

## The CMS trigger system

---

**J. Varela**\*

*LIP - Lab. de Instrumentação e Física Exp. de Partículas*

*Av. Elias Garcia 14, 1000, Lisboa, Portugal*

*Also at IST - Inst. Superior Técnico, Univ. Técnica de Lisboa, Portugal, and CERN, Geneva, Switzerland*

*E-mail: joao.varela@cern.ch*

on behalf of the CMS Collaboration

This paper presents an overall description of the CMS trigger system which is composed by the Level-1 Trigger and the High-Level Trigger sub-systems. Details on the algorithms and selection criteria at a luminosity of  $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$  are provided. High efficiencies for benchmark physics channels are attainable with an inclusive selection that avoids detailed topological requirements on the event.

*European Physical Society*

*HEP2005 International Europhysics Conference on High Energy Physics*

*EPS (July 21<sup>st</sup>-27<sup>th</sup> 2005) in Lisboa, Portugal*

---

\* Speaker

## 1. Introduction

The CMS detector will operate in the LHC machine where the bunch-crossing frequency will be 40 MHz and the maximum luminosity will be  $10^{34} \text{ cm}^{-2}\text{s}^{-1}$  originating about 20 inelastic interactions per bunch crossing. The CMS trigger system has the enormous task of reducing the input data rate by a factor of the order of  $10^6$  to a rate of  $0(10^2)$  Hz adequate for permanent storage, preserving high efficiency for the interesting and rare physics events. In order to be open to new physics, the selection criteria are mostly based on inclusive production of high energy leptons or jets avoiding the use of topological or other event conditions specific of known physics. In these conditions the emphasis has to be placed on the reconstruction of trigger objects (e.g. electrons, muons) with high quality. To achieve the required high background rejection factors and to measure the trigger objects with enough precision sophisticated pattern recognition algorithms have to be implemented in dedicated hardware processors in the Level-1 Trigger and complete reconstruction algorithms with “off-line quality” need to run on the High-Level Trigger processors.

The performance of this system at both low luminosity,  $2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ , and at high luminosity,  $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ , was investigated using a full detector GEANT simulation [1]. Simulated digitization, including both in-time and out-of-time pileup, was performed at both luminosities.




Fig. 1: Data flow architecture of the CMS trigger and data acquisition system.

## 2. Overview of the CMS trigger and data acquisition system

The CMS trigger and data acquisition system design is shown in Fig. 1. The Level-1 trigger processor uses coarse grained detector data to decide whether to keep the event while the fine grained data is stored in pipeline memories in the detector front-ends. In case of a Level-1 accept (L1A), the readout builder network assembles complete events collecting event fragments in the front-ends and transferring them to a processing unit in the filter farm.

The CMS data acquisition system (DAQ) is designed to accept an input L1A rate of 100 kHz. A custom Level-1 processor is used to select this 100 kHz of events from the input 40 MHz bunch-crossing rate. The High-Level trigger is made in a farm of standard processors, after data readout through the event-builder switch. No dedicated Level-2 processor before the switch network is used.

A commercial processor farm for all selection beyond Level-1 allows to benefit from the evolution of computing technology, having the maximum flexibility for future upgrades of trigger algorithms. A further notable feature of the CMS DAQ system is its modularity: it is built up of eight 12.5 kHz units, not all of which need be installed at start-up.

## 3. The Level-1 trigger

The CMS Level-1 trigger uses coarse local data from the calorimeter and muon systems to make electron/photon triggers, jet and energy sum triggers, and muon triggers. The Level-1 trigger is a synchronous pipelined system working at the LHC clock frequency (40 MHz). The trigger latency, defined as the time needed for a decision and its propagation to the front-end detector electronics, is specified to be smaller than 3.2  $\mu$ s. The latency requirement is established to keep the size of the  $O(10^7)$  pipeline analog memories necessary in the CMS silicon tracker front-end electronics at a manageable level.

In order to cope with the large jet background to the lepton triggers expected at LHC, the CMS Level-1 system uses rather sophisticated pattern recognition trigger algorithms that allow high selectivity and steep efficiency thresholds curves. To achieve this goal a highly complex electronics systems was designed and built and it is now in the final installation and commissioning phases.

The system is organized in distinct and separate calorimeter and muon systems (see Fig.2). Each of these sub-systems has a similar vertical architecture that starts with the computation of trigger primitives (e.g. energy deposits in the calorimeter trigger towers or track segments identified in the muon stations). The trigger primitives are processed by regional triggers that implement the spatial pattern recognition algorithms (e.g. electromagnetic cluster identification or muon track reconstruction). The Level-1 muon trigger has a global muon trigger that combines the trigger information from the muon detectors trigger systems (Drift Tubes in the barrel, Cathode Strip Chambers in the endcaps and Resistive Plate Chambers). The resulting trigger objects (e.g. electron/photon candidate) are identified by the kinematical variables (transverse momentum,  $\eta$ - $\phi$  location), topological information (e.g. isolation criteria), and

dedicated quality bits. This information is used to sort the candidate trigger objects according to a rank classification that takes into account all the object variables. The four highest ranking objects of the various categories (electrons/photons, isolated electrons/photons, muons, tau-jets, barrel and forward jets) are combined in any one of 128 possible algorithms of the global trigger system.

The trigger primitives are implemented in about 5000 electronics boards of 7 different types installed on-detector and the regional and global trigger system is implemented in 45 9U-crates housing about 630 boards of 32 different types. The system has large flexibility provided by a large number of electronics programmable parameters and most importantly most algorithms are implemented in re-programmable FPGAs. In particular the logical combinations implemented in the global trigger algorithms can be re-defined.

The Level-1 trigger rejection is sufficiently large to reduce the Level-1 accept rate to a maximum of 100 kHz so that the data flow matches the event builder switch network bandwidth.

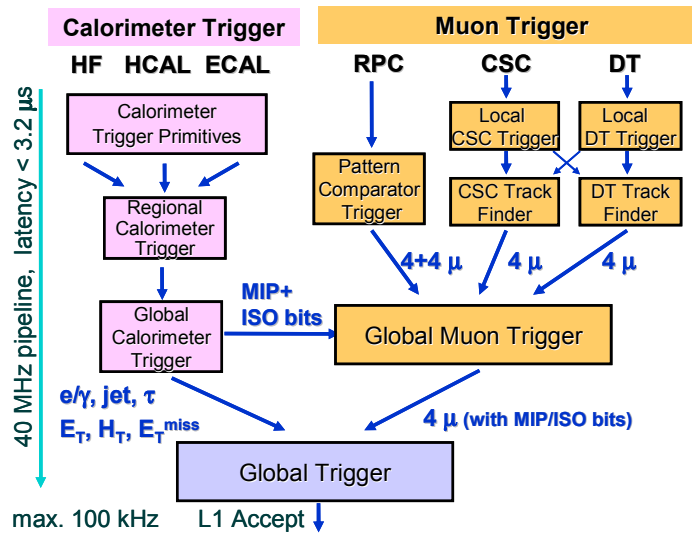


Fig. 2: Architecture and data flow in the synchronous pipelined Level-1 Trigger system.

### 3.1 The calorimeter trigger

The calorimeter trigger is based on trigger towers of size  $0.087 \times 0.087$  in  $\eta$ - $\phi$  space in the central region, and somewhat larger for  $|\eta| > 2$ . This size corresponds to a single readout tower in the hadron calorimeter (HCAL), and  $5 \times 5$  crystals in the electromagnetic calorimeter (ECAL) barrel. In the ECAL endcap crystals are mechanically arranged in a regular x-y array but are combined in trigger towers of about the same dimension in  $\eta$ - $\phi$  space.

The trigger tower energy sums are formed by the ECAL and HCAL Trigger Primitive Generator (TPG) circuits from the individual calorimeter cell energies. For the ECAL, these energies are accompanied by a bit indicating the transverse extent of the electromagnetic energy

deposit. For the HCAL, the energies are accompanied by a bit indicating the presence of minimum ionizing energy. FIR filters associated to peak finders are used in the TPG circuits to assign the trigger data to the correct bunch crossing. The filter coefficients are optimized as a function of the pulse shape and noise to achieve efficiencies larger than 99% for trigger tower energies above 1 GeV.

The TPG information is transmitted over high speed copper links to the Regional Calorimeter Trigger (RCT), which finds candidate electrons/photons, taus, and jets. The RCT separately finds both isolated and non-isolated electron/photon candidates. The RCT transmits the candidates along with sums of transverse energy to the Global Calorimeter Trigger (GCT). The GCT sorts the candidate electron/photons, taus, and jets and forwards the top 4 of each type to the global trigger. The GCT also calculates the total transverse energy and total missing energy vector. It transmits this information to the global trigger as well. The RCT also transmits an  $(\eta, \phi)$  grid of quiet regions to the global muon trigger for muon isolation cuts.




Fig. 3: Illustration of the Level-1 electron/photon sliding window algorithm used to identify electron/photon trigger objects. The transverse energy is given by the sum of the central tower with the highest of the four neighbors. Electron/photon candidates pass a cut on the H/E ratio in the central tower. Isolated electron/photon candidates have quiet E and H neighbor towers and have more than 90% of the energy localized in two adjacent crystal strips of the central tower. Strips are defined along  $\phi$  to allow for the energy spread due to the magnetic field.

The electromagnetic trigger works with fully overlapping windows of  $3 \times 3$  trigger towers, with the electron/photon candidate energy given by the sum of two adjacent ECAL towers. Cuts may be put on the hadronic/electromagnetic fraction, on electromagnetic or hadronic isolation, and on the fine-grain lateral shape in the ECAL. The four highest rank electron/photon objects and the four highest rank isolated electron/photon objects are transmitted to the global trigger. The requirements on the objects of a multi electron/photon trigger, namely the energy threshold, the cluster shape and isolation cuts and the  $(\eta, \phi)$  region, are set individually. Requirements on the  $(\eta, \phi)$  separation between objects can also be defined. The electron/photon algorithm is illustrated in Fig. 3.

Figure 4 (left) shows efficiency turn-on curves, for different threshold cuts, for isolated electron trigger as a function of electron  $p_T$ . Also shown (right) is the background rate as a function of the threshold on the isolated single electron trigger at low luminosity.




Fig. 4: (Left) Efficiency turn-on curves different threshold cuts, for isolated electron trigger as a function of electron  $p_T$  and (right) background rate as a function of the threshold on the isolated single electron trigger at  $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ .

The jet trigger is also based on  $3 \times 3$  windows, but for jets the elements of these windows are  $4 \times 4$  arrays of trigger towers, called a calorimeter trigger region, except in the Hadronic Forward (HF) region where it is the trigger tower itself.

Thus the jet algorithm sums transverse energy in a  $12 \times 12$  array of trigger towers, approximately corresponding to a unit square in  $\eta$ - $\phi$  space. The central region  $E_T$  is required to be higher than the eight neighbor region  $E_T$  values. In addition, the central region  $E_T$  is required to be greater than a fixed value. Separate lists are made of central jets and forward jets. Figure 5 shows the Level-1 jet trigger rates as a function of the threshold.

The tau-jet trigger is intended to capture single and three-prong decays of  $\tau$  leptons forming narrow clusters of energy deposits in the calorimeter. It requires a narrow 'tau-like' energy deposit shape in the central region and stringent isolation in the eight surrounding regions. A tau-like pattern is obtained if the active ECAL and HCAL towers (towers above a threshold) are confined within  $2 \times 2$  contiguous trigger towers of the central calorimeter region.

Three separate classes of jet - central, tau-jet, and forward - provide flexibility for the definition of combined triggers. Jets and tau-jets occurring in a calorimeter region where an electron is identified are not considered. The top four candidates of each class of calorimeter trigger are sent to the global trigger.

In addition counters of the number of jets above programmable thresholds in various  $\eta$  regions are provided to give the possibility of triggering on events with a large number of low energy jets. Jets in the forward and backward HF calorimeters are sorted and counted separately.

Missing  $E_T$  is computed from the sums of calorimeter region values of  $E_x$  and  $E_y$ , and the sum extends to the end of the forward calorimeter, i.e.  $|\eta| = 5$ . The total  $E_T$  is given by the scalar sum of the calorimeter regions  $E_T$ . The sum extends up to the end of forward calorimeter as well. The  $H_T$  trigger is defined as the scalar sum of the  $E_T$  of jets above a programmable threshold. This trigger is not as susceptible as the total  $E_T$  trigger to noise and pileup effects. The  $H_T$  trigger can capture high jet multiplicity events such as those from fully hadronic top decay, and hadronic decays of squarks and gluinos.




Fig. 5: Level-1 jet trigger rates for low and high luminosity.

### 3.2 The muon trigger

As the rate of real muons at LHC is huge the main issue for the muon trigger system is to provide an accurate measurement of the muon  $p_T$  allowing for the definition of sharp efficiency thresholds. The Level-1 muon trigger uses information from fast dedicated muon trigger detectors, resistive plate chambers (RPCs), complemented by the precise position measurements of the muon chambers drift tubes in the barrel and cathode strip chambers in the end-cap. This association provides a robust and redundant system that combines excellent efficiency for bunch crossing identification given by the RPC trigger with the precise  $p_T$  measurements given by the other muon trigger sub-systems. The bending in the successive layers of the iron yoke (which completes the magnetic circuit of the CMS field) is measured by first assembling local vectors in the measurement stations and then assembling tracks by linking these vectors across the iron.

Each of the Level-1 muon trigger systems has its own trigger logic. The Cathode Strip Chambers form Local Charged Tracks (LCT) from the cathode strips, which are combined with the anode wire information for bunch crossing identification. The LCT pattern logic assigns a  $p_T$  and quality bits to the local track segment which are used to sort the tracks. The top three LCTs in a CSC sector are transmitted to the CSC Track Finder, which combines the LCTs into full

muon tracks and assigns  $p_T$  values to them. The CSC and Drift Tube Track-Finders exchange track segment information in the region where the chambers overlap.

The Barrel Muon Drift Tubes are equipped with Bunch and Track Identifier (BTI) electronics that finds track segments from coincidences of aligned hits in 4 layers of one drift tube superlayer. The track segments positions and angles are sent to the Track Correlator, which attempts to combine the segments from the two Super Layers in a DT station measuring the  $\phi$  coordinate. The best combinations from all TRACOs of a single chamber together with the Super Layers  $\eta$  segments are collected and the best two segments (if found) are sent to the Track Finder. The Track Finder combines the segments from different stations into full muon tracks.

The RPC strips are connected to a Pattern Comparator Trigger (PACT), which is projective in  $\eta$  and  $\phi$ , and that combines up to 6 layers of RPC detectors to identify allowed muon patterns. Each pattern is associated to a given  $p_T$ .

The Global Muon Trigger sorts the RPC, DT and CSC muon tracks, and then attempts to correlate the CSC and DT tracks with RPC tracks. It contains logic to cancel ghost tracks that arise when a single muon is found by more than one muon system and is not otherwise matched. The GMT also correlates the found muon tracks with a  $\eta$ - $\phi$  grid of quiet calorimeter towers to determine if the muons are isolated. The final ensemble of muons are sorted based on their initial quality, correlation and  $p_T$  and then the 4 top muons are sent to the Global Trigger.




Fig. 6: Block diagram of the Level-1 muon trigger.






Fig. 7: Efficiency of Level-1 muon trigger as a function of  $\eta$ , for muons from  $W \rightarrow \mu\nu$ .

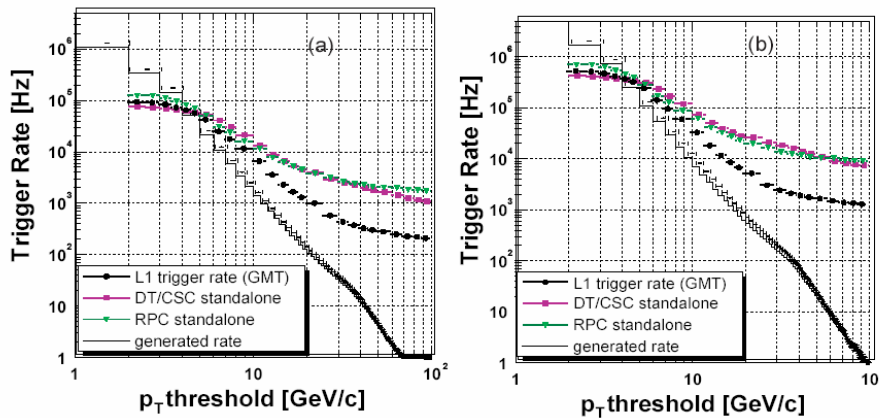


Fig. 8: Level-1 muon trigger rate as a function of  $p_T$  threshold for low (a) and high (b) luminosity.

The overall muon trigger scheme is illustrated in Fig. 6. Figure 7 shows the resulting muon trigger efficiency, as a function of  $\eta$ , for muons coming from W-boson decay. Figure 8 illustrates the expected muon trigger rates obtained from a full detector simulation.

### 3.3 The Global Level-1 Trigger

The Level-1 Trigger system sorts ranked trigger objects, rather than histogramming objects over a fixed threshold. This allows all trigger criteria to be applied and varied at the Global Trigger level rather than earlier in the trigger processing. All trigger objects are accompanied by their coordinates in  $\eta$ - $\phi$  space. This allows the Global Trigger to vary thresholds based on the location of the trigger objects. It also allows the Global Trigger to require trigger objects to be close or opposite from each other. In addition, the presence of the trigger object coordinate data in the trigger data, which is read out first by the DAQ after a L1A,

permits a quick determination of the regions of interest where the more detailed HLT analyses should focus.

The Global Level-1 Trigger transmits a decision to either accept (L1A) or reject each bunch crossing. This decision is transmitted by the Trigger Control System (TCS) to sub-detector frontends using a fiber optic distribution system. Fast commands for trigger synchronization and calibration triggers distributed in the LHC gaps are also available. The Trigger Throttling System (TTS) allows the rate reduction by prescaling or shutting off of L1A signals in case the detector readout or DAQ buffers are at risk of overflow.

### 3.3 Level-1 trigger table

In order to construct a complete table of the Level-1 selection it is necessary to allocate the available DAQ bandwidth between the various triggers. The full design bandwidth can accommodate 100 kHz of 1 MB events, however the CMS plan is to use the flexibility of the modular DAQ system and at startup install only a 50 kHz capacity.

In the allocation optimization a safety factor of three is taken to account for simulation uncertainties and unexpected backgrounds. Thus 16 kHz is allocated for low luminosity running and 33 kHz for high luminosity. The optimization will only be completed when real data is taken. Presently an equal sharing of rate to four classes of trigger: electron/photon triggers, muon triggers, tau-jet triggers, and jets and missing energy triggers, is allocated. The priority in this allocation has been to guarantee discovery physics while at the same time maintaining a sufficiently wide and general suite of channels so as to remain inclusive and be open to unexpected physics. The Level-1 trigger table for low luminosity is shown in Table 1.

Table 1: Level-1 Trigger table for  $2 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$ . Thresholds correspond to values with 95% efficiency. The combined rate for the three different jet triggers is given on a single line, but the three thresholds are shown (the two jet trigger is found to be redundant).

Trigger	Threshold (GeV or GeV/c)	Rate (kHz)	Cumulative Rate (kHz)
Inclusive isolated electron/photon	29	3.3	3.3
Di-electrons/di-photons	17	1.3	4.3
Inclusive isolated muon	14	2.7	7.0
Di-muons	3	0.9	7.9
Single tau jet trigger	86	2.2	10.1
Two tau jets	59	1.0	10.9
one jet, three jets, four jets	177, 86, 70	3.0	12.5
Jet * $E_T^{\text{miss}}$	88 * 46	2.3	14.3
Electron * Jet	21 * 45	0.8	15.1
Minimum bias (calibration)		0.9	16.0
TOTAL			16.0

## 4 The High-Level Trigger

The CMS High-Level trigger runs on a farm of mass-market processors using code that is as close as possible to offline code. The final output rate of the HLT must remain manageable, and the target rate is taken as  $0(10^2)$  Hz.

Various strategies guide the development of the HLT code. Regional reconstruction and reconstruction on demand strategies are used: rather than reconstruct all possible objects in an event, whenever possible only those objects and regions of the detector that are needed are reconstructed. Events are to be discarded as soon as possible, this leads to the idea of partial reconstruction, and also to the development of virtual 'trigger levels': at Level 2 calorimeter and muon trigger information is used, Level 2.5 is the term used to describe the additional use of tracker pixel information, and Level 3 refers to the use of the full event information including the complete tracker.

### 4.1 Electron and photon trigger

The first step of the HLT selection process for electrons is the reconstruction of clusters in the ECAL matched to the Level-1 electron/photon triggers using its full granularity. The key issue here is the recovery of the energy radiated as bremsstrahlung in the tracker and spread in  $\varphi$  due to the bending of the electrons in the 4 T magnetic field. The energy is collected in clusters of clusters, termed super-clusters. An  $E_T$  threshold is applied to the reconstructed super-clusters.

The Level-1 electron and photon trigger rate is entirely dominated by the decay of neutral hadrons in jets (mainly  $\pi^0$ 's) to photons. The most important step in the electron selection comes at Level 2.5 where super-clusters are propagated back in the magnetic field from the ECAL to the pixel detector layers and matching hits are sought. Searching for two matching hits, out of three possible, within a small region, provides a large rejection factor with only a small efficiency loss. The unmatched clusters become photon candidates, the rate of which is reduced by much higher threshold cuts than are used in the electron channels.

The electron and photon rates output by the HLT at low luminosity, broken down by contribution, are listed in Table 2. A loose calorimetric isolation has been applied to the photon streams, but no isolation (beyond that of the Level-1 Trigger) has been applied to the electron streams. To control the two-photon rate the thresholds have been raised to  $E_T^1 > 40$  GeV,  $E_T^2 > 25$  GeV (equal to the final offline cuts envisaged for  $H \rightarrow \gamma\gamma$ ). These cuts reduce the rate from 11 Hz to 5Hz, and has a negligible effect on the efficiency.

Table 2: Electron and photon stream output from HLT selection at a luminosity of  $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ 

	Signal	Background	Total
Single electron	$W \rightarrow e\nu$ : 10Hz	$\pi^\pm/\pi^0$ overlap: 5Hz $\pi^0$ conversions: 10Hz $b/c \rightarrow e$ : 8Hz	33Hz
Double electron	$Z \rightarrow ee$ : 1Hz	$\sim 0$	1Hz
Single photon	2Hz	2Hz	4Hz
Double photon	$\sim 0$	5Hz	5Hz
TOTAL:			43Hz

## 4.2 Muon trigger

The muon selections works by successive refinement of the muon  $p_T$  measurement. At Level 2 the muons are reconstructed in the muon system alone, with the additional requirement that the track segments have a valid extrapolation to the interaction region. The  $p_T$  resolution obtained for muons from W decays is 10% in the region  $|\eta| < 0.8$ , and between 15% and 16% for the remaining fiducial region ( $0.8 < |\eta| < 2.1$ ).

At Level 3 full track reconstruction, including the inner tracker, is used. Starting from the regional seeds, a track reconstruction algorithm based on the Kalman filter technique is used to reconstruct tracks within the selected regions of interest. The gain in momentum resolution is substantial: for muons from W decays the  $p_T$  resolution is 1.0% in the region  $|\eta| < 0.8$ , 1.4% for ( $0.8 < |\eta| < 1.3$ ) and 1.7% for ( $1.3 < |\eta| < 2.1$ ). The algorithmic efficiency for the Level-3 muon tracking is typically 99%, except in the pseudorapidity interval  $0.8 < |\eta| < 1.2$  where the drift tube and cathode strip chamber systems overlap and the efficiency is about 97%. Isolation cuts can be used to suppress muons from b, c, K and  $\pi$  decays.

## 4.3 Tau trigger

The High Level Trigger algorithms for tau identification are designed to be used in the selection of isolated  $\tau$ 's such as those expected in the MSSM Higgs decays  $A/H \rightarrow \tau^+ \tau^-$  and  $H^\pm \rightarrow \tau \nu$ . The final-state signatures involve events with a lepton plus a taujet, two tau jets or only one tau jet. For taujets with  $E_T > 50$  GeV about 90% of the energy is contained in a very small region in  $\eta$ - $\phi$  space of radius 0.15 to 0.20, and about 98% in a radius of 0.4.

At Level 2 rejection of background to hadronic tau decays is obtained by looking for very narrow jets in the calorimeters ( $\Delta R = 0.13$ ) surrounded by an isolation region ( $\Delta R = 0.4$ ). Both pixel isolation and full track isolation can be used to tighten the selection.

#### 4.4 Jets and missing $E_T$ trigger

Global jet finding is done using a simple iterative cone algorithm. In this algorithm, a list of towers is made, and a "protojet" is formed using the direction of the tower from the list with the highest  $E_T$  (the "seed tower") as the protojet direction. The direction of the protojet is calculated from the transverse-energy-weighted angles of the towers in a cone around the protojet direction in  $\eta$ - $\phi$  space, and the transverse energy of the protojet is calculated using the direction of the protojet and the sum the energies of the towers in the cone. The direction of the protojet is used to seed a new protojet. The procedure is repeated until the energy of the protojet changes by less than 1% between iterations and the direction of the protojet changes in  $\eta$ - $\phi$  space by less than 0.1, or until 100 iterations is reached.

To identify neutrinos in the HLT, the calorimeter information is used to look for missing transverse energy ( $E_T^{\text{miss}}$ ). The current algorithm calculates  $E_T^{\text{miss}}$  as a simple vector sum of the towers over a threshold of 500 MeV.

#### 4.5 HLT trigger table and performance summary

The cuts and thresholds described in the preceding sections must be chosen to provide a final physics selection. The known discovery channels provide guidance, but the selection should remain sufficiently inclusive. Table 3 shows the current set of thresholds and the corresponding rates to storage and provide an indication of the kind of event mixture that an output rate of  $0(10^2)$  Hz at a luminosity of  $2 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$  would yield. Table 4 shows the efficiency of the selection for some representative channels. The values shown include the effect of both the Level-1 trigger and the HLT. The numbers give the efficiency for selecting fiducial objects.

A key issue for the High-Level Trigger selection is the CPU power required for the execution of the algorithms. The time taken by the selection algorithms has been measured on a Pentium-III 1GHz processor, and the results vary from a very fast 50 ms for jet reconstruction to the longer 700ms for muon reconstruction. On average we obtain a mean of 271 ms per event passing the Level-1 triggers specified in Table 1.

Taking the start-up scenario of a DAQ system capable of reading a maximum of 50 kHz of events accepted by the Level-1 trigger, the average of 271 ms per event translates to 15,000 Pentium-III 1GHz CPUs. Assuming a factor of eight increase in computing power yields 40 ms per event, and a need for 2,000 CPUs. This figure comfortably matches our target estimate of 1,000 dual-CPU PCs for the HLT farm.

Table 3: Set of thresholds and the corresponding rates to storage at a luminosity of  $2 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$ .

Trigger	Threshold (GeV or GeV/c)	Rate (Hz)	Cumulative Rate (Hz)
Inclusive electron	29	33	33
Di-electrons	17	1	34
Inclusive photons	80	4	38
Di-photons	40, 25	5	43
Inclusive muon	19	25	68
Di-muons	7	4	72
Inclusive $\tau$ jets	86	3	75
Di- $\tau$ jets	59	1	76
1 jet * $E_T^{\text{miss}}$	180 * 123	5	81
1 jet OR 3 jets OR 4 jets	657, 247, 113	9	89
Electron * Jet	19 * 45	2	90
Inclusive b jets	237	5	95
Calibration etc. (10%)		10	105
<b>TOTAL</b>			<b>105</b>

Table 4: Efficiency for typical physics channels to pass the complete Level-1 and HLT selection (geometric acceptance factors are not included here: the selected physics objects are within the detector fiducial regions).

Channel	Efficiency
$H(115 \text{ GeV}/c^2) \rightarrow \gamma\gamma$	77%
$H(160 \text{ GeV}/c^2) \rightarrow WW^* \rightarrow \mu\nu\mu\nu$	92%
$H(150 \text{ GeV}/c^2) \rightarrow ZZ^* \rightarrow \mu\mu\mu\mu$	98%
$A/H(200 \text{ GeV}/c^2) \rightarrow \tau\tau$	45%
SUSY ( $\sim 0.5 \text{ TeV}/c^2$ sparticles)	$\sim 60\%$
With RP-violation	$\sim 20\%$
$W \rightarrow e\nu$	42%
$W \rightarrow \mu\nu$	69%
Top $\rightarrow \mu X$	72%

## 5 Conclusions

The CMS experiment developed an innovative trigger system based on two trigger levels to perform the difficult task of triggering on new interesting events produced in high-energy proton-proton collisions at LHC.

We have shown that selection criteria mostly based on the inclusive production of high energy leptons or jets avoiding the use of topological or other event conditions specific of known physics provide high efficiencies for benchmark physic channels and reduce the collision rate by six orders of magnitude as required by the available storage capacity.

The Level-1 trigger algorithms are implemented in one of the most complex electronics systems ever built for a high-energy physics experiment. Dedicated FPGA based processors implementing a synchronous and pipelined architecture still provide large flexibility for new upgrades in the future if required by unexpected operation conditions or different physics requirements.

The High-Level triggers are implemented in software running in a farm of standard processors. This solution provides the maximum flexibility for adjusting to unforeseen circumstances resulting from bad beam conditions, high background levels or new physics channels not previously studied.

A possible trigger table for the Level-1 and the High-Level Trigger selection at a luminosity of  $2 \times 10^{33} \text{cm}^{-2} \text{s}^{-1}$  assuming a total DAQ bandwidth of 50 kHz was developed and validated by full detector simulation. The overall CPU requirement is approximately 300 ms per event on an Intel 1 GHz Pentium-III CPU.

## References

- [1] CMS Collaboration, *Data Acquisition & High-Level Trigger Technical Design Report*, CERN/LHCC 2002-26, 15 December 2002