, A. Savin, A. Sharma, N. Smith, W.H. Smith, D. Taylor, N. Woods

†: Deceased

- 1: Also at Vienna University of Technology, Vienna, Austria
- 2: Also at CERN, European Organization for Nuclear Research, Geneva, Switzerland
- 3: Also at State Key Laboratory of Nuclear Physics and Technology, Peking University, Beijing, China
- 4: Also at Institut Pluridisciplinaire Hubert Curien, Université de Strasbourg, Université de Haute Alsace Mulhouse, CNRS/IN2P3, Strasbourg, France
- 5: Also at National Institute of Chemical Physics and Biophysics, Tallinn, Estonia
- 6: Also at Skobeltsyn Institute of Nuclear Physics, Lomonosov Moscow State University, Moscow, Russia
- 7: Also at Universidade Estadual de Campinas, Campinas, Brazil
- 8: Also at Centre National de la Recherche Scientifique (CNRS) - IN2P3, Paris, France
- 9: Also at Laboratoire Leprince-Ringuet, Ecole Polytechnique, IN2P3-CNRS, Palaiseau, France
- 10: Also at Joint Institute for Nuclear Research, Dubna, Russia
- 11: Now at Suez University, Suez, Egypt
- 12: Also at Beni-Suef University, Bani Sweif, Egypt
- 13: Now at British University in Egypt, Cairo, Egypt
- 14: Also at Cairo University, Cairo, Egypt
- 15: Also at Fayoum University, El-Fayoum, Egypt
- 16: Also at Université de Haute Alsace, Mulhouse, France
- 17: Also at Tbilisi State University, Tbilisi, Georgia
- 18: Also at RWTH Aachen University, III. Physikalisches Institut A, Aachen, Germany
- 19: Also at Indian Institute of Science Education and Research, Bhopal, India
- 20: Also at University of Hamburg, Hamburg, Germany
- 21: Also at Brandenburg University of Technology, Cottbus, Germany
- 22: Also at Institute of Nuclear Research ATOMKI, Debrecen, Hungary
- 23: Also at Eötvös Loránd University, Budapest, Hungary
- 24: Also at University of Debrecen, Debrecen, Hungary
- 25: Also at Wigner Research Centre for Physics, Budapest, Hungary
- 26: Also at University of Visva-Bharati, Santiniketan, India
- 27: Now at King Abdulaziz University, Jeddah, Saudi Arabia
- 28: Also at University of Ruhuna, Matara, Sri Lanka

- 29: Also at Isfahan University of Technology, Isfahan, Iran
- 30: Also at University of Tehran, Department of Engineering Science, Tehran, Iran
- 31: Also at Plasma Physics Research Center, Science and Research Branch, Islamic Azad University, Tehran, Iran
- 32: Also at Università degli Studi di Siena, Siena, Italy
- 33: Also at Purdue University, West Lafayette, U.S.A.
- 34: Now at Hanyang University, Seoul, Korea
- 35: Also at International Islamic University of Malaysia, Kuala Lumpur, Malaysia
- 36: Also at Malaysian Nuclear Agency, MOSTI, Kajang, Malaysia
- 37: Also at Consejo Nacional de Ciencia y Tecnología, Mexico city, Mexico
- 38: Also at Warsaw University of Technology, Institute of Electronic Systems, Warsaw, Poland
- 39: Also at Institute for Nuclear Research, Moscow, Russia
- 40: Now at National Research Nuclear University 'Moscow Engineering Physics Institute' (MEPhI), Moscow, Russia
- 41: Also at St. Petersburg State Polytechnical University, St. Petersburg, Russia
- 42: Also at California Institute of Technology, Pasadena, U.S.A.
- 43: Also at INFN Sezione di Padova; Università di Padova; Università di Trento (Trento), Padova, Italy
- 44: Also at Faculty of Physics, University of Belgrade, Belgrade, Serbia
- 45: Also at National Technical University of Athens, Athens, Greece
- 46: Also at Scuola Normale e Sezione dell'INFN, Pisa, Italy
- 47: Also at National and Kapodistrian University of Athens, Athens, Greece
- 48: Also at Institute for Theoretical and Experimental Physics, Moscow, Russia
- 49: Also at Albert Einstein Center for Fundamental Physics, Bern, Switzerland
- 50: Also at Gaziosmanpasa University, Tokat, Turkey
- 51: Also at Mersin University, Mersin, Turkey
- 52: Also at Cag University, Mersin, Turkey
- 53: Also at Piri Reis University, Istanbul, Turkey
- 54: Also at Adiyaman University, Adiyaman, Turkey
- 55: Also at Ozyegin University, Istanbul, Turkey
- 56: Also at Izmir Institute of Technology, Izmir, Turkey
- 57: Also at Marmara University, Istanbul, Turkey
- 58: Also at Kafkas University, Kars, Turkey
- 59: Also at Istanbul Bilgi University, Istanbul, Turkey
- 60: Also at Yildiz Technical University, Istanbul, Turkey
- 61: Also at Hacettepe University, Ankara, Turkey
- 62: Also at Rutherford Appleton Laboratory, Didcot, United Kingdom
- 63: Also at School of Physics and Astronomy, University of Southampton, Southampton, United Kingdom
- 64: Also at Instituto de Astrofísica de Canarias, La Laguna, Spain
- 65: Also at Utah Valley University, Orem, U.S.A.
- 66: Also at University of Belgrade, Faculty of Physics and Vinca Institute of Nuclear Sciences, Belgrade, Serbia
- 67: Also at Facoltà Ingegneria, Università di Roma, Roma, Italy
- 68: Also at Argonne National Laboratory, Argonne, U.S.A.
- 69: Also at Erzincan University, Erzincan, Turkey
- 70: Also at Mimar Sinan University, Istanbul, Istanbul, Turkey
- 71: Also at Texas A&M University at Qatar, Doha, Qatar
- 72: Also at Kyungpook National University, Daegu, Korea