Eur Phys J C **34**, s01, s151–s159 (2004)
Digital Object Identifier (DOI) 10.1140/epjcd/s2004-04-016-8

EPJ C direct
electronic only

The CMS trigger system

C. Seez, on behalf of the CMS collaboration

Imperial College of Science Technology and Medicine, London SW7, UK, e-mail: Chris.Seez@cern.ch

Received: 4 August 2003 / Accepted: 4 August 2003 /
Published Online: 13 July 2004 –

Abstract. The CMS trigger system must reduce an input data rate from the LHC bunch-crossing frequency of 40 MHz to a rate which will be written to permanent storage. A detailed study has recently been made of the performance of this system. This paper presents key elements of the results obtained and gives details of a draft "trigger table" for the Level-1 Trigger and the High-Level Trigger selection at a "start-up" luminosity of 2×10^{33} cm^{-2}s^{-1}. High efficiencies for most physics objects are attainable with a selection that remains inclusive and avoids detailed topological or other requirements on the event.

1 Introduction

The CMS experiment will operate a general purpose detector in the LHC machine where the bunch-crossing frequency will be 40 MHz and the luminosity will range from around 10^{33} cm^{-2}s^{-1} to the design luminosity of 10^{34} cm^{2}s^1, at which luminosity there will be about 20 inelastic interactions per bunch crossing. The CMS trigger system has the formidable task of reducing this input data rate to a rate of O(10^2) Hz which will be written to permanent storage.

A detailed study has recently been made of the performance of this system at both low luminosity, 2×10^{33} cm^{-2}s^{-1}, and at high luminosity, 10^{34} cm^{2}s^1 [1]. This study involved the full detector simulation of more than 7 M events using GEANT 3. Simulated digitization, including both in-time and out-of-time pileup, was performed at both luminosities for much of this sample. Digitization and reconstruction were done within the CMS OO environment with C++ code. This paper presents key elements of the results of this study.

2 CMS trigger and data acquisition system design

The CMS data acquisition system (DAQ) is designed to accept an input rate of 100 kHz events having a size of 1 MB. The trigger system uses a custom Level-1 processor to select this 100 kHz of events from the input 40 MHz bunch-crossing rate. During the 3μs latency of the Level-1 trigger the event data is stored in front-end pipelines. The remaining selection process is made in a farm of standard commercial processors, on data after it has been read out through the event-builder switch network. The CMS design is illustrated in Fig. 1 and compared to a more conventional architecture with a dedicated Level-2 processor before the switch network. By using a processor farm for all selection beyond Level-1 CMS is able to benefit maximally from the evolution of computing technology. Flexibility is maximized since there are no built-in design or architectural limitations; there is complete freedom in what data to access and in the sophistication of algorithms. Evolution is possible, allowing response to unforeseen backgrounds. The minimization of in-house elements has benefits both in terms of cost and maintainability.

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. A decision is returned to the front-end detector electronics after a latency of about 3 μs, of which nearly 2 μs is taken by transmission delay. It is required that the rejection is sufficiently large to reduce the Level-1 accept rate so that the data flow matches the switch network bandwidth. The hardware is custom made, largely using ASICs, but with widespread use of FPGAs where appropriate. It is organized in distinct and separate calorimeter and muon systems, and the results are combined and organized in a global trigger, where the final binary accept/reject decision is made.

3.1 The calorimeter trigger

The calorimeter trigger is based on trigger towers of size 0.087×0.087 in $\eta-\varphi$ space in the central region, and somewhat larger for $|\eta| > 2$. This size represents a single readout tower in the hadron calorimeter (HCAL), and

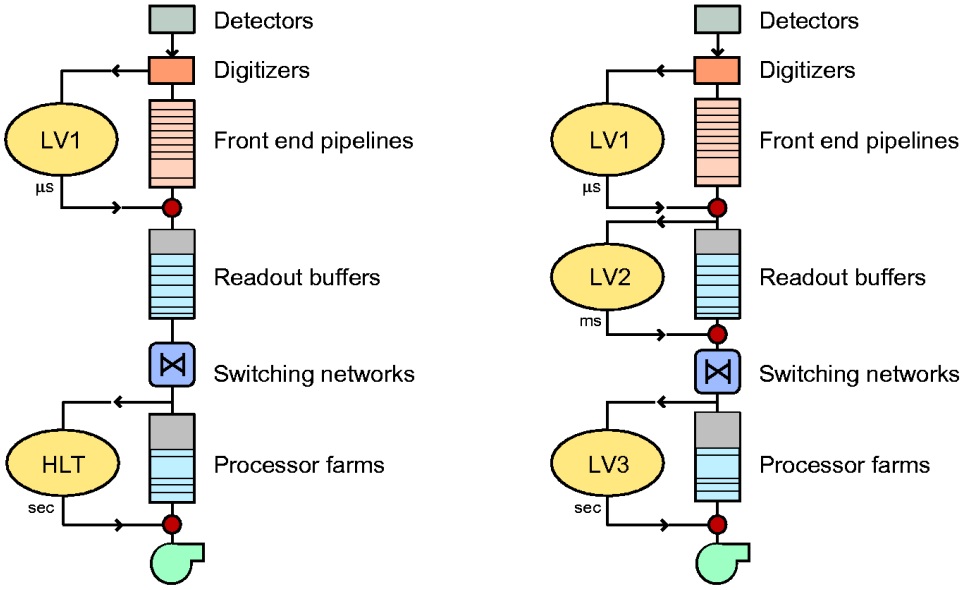

Fig. 1. Data flow architecture of the CMS trigger and data acquisition system (*on the left*), compared to a more conventional architecture using a dedicated Level-2 trigger

5×5 crystals in the electromagnetic calorimeter (ECAL) barrel. The tower energy sums are formed by the ECAL, HCAL and forward hadron calorimeter (HF) trigger primitive generator 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.

The electromagnetic trigger works with fully overlapping windows of 3×3 trigger towers, applying a threshold to the sum of two adjacent ECAL towers. Cuts may be put on isolation, on the hadronic/electromagnetic fraction, and on the fine-grain lateral shape in the ECAL (which acts as a sort of local isolation). Figure 2 (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.

The jet trigger is also based on 3×3 windows, but for jets the elements of these windows are 4×4 arrays of trigger towers. Thus the jet algorithm sums transverse energy in a 12×12 array of trigger towers, approximately corresponding to a unit square in $\eta - \varphi$ space. Separate lists are made of central jets and forward jets. The taujet trigger is obtained by demanding a narrow 'tau-like' shape in the central region. This is illustrated in Fig. 3. Single, double, triple and quad jet triggers are possible. The three separate classes of jet — central, tau-jet, and forward — provide flexibility for the definition of combined triggers.

The top four candidates of each class of calorimeter trigger are sent to the global trigger. Figure 4 shows the Level-1 jet trigger rates as a function of the threshold.

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. $|\eta| = 5$.

3.2 The muon trigger

The Level-1 muon trigger receives 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.

Each of the Level-1 muon trigger systems has its own trigger logic. The RPC strips are connected to a Pattern Comparator Trigger (PACT), which is projective in η and φ. 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 on a Trigger Motherboard. The Barrel Muon Drift Tubes are equipped with Bunch and Track Identifier (BTI) electronics that find track segments from coincidences of aligned hits in four layers of one drift tube superlayer.

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. Finally the information from the three types of muon detector are combined and the four best muon candidates identified and sent to the global trigger. The overall scheme is illustrated in Fig. 5. Figure 6 shows the resulting muon trigger efficiency, as a function of η, for muons coming from W-boson decay.

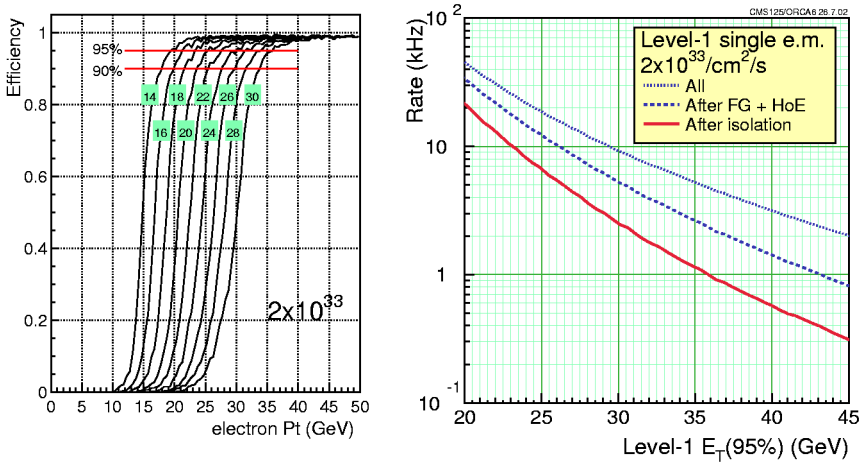

Fig. 2. (*Left*) Efficiency turn-on curves different threshold cuts, for isolated electron trigger as a function of electron p_T, and (*right*) background rate as a function of the threshold on the isolated single electron trigger at 2×10^{33} cm^{-2}s^{-1}

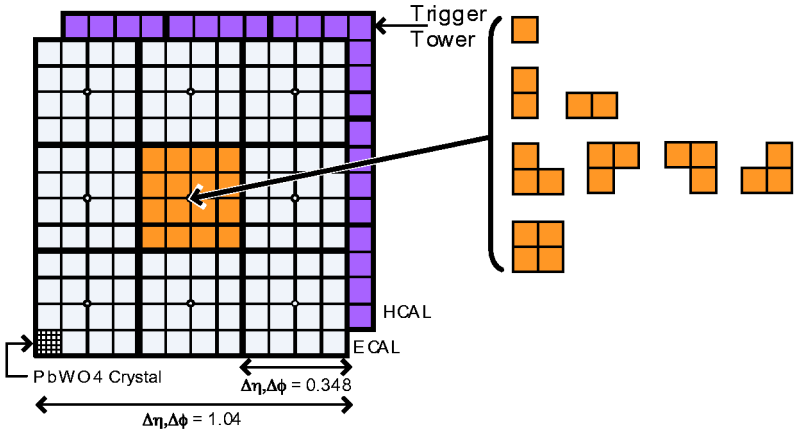

Fig. 3. Illustration of the Level-1 jet trigger algorithm, showing the 'tau-like' shapes demanded in the centre region for the taujet trigger

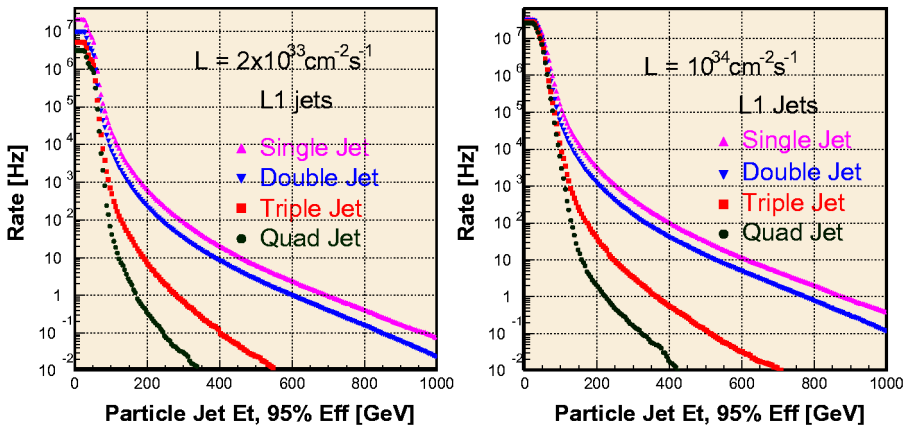

Fig. 4. Level-1 jet trigger rates for low and high luminosity

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 un-

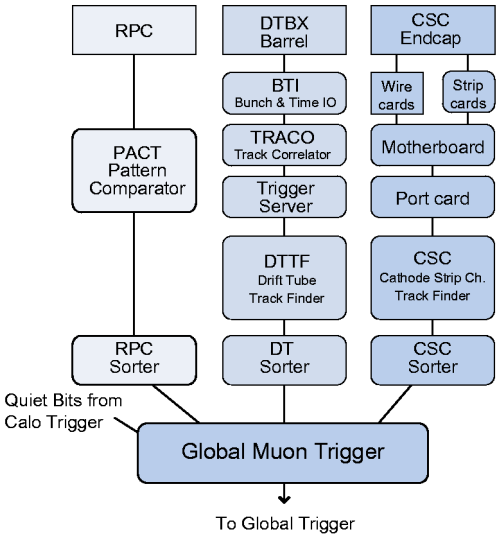

Fig. 5. Block diagram of the Level-1 muon trigger

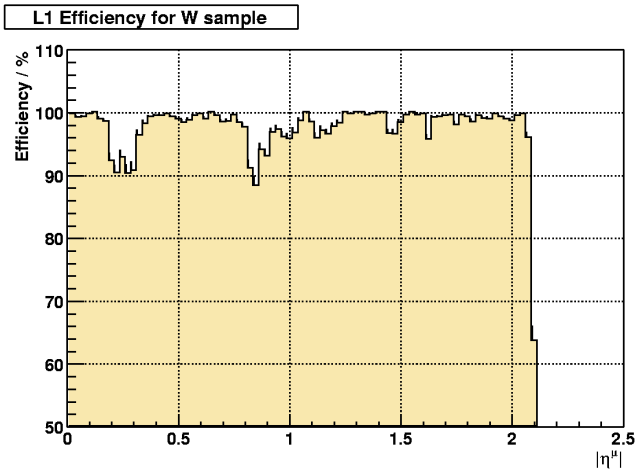

Fig. 6. Efficiency of Level-1 muon trigger as a function of η, for muons from W$\rightarrow \mu\nu$

expected backgrounds. Thus 16 kHz is allocated for low luminosity running and 33 kHz for high luminosity. The optimization, which is an iterative process which will only be completed when real data is taken, is begun by allocating an equal sharing of rate to four classes of trigger: electron/photon triggers, muon triggers, tau-jet triggers, and jets and missing energy triggers. Then the rate must be shared between triggers within the classes — for example between single and double triggers. The priority in this allocation procedure 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, which has been arrived at by the procedures described and satisfies the constraints outlined, is shown in Table 1.

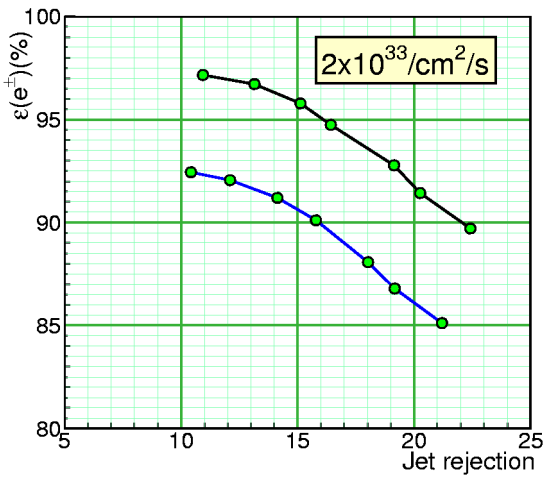

Fig. 7. Efficiency to pass electrons versus rejection of jet background, at 2×10^{33} cm^{-2}s^{-1} using the Level-2.5 pixel matching. The upper curve is the performance with the full pixel detector system; the lower curve refers to a reduced (staged) startup installation which was considered and rejected. The different points on the lines represent refer to different sizes of window used for the search area

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. This strategy eases maintenance and allows offline code development to be rapidly exploited in the trigger. The final output rate of the HLT must remain manageable, and the target rate is taken as O(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.

The HLT for a given event runs on a single processor, which deals with a single event at a time. It has access to the full event information, the full granularity and resolution is available, and the only limitations are the CPU time usage, the limited output rate and the imprecision of the calibration constants available online. To satisfy the physics requirements of the experiment, the selection must be efficient, it must be sufficiently inclusive to be sensitive to unexpected physics, and it must not rely on a very precise knowledge of run conditions or calibration. Monitoring of the HLT performance, the algorithms and the processors, is another important issue that merits careful attention.

Table 1. Level-1 Trigger table for $2\times10^{33}\mathrm{cm}^{-2}\mathrm{s}^{-1}$. 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^{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

4.1 HLT for electrons and photons

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. The spray of energy extends in φ beyond the boundaries of a single shower 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 π^0s) to photons. The most important step in the electron selection comes at Level 2.5 where super-clusters are propagated back in the field from the ECAL to the pixel detector layers and matching hits are sought. The pixel layers are situated just outside the beam-pipe before most of the tracker material and hence before most electrons have radiated significantly and before photons have had a large probability to convert. Searching for two matching hits, out of three possible, within a small region, provides a large rejection factor with only a small efficiency loss (see Fig. 7). The unmatched clusters become photon candidates, the rate of which are 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 11Hz to 5Hz, and has a negligible effect on the efficiency. A fully optimized selection would probably involve track isolation on the photon streams (wholly or partly replacing the calorimetric isolation and the raised threshold) and track isolation in the single electron stream. This selection would reduce the total rate to about 26 Hz, of which only half is background, with the introduction of only a small further inefficiency.

4.2 HLT for muons

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 (defined by the beam spot size: σ_{xy}=15 μm and σ_z=5.3 cm). The momentum estimate from the Level-1 muon trigger is used initially for propagation in the magnetic field. 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 π decays. Three isolation techniques have been studied: calorimeter isolation, which can be applied at

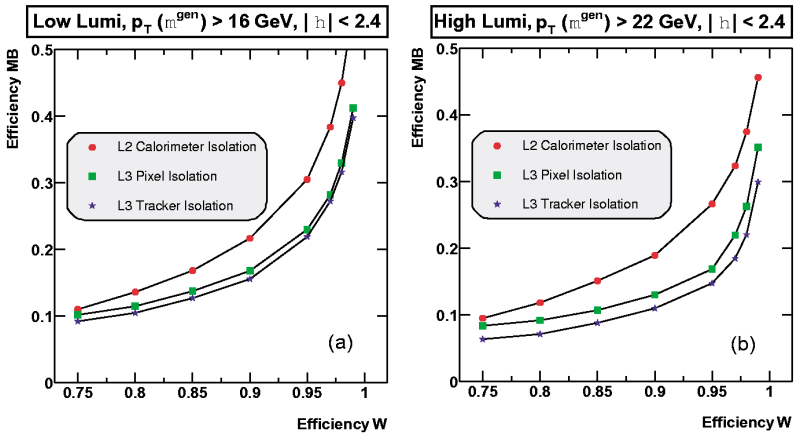

Fig. 8. Efficiency of the three isolation algorithms on the 'reference background' muons as a function of efficiency for signal muons from W decay at **a** low and **b** high luminosity

Table 2. Electron and photon stream output from HLT selection at a luminosity of 2×10^{33} cm^{-2}s^{-1}

	Signal	Background	Total
Single electron	W→eν: 10Hz	$\pi^{\pm}/\pi^0$ overlap: 5Hz π^0 conversions: 10Hz b/c→e: 8Hz	33Hz
Double electron	Z→ee: 1Hz	~0	1Hz
Single photon	2Hz	2Hz	4Hz
Double photon	~0	5Hz	5Hz
TOTAL:			43Hz

Table 3. Set of thresholds and the corresponding rates to storage at a luminosity of 2×10^{33} cm^{-2}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 τ jets	86	3	75
Di-τ jets	59	1	76
1 jet * E_T^{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
TOTAL			**105**

Level 2; pixel isolation, using track stubs reconstructed in the pixel detector; and isolation using fully reconstructed tracks. In all three techniques, jet activity is sought in a circular region around the muon in $\eta - \varphi$ space — for all three techniques the optimum size of this region is found to have a radius, ΔR, between 0.2 and 0.3. Figure 8 shows the efficiency for background versus the efficiency for signal (muons from W decays) for these isolation techniques at both low and high luminosity.

Table 4. Efficiency for typical physics channels to pass the complete Level1 and HLT selection (geometric acceptance factors are not included here: the selected physics objects are within the detector fiducial regions)

Channel	Efficiency
H(115 GeV/c^2)→γγ	77%
H(160 GeV/c^2)→WW* →μνμν	92%
H(150 GeV/c^2)→ZZ*→μμμμ	98%
A/H(200 GeV/c^2)→ττ	45%
SUSY (~0.5 TeV/c^2 sparticles)	~60%
With RP-violation	~20%
W→eν	42%
W→μν	69%
Top→μ X	72%

4.3 HLT for taus

The High Level Trigger algorithms for τ identification are designed to be used in the selection of isolated τ's such as those expected in the MSSM Higgs decays A/H→ $\tau^{\rm r}\tau^{\rm t}$ and $H^{\pm} \rightarrow \tau\nu$. The final-state signatures involve events with a lepton plus a tau jet, two tau jets or only one tau jet. The τ's decays hadronically 65% of the time, producing a narrow jet containing a relatively small number of charged and neutral hadrons. For tau jets with $E_T > 50$ GeV about 90% of the energy is contained in a very small region in η, φ 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 (ΔR=0.13) surrounded by an isolation region (ΔR=0.4). Both pixel isolation and full track isolation can be used to tighten the selection.

4.4 HLT for jets and missing E_T

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 η, φ 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 ($E_T^{\rm miss}$). The current algorithm calculates E_T^{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. In principle there are difficult choices to be made in optimizing the use of available bandwidth: choices must be made between unlike channels; the purity of channels is in many cases quite dissimilar (e.g. the muon streams contain little background whereas the single electron stream tends to contain more background) and this needs to be balanced against the complementarity of different streams. 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 $O(10^2)$ Hz at a luminosity of 2×10^{33} cm^{-2}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 Level1 trigger and the HLT. The numbers give the efficiency for selecting fiducial objects.

4.6 CPU time usage of HLT selection

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. The CPU needs of the algorithms must be weighted by the frequency of their application, which is the Level-1 rate of the corresponding channel. A total of 4092 seconds is needed to cover the 15.1 kHz of events allocated in the Level-1 trigger table (Table 1), as shown in Table 5, and corresponds to a mean of 271 ms per event passing Level 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 CPUs such as are currently available in a standard commercial Personal Computer (PC). It is not believed that a detailed extrapolation of these figures to the year 2007 would be sufficiently reliable to justify the effort required to make it, but an estimate can be made assuming Moore's Law, i.e. the doubling of CPU power every 18 months. 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.

There are, naturally, large uncertainties in the above estimate, listed in order of importance:

– A major uncertainty is related to all the samples that have not been simulated. Inherent in the above estimate is the assumption that all events in the 50 kHz of events accepted by the Level-1 Trigger will require the full 300 ms. However, only 16 kHz of the total of 50 kHz has really been simulated, with the rest of the events being included in the "safety factor" in the allocation of the Level-1 Trigger bandwidth. It is clear

Table 5. Summary of CPU time required for the selection of each physics objects in the HLT. The CPU figures refer to a 1 GHz Intel Pentium-III CPU

Physics Object	CPU time per Level-1 event (ms)	Level-1 Trigger rate (kHz)	Total CPU time (s)
Electrons/photons	160	4.3	688
Muons	710	3.6	2556
Taus	130	3.0	390
Jets and E_T^{miss}	50	3.4	170
Electron + Jet	165	0.8	132
B-jets	300	0.5	150

that much of this safety factor will be used, since basic processes from the beam in the machine, e.g. beam halo, have not been included. It is expected, however, that such processes will require far less CPU to identify and reject.

- The figure of ∼300 ms per event accepted by the Level-1 Trigger does not include all the processing needed to go from the actual raw data to the final decision by the HLT. The overhead from the software framework for accessing the DAQ packets of raw data, and the time needed to unpack the raw data in the form needed by the reconstruction algorithm are not included. These times cannot be measured at present, since the raw data formats are not finalized. First estimates suggest that the CPU requirement can be significant, taking up to one quarter of the total reconstruction time.
- The selection of Table 3 is only an example of the type of requirements and rates involved in the HLT selection, and the actual trigger table will very likely include additional selections. In particular, it is expected that more selections involving combined objects will be introduced, which will increase the required CPU time.
- The time needed to process a single combination depends on the actual implementation of the reconstruction algorithm and the underlying libraries, e.g. the mathematical and memory allocation ones. It also depends on the quality of the compiler. In these respects the current CPU time estimates are an upper limit. The current time estimates are dominated by the muon reconstruction step which is in turn dominated by the GEANE [2] extrapolation step. A significant speed-up is expected when the extrapolation package is upgraded. Moreover, the current package used for matrix algebra will be replaced with a better one, and specialized memory allocators will be used for cases where the memory management overhead has been measured to be significant.
- There are uncertainties arising from the actual occupancy in the CMS sub-detectors. The time needed to perform a given pattern recognition and/or reconstruction task depends on the number of elements or combinations that need to be considered in order to reach a given reconstruction efficiency, as well as on the time needed to treat each element. The minimum number of combinations needed to achieve a given efficiency depends on the occupancy of the detectors, and on the accuracy of each measurement. The occupancy, in turn, has uncertainties arising from the event generators description of the hard interaction, the noise level in the detector, including the full readout chain, the neutron background in the detector, and finally from beam backgrounds. The accuracy of each measurement has an uncertainty related to the current understanding of the detector and accuracy of the detector response simulation knowledge of the calibration and alignment at HLT time. All these uncertainties have reasonable estimates or are expected to be sufficiently small to be neglected in the present studies.

The CMS HLT reconstruction algorithms underwent rapid development in the context of the DAQ TDR, and this development will continue. It is expected that the CMS reconstruction software will continue to evolve considerably between now and the first physics run. The HLT filtering step should benefit from all improvements, since the code used in the HLT, with the exception of the input and output mechanisms, is identical to the code to be used offline.

5 Summary

A draft "trigger table" for the Level-1 Trigger and the High-Level Trigger selection at a "start-up" luminosity of 2×10^{33} cm^{-2}s^{-1} has been shown. The assumption of this table is a total DAQ bandwidth of 50 kHz. High efficiencies for most physics objects are attainable with a selection that remains inclusive and avoids detailed topological or other requirements on the event. The overall CPU requirement of this selection is approximately 300 ms on a an Intel 1 GHz Pentium-III CPU.

Much more sophisticated trigger requirements can, and most likely will, be employed. As an example, at a minimum, as the instantaneous luminosity drops throughout a fill of the LHC, some bandwidth will be freed from the triggers discussed here. This additional bandwidth can be reallocated to the same triggers by decreasing the thresholds.

The additional bandwidth may also be used in introducing new triggers, e.g. for B-physics. Introduction of such triggers is then purely an issue of whether there are adequate CPU resources for the selection of the relevant events. The systematic optimization of the track reconstruction code and the extensive use of "regional" and "conditional" track reconstruction allow for the very fast search for and the full reconstruction of B-meson decays. Further, ongoing optimization of the tracking code indicates that it can be applied to the full Level-1 event rate at both low and high luminosity. This would extend and complement the current "Level-2" selections.

The selection presented in this paper indicates that it is possible to provide the HLT selection of 1:1000 in a single processor farm. Furthermore, the full event record is available, and the software that implements all algorithms can be changed and extended. The CMS HLT system has great flexibility and provides much room both for improving the selection of the various physics channels, as well as for adjusting to unforeseen circumstances resulting from bad beam conditions, high background levels or new physics channels not previously studied.

References

1. CMS Collaboration: Data Acquisition & High-Level Trigger Technical Design report, CERN/LHCC 2002-26, 15 December 2002
2. V. Innocente, M. Maire, and E. Nagy: CERN program library, IT-ASD W5013, 1991