CMS Conference Report

1 December 2003

Online Event Selection at the CMS experiment

M. Konecki

University of Basle, Switzerland

CMS collaboration

Abstract

Triggering in the high-rate environment of the LHC is a challenging task. The CMS experiment has developed a two-stage trigger system. The Level-1 Trigger is based on custom hardware devices and is designed to reduce the 40 MHz LHC bunch-crossing rate to a maximum event rate of 100 kHz. The further reduction of the event rate to $O(100\text{ Hz})$, suitable for permanent storage, is performed in the High-Level Trigger (HLT) which is based on a farm of commercial processors. The methods used for object identification and reconstruction are presented. The CMS event selection strategy is discussed. The performance of the HLT is also given.

Presented at *9th International Conference on B-Physics at Hadron Machines - BEAUTY 2003*,
Pittsburgh, Pennsylvania, USA, October 14 – 18, 2003.

Online Event Selection at the CMS experiment

Marcin Konecki

on behalf of the CMS Collaboration

University of Basle, Switzerland

Abstract. Triggering in the high-rate environment of the LHC is a challenging task. The CMS experiment has developed a two-stage trigger system. The Level-1 Trigger is based on custom hardware devices and is designed to reduce the 40 MHz LHC bunch-crossing rate to a maximum event rate of 100 kHz. The further reduction of the event rate to O(100 Hz), suitable for permanent storage, is performed in the High-Level Trigger (HLT) which is based on a farm of commercial processors. The methods used for object identification and reconstruction are presented. The CMS event selection strategy is discussed. The performance of the HLT is also given.

1. INTRODUCTION

The Large Hadron Collider (LHC) is the next accelerator at CERN. The main motivation [1] for building the LHC is to explore the energy region up to a few TeV range, in view of improving the understanding of the Electroweak Symmetry Breaking mechanism, and of searching for Supersymmetry and any kind of New Physics. Precise tests of the Standard Model are also important topics.

At the LHC protons will collide with centre of mass energy $\sqrt{s}=14$ TeV, about seven times higher than Tevatron II at Fermilab. The total inelastic, non-diffractive p - p cross section at LHC energies exceeds 50 mb. On the other hand the effective cross section in some Higgs boson discovery channels can be as low as at fb-pb level. A comparison shows that the event selection of the order of 1 in 10^{13} is necessary. In order to collect significant statistics a large number of p - p collisions is essential. It is obtained by a strong beam focusing, a large number of protons per bunch and high bunch collision frequency (40 MHz).

It is expected that soon after the startup (first physics runs are planned in year 2007) the luminosity will reach $2 \cdot 10^{33} \text{cm}^{-2} \text{s}^{-1}$ (so called *low luminosity* mode) and will be increased later to the designed luminosity of $10^{34} \text{cm}^{-2} \text{s}^{-1}$ (*high luminosity* mode).

The Compact Muon Solenoid (CMS) is a general purpose experiment for physics discoveries at the highest luminosities of LHC. The present construction status is described in [2]. The CMS design principles [3] include efficient lepton and photon selection and measurements. These particles are of key importance for low statistics discovery physics. The event selection system has to be flexible to preserve signal signatures and efficiently re-

duce the background.

The CMS is capable to handle these requirements keeping the output data volume rate at the level of O(100 MB/s), suitable for permanent storage and further, offline analysis. It has been demonstrated and documented in the CMS Trigger and Data Acquisition reports [4] and [5]. The key elements of the CMS event selection strategy are presented in this paper.

2. THE CMS TRIGGER AND DATA ACQUISITION SYSTEM

Each LHC beam crossing at high luminosity produces about 20 p - p collisions. Each event recorded in the CMS detector results in approximately 1 MB of zero suppressed data. This limits the final data storage rate to about 100 Hz. The event selection at CMS is done in two triggering steps only.

The Level-1 Trigger reduces the 40 MHz beam collision rate below 100 kHz. It is based on custom, partially programmable hardware devices (mostly special-purpose ASICs but also FPGAs where appropriate). It analyses coarsely segmented data from the calorimeter and muon systems only. During the Level-1 Trigger latency the full granularity data are stored in detector front-end pipelines.

The 100 kHz event rate of 1 MB event data is a design constraint for the CMS Data Acquisition (DAQ) and the High-Level Trigger (HLT) system. A schematic view of the central part of CMS DAQ is shown in Fig. 1.

Data from the detectors are stored in modules of the detector Front-End system upon the Level-1 accept is issued by the Global Trigger. Then data are read out and

buffered in Readout Units. At this stage data from each event are spread out over 64 units. The Builder Network is responsible for collecting data belonging to one event and providing it to the Filter Systems. The Builder Network itself is a large switching fabric with 100 GB/s bandwidth, imposed by the Level-1 rate and event size. In the Filter Systems each event is buffered in one Builder Unit and the data from a complete event is provided to a Filter Unit. Each Filter Unit contains a set of commercial CPU units and the Filter Units form a Filter Farm. Each event is assigned to a single processor where the HLT algorithms are executed and the actual event selection takes place. The HLT reduces the Level-1 output rate below the level of the maximal event storage rate of the order of 100 Hz. The final treatment of the events selected at HLT is addressed to Computing Services that forward the data to mass storage and perform some monitoring tasks. The DAQ system is completed with the Event Manager, responsible for the data flow control, and Control System which takes care of configuration, DAQ monitoring and various controlling tasks.

FIGURE 1. Schematic view of the central part of the CMS Data Acquisition System.

The CMS Trigger system does not have any intermediate step between Level 1 and HLT. The traditional "Level 2" trigger is often based on dedicated hardware and uses limited data granularity. The CMS choice to skip this intermediate filter and execute HLT algorithms directly after Level 1 provides a lot of advantages. This design benefits maximally from the computer technology and its developments; it has maximal flexibility without design and architectural limitations of dedicated hardware; there is no limitation on data accessing type and granularity; there is maximal freedom in the choice of the selection algorithms. On the other hand the maintenance of a huge amount of data makes the system challenging

and the HLT algorithms are required to be efficiently performing.

Another important feature of the CMS DAQ is that the system is build out of eight 12.5 kHz units (Fig. 1). This modularity makes the DAQ system staging possible. It is especially important because the expected initial event rate at LHC will be substantially lower than the design rate. Only a part of the system has therefore to be ready at startup, allowing overall cost optimization. The CMS choice, called *50 kHz DAQ scenario* is to build only four DAQ units at startup, providing only half of the design bandwidth.

3. LEVEL-1 TRIGGER

The Level-1 Trigger is based on coarse-grained trigger data from the muon system and calorimeters. The main components and their dependence are shown in Fig. 2.

FIGURE 2. The collaboration between the Level-1 Trigger components.

The Level-1 Calorimeter Trigger uses the information from the electromagnetic calorimeter (ECAL), the hadron calorimeter (HCAL) and the hadron forward calorimeter (HF). The following categories of Level-1 objects are delivered: a maximum of four isolated and non-isolated electrons/photons; at most four taus, central and forward jets; total and missing transverse energy and jet multiplicity. The electron/photon, jet and tau candidates are found locally by the Regional Calorimeter Trigger and combined by the Global Calorimeter Trigger that also make use of energy sums in the calorimeter regions to define global objects. The Global Calorimeter Trigger calculates also additional information which can be used for muon isolation.

The Level-1 Muon Trigger provides up to four Level-1 muons. Trigger data are provided by Resistive Plate Chamber (RPC), Drift Tubes (DT) and Cathode Strip

Chamber (CSC) detectors. Each of these detectors has a dedicated triggering system which builds and defines muon candidates. The final Level-1 muons are selected by the Global Muon Trigger supplied by RPC, DT and CSC subtriggers and calorimeter information needed for optional isolation conditions.

The outputs of the Global Calorimeter Trigger and the Global Muon Trigger are sent to the Level-1 Global Trigger, where the decision whether to accept or reject the event is made. The Global Trigger comprises a novel concept of selection based not only on objects exceeding energy or momentum thresholds but also complex event topology. This is made possible by the availability of candidate position, charge and quality measurement information in the Global Trigger algorithms calculations.

The Global Trigger accept bit pattern is forwarded to the DAQ/HLT system and to the Trigger Timing and Control system in order to signal all detector front-end and readout subsystems. The overall Level-1 Trigger latency is $3.2\mu\text{s}$ and is dominated by the signal propagation from the detector to the control room and back.

The Calorimeter algorithms. The electron/photon identification algorithms are defined in a sliding window of 3×3 towers. Trigger towers in the central part of the detector have a $\Delta\eta \times \Delta\phi$ size approximately equal to 0.087×0.087 following the HCAL segmentation. The ECAL segmentation is much tighter and multiple crystals (5×5 in barrel) are assigned to one tower. The size of the trigger towers varies with the pseudorapidity in the endcaps.

The transverse energy of the electron/photon candidate is defined as the transverse energy deposit in the central tower of the sliding window summed with the largest deposit in one of its four nearest neighbours. In order to reduce the background (Fig. 3) additional identification criteria are introduced. The object is classified as non-isolated electron/photon if, in the central tower, most of the energy is deposited in the narrow strip ($n_\eta \times n_\phi = 2 \times 5$) of crystals (Fine-Grain energy profile) and the HCAL contribution is below a given threshold. The isolated electron/photon has to fulfil the above criteria in all towers of the sliding window supplemented with additional isolation cut. This isolation cut requires that at least one corner of the sliding window to which five towers are associated is quiet, i.e. all the associated towers are below a (programmable) threshold.

The Jet and Tau identification algorithms work in the sliding window made of 3×3 trigger regions where the trigger region is made of 4×4 towers, except for HF where the region corresponds just to a tower. The sliding window is centred on the region with the largest transverse energy, required to be above a fixed value threshold. This central region defines the position of the jet. The jet transverse energy is a sum of transverse

FIGURE 3. Level-1 electron/photon rates. The candidates are selected by energy thresholds, non-isolated electron/photon identification and isolation criteria

energies in a sliding window. Additional requirements are necessary for the identification of jets from hadronic decays of taus (tau-jets). The hadronic decays of tau are dominated by modes with one or three charged particles accompanied by a small number of neutral particles. The products of decay form a low-mass, thin jet with the energy deposited both in ECAL and HCAL. Since the calorimeter response to single particles is often confined in a small cluster extending to no more than 2×2 towers (tau-jet patterns) the sliding window provides a tau-jet tag if all its active regions (ECAL or HCAL energy above a threshold) contain only tau-jet patterns.

On top of the described jet objects the Level-1 Calorimeter Trigger computes, in the pseudorapidity region $|\eta| < 5$, the sum of the transverse energy (E_T) and total missing transverse energy (E_T^{MISS}) by summing E_T components and the information on jet multiplicity (to allow triggering on events with a large number of low energetic jets).

The Muon algorithms. Each of the CMS muon subtriggers (RPC, DT, CSC) has its own trigger algorithm. The RPC trigger is based on a Pattern Comparator. The hit positions from all RPC detectors in a certain η and ϕ range are assigned to a single comparator to form a pattern. It is compared with a set of predefined patterns corresponding to tracks with different momenta. In the case of the DT and the CSCs the multi-hit information in one chamber is used by the local triggers to deliver a track segment. The segments are then combined by a track finder and muon candidates are formed. The three

trigger systems deliver the charge, the position and the estimate of the maximal transverse momentum of each muon candidate, as well as the quality of the measurement.

The role of the Global Muon Trigger is to combine the information from its suppliers in the optimal way. The DT are located in the barrel part of the CMS detector, the CSCs in the forward region while RPCs extend¹ over the full Level-1 Trigger pseudorapidity range $|\eta| < 2.1$. They have different trigger logics so that the three sub-triggers are complementary. This feature is exploited by the Global Muon Trigger to gain both in rate reduction (shown in Fig. 4) and in increase of efficiency.

FIGURE 4. The Level-1 single muon rate (low luminosity) provided by the Global Muon Trigger (GMT). The standalone RPC trigger rate and the combined DT and CSC trigger rates are shown separately. The simulated rate is also shown.

The Level-1 Selection. The trigger thresholds are defined by the total available bandwidth and the allocation to different trigger streams which has to compromise various discovery scenarios. The CMS DAQ system is designed to accept 100 kHz Level-1 output rate, but due to staging at startup the maximum input rate is limited to 50 kHz. The CMS strategy is to apply a safety factor of three on the available bandwidth to take into account simulation uncertainties and non-simulated backgrounds. Thus the trigger tables are designed for the output rate of 16 kHz and 33 kHz for low and high luminosity mode respectively. Since the bandwidth allocation

cannot be optimized without data, it has been divided in equal parts to electron/photon triggers, muons, taus and jet plus combined triggers (as shown in Table 1). This strategy provides inclusive selection optimal for discoveries.

TABLE 1. The Level-1 Trigger table for $2 \cdot 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity and 50 kHz DAQ scenario with safety factor three superimposed. The listed thresholds correspond to 90-95% efficiency.

	Threshold (GeV, GeV/c)	Individual rate (kHz)	Cumulative rate (kHz)
$1e/\gamma, 2e/\gamma$	29, 17	4.3	4.3
$1\mu, 2\mu$	14, 3	3.6	7.9
$1\tau, 2\tau$	86, 59	3.2	10.9
1,3,4-jet(s)	177, 86, 70	3.0	12.5
$\text{Jet} \oplus E_T^{\text{MISS}}$	$88 \oplus 46$	2.3	14.3
$e/\gamma \oplus \text{Jet}$	$21 \oplus 45$	0.8	15.1
min. bias		0.9	16.0

4. HIGH-LEVEL TRIGGER

The HLT algorithms run on a computer Filter Farm, where each event is processed by a single processor. The reconstruction and selection code is written in Object Oriented C++. To cope with the constraints on the event storage resources, the HLT algorithms must achieve a factor ~ 1000 reduction on the event rate. On top of the hardware constraints the HLT selection must satisfy physics requirements to ensure the most inclusive selection and the highest efficiency for "signal" signatures.

The strategy to eliminate unwanted events may explore fast, simplified but not accurate reconstructions or use of minimal amount of precise information. Both ways are used to optimize the event rejection speed although the latter is preferred. The CMS guideline is to keep the HLT reconstruction code as close as possible to that used for the offline reconstruction. The only restriction is that the HLT reconstruction can not rely on the detailed knowledge of the calibration constants which are not available online. The following features are included in the HLT reconstruction code:

- **reconstruction on demand ("when needed"):** The reconstruction is started only if it is required by the trigger decision, rather than always reconstruct particular objects,
- **regional reconstruction ("where needed"):** Data are reconstructed only in the region of interest, defined by a trigger condition. Data from other "partitions" are not accessed. This strategy does not hold for non-seeded objects (for example E_T^{MISS}) where the **global** reconstruction is necessary.

¹ The initial pseudorapidity coverage of RPC trigger has been recently limited to approximately $|\eta| < 1.6$ (staging). This change is not included in the results presented here.

- **partial and conditional track reconstruction ("if needed"):** Tracking is stopped if the combinatorial ambiguities are solved and/or the precision is sufficient for further triggering steps. The reconstruction is abandoned if the track fit parameters disagree with the kinematical requirements (for example momentum below threshold).

The event reconstruction and selection takes place in several steps. Because the reconstruction is triggered by Level-1 Trigger based on calorimetry and muon system, the natural way is to fully explore the provided information with full granularity reconstruction therein. The calorimeter and muon system based reconstruction is called *Level 2*. The reconstructed objects are supplemented with limited information from the Tracker System (usually by Pixel detector hits) at *Level 2.5* and finally reconstructed using more Tracker data at *Level 3*. The *virtual trigger levels* introduced are a matter of convention. The HLT algorithms have no data access restrictions and the order of their execution can be optimized to follow the master principle to reject the event as soon as possible. The description of main algorithms used at HLT is given below.

Electron and Photon reconstruction. The Level-1 electron/photon trigger rate is dominated by fake candidates caused by jets in which a single π^0 carries a substantial fraction of energy. The first step at HLT is to find the ECAL clusters and confirm the Level-1 candidates using the full granularity. Because of the sizeable amount of material ($0.4-1.4 X_0$) in front of the electromagnetic calorimeter, electrons radiate a large fraction of their energy. The recovery of the energy lost in Bremsstrahlung photons is done by measuring the energy in cluster of clusters (*super clusters*) spread in ϕ due to electron bending in magnetic field. The rejection of fake electrons is done at Level 2.5, where the electron candidate is propagated back to the Pixel detector layers and hit matching is performed. The rejection power is shown in Fig. 5. This selection does not affect photons due to much higher energy thresholds required for them. At Level 3 the electron track is reconstructed in the tracker. The cuts on the ratio of the measured energy and the measured momentum as well as track-cluster matching are applied.

Muon reconstruction. In contrast with the electrons/photons the Level-1 muon candidates are real muons and the high rate is caused by momentum mis-measurement. The Level-2 muon algorithm works in a region of interest defined by the Level 1. It uses DT segments, CSC and RPC hits reconstructed locally on chambers. The muon trajectory is built and updated by collecting measurements using the Kalman Filter Fit. Since the most interesting muons originate from the beam spot or

FIGURE 5. The electron selection at Level 2.5 achieved by matching the electron candidate with Pixel detector hits. The upper curve corresponds to the baseline Pixel detector layout with three layers on pixels and two disks. The bottom curve corresponds to only two layers of pixels available.

nearby, the momentum measurement is improved by updating the muon state with the vertex constraint.

The Level-3 muon reconstruction starts with the determination of a region of interest, with parameters and uncertainties defined by the Level-2 muon state at the vertex. The seeds for the regional reconstruction are given by Pixel hit pairs. The Kalman Fitting technique is employed to build a Tracker based muon state which is then updated with muon system measurements (collected by original Level-2 muon). The typical resolution of $(p_T^{gen} - p_T^{rec})/p_T^{rec}$ is 10-16% at Level 2 improved to 1.0-1.7% at Level 3.

The muon rate in the p_T range below 30 GeV/c is dominated by muons from b, c and K, π decays. In order to separate them from those produced in heavy object decays, isolation algorithms are used (Fig. 6). Three techniques were studied: *calorimeter isolation*, with combined ECAL and HCAL energies, *pixel isolation*, with tracks from simplified standalone reconstruction in the Pixel detector and *tracker isolation*, with tracks reconstructed by standard regional reconstruction algorithm. The isolation algorithms apply a threshold (which may depend on η) to the deposits ($\sum E_T$ or $\sum p_T$) in a cone around the muon.

Jet and Missing energy reconstruction. The HLT jet finding is a global algorithm based on iterative grouping of calorimeter towers. Initially, the most energetic tower defines the direction of a "protojet". The towers in a fixed size cone around a protojet direction are as-

FIGURE 6. Contribution to the Level-3 muon rate (high luminosity) before (a) and after (b) isolation algorithms applied. The b, c and K, π components are reduced while W, Z well preserved. The typical isolation cone has a size $\Delta R = 0.2-0.3$ with $\sim 3\text{GeV} \sum p_T$ cut.

signed to the protojet. The size of the jet cone is chosen to be $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} \approx 0.5$. The protojet energy and position are recalculated providing a new protojet definition. The protojet position is determined by weighting the tower positions with their energies. The procedure is repeated until the variation of the protojet energy and position between two iteration steps is small enough. Subsequently the towers contributing to a jet are removed from the list of active towers and the procedure is repeated. The simple calculation of the vector sum of the tower energies is used to calculate the total missing transverse energy E_T^{MISS} , useful for neutrinos identification.

Tau-jet Reconstruction. The CMS tau-jet finding algorithm is optimized for the identification of isolated high energetic taus (e.g. from SUSY Higgs decays). At Level 2 a jet is tagged as a tau jet if the jet energy is con-

centrated in a narrow cone ($\Delta R = 0.13$). The fake tau-jet rejection is improved at Level 2.5 and Level 3 by improving the isolation with reconstructed tracks.

Primary vertex reconstruction. Because of the large number of collisions in a single bunch crossing the knowledge of the position of the signal vertex is useful. Recently updated ([6]) HLT primary vertex finding algorithms are based on fast, three-hit track reconstruction in the Pixel detector. Since the Pixel hit position measurement is very accurate and the track bending is only in transverse plane, the pixel tracks provide a good measurement of the track intercept with the beam axis (Fig. 7) on which the primary vertex finding algorithms rely. The vertex finding efficiency is 95-100%. There are,

FIGURE 7. The resolution of impact parameter along beam line for single track measured in pixel detector only

however several vertices found per event, due to event pileup. The efficiency to select correctly the primary vertex is thus lower (50-97%, event type and luminosity dependant) for general algorithms. It can be however optimized for a particular channel.

HLT selection and efficiency performance. The HLT threshold values for the low luminosity case are presented in Table 2. The thresholds are optimized to achieve the 100 Hz output rate. The initial inclusive event selection principle is followed. The efficiency for some signals, obtained using the above triggering conditions, are given in Table 3.

The timing performance. Due to limited computing resources the execution time of the HLT algorithms is a crucial parameter. The CPU time needed to select HLT physics objects was measured on 1 GHz Pentium III

TABLE 2. High-Level Trigger table for the $2 \cdot 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ luminosity case.

	Threshold (GeV, GeV/c)	Individual rate (Hz)	Cumulative rate (Hz)
1 e , 2 e	29, 17	34	34
1 γ , 2 γ	80, (40 $\oplus$ 25)	9	43
1 μ , 2 μ	19, 7	29	72
1 τ , 2 τ	86, 59	4	76
Jet $\oplus$ E_T^{MISS}	180 $\oplus$ 123	5	81
1,2,3-jet(s)	657, 247, 113	9	89
$e/\gamma \oplus$ Jet	19 $\oplus$ 52	1	90
incl. b -jets	237	5	95
others		10	1005

TABLE 3. HLT selection efficiencies for some types of signal. The numbers are given for events in detector fiducial volume only.

Channel	Efficiency
$H(115\text{GeV}) \rightarrow \gamma\gamma$	77%
$H(160\text{GeV}) \rightarrow WW^* \rightarrow 2\mu$	92%
$H(150\text{GeV}) \rightarrow ZZ \rightarrow 4\mu$	98%
$A/H(200\text{GeV}) \rightarrow 2\tau$	45%
SUSY (~ 0.5 TeV sparticles)	$\sim 60\%$
with R_p - violation	$\sim 20\%$
$W \rightarrow e\nu$	42%
$W \rightarrow \mu\nu$	69%
$t \rightarrow \mu X$	72%

machines (Table 4). The average selection time, defined as the object selection time weighted by its Level-1 input rate, is approximately 300 ms/event.

The CMS Collaboration has assumed that the computer CPU power will follow the Moore's law and double the CPU power each 18 months. This assumption implies that in 2007 eight times² faster machines will be widely available on the market. Currently available algorithms will be executed in $\simeq 40$ ms at LHC startup. Thus, for the 50 kHz DAQ scenario the required number of processors is about 2000 corresponding to a well affordable Filter Farm of about 1000 dual CPU boxes.

This estimate is not accurate. The uncertainties are coming from the following sources: not yet estimated delay time caused by raw data accessing and formatting; assumption that the superimposed DAQ safety factor does not change cpu average; non simulated sources of backgrounds; possible underestimated occupancies in detector that may slow down for example pattern recognition. On the other hand the performance of the algorithms is expected to improve.

TABLE 4. The average CPU time (1 GHz Pentium III) for the HLT algorithms

	CPU (ms)	Rate (kHz)	Time (s)
1 e/γ , 2 e/γ	160	4.3	688
1 μ , 2 μ	710	3.6	2556
1 τ , 2 τ	130	3.0	390
Jets, Jet $\oplus$ E_T^{MISS}	50	3.4	170
$e/\gamma \oplus$ Jet	165	0.8	132
b -jets	300	0.5	150

5. SUMMARY

The CMS experiment has two trigger levels to reduce the 40 MHz LHC bunch crossing rate to 100 kHz (50 kHz for initial running conditions) in the first step and to the order of 100 Hz appropriate for mass storage, in the second step. The algorithms for the event selection for both trigger stages were developed. The current average execution time of the HLT algorithms is 300 ms/event (1 GHz Pentium III). This execution time is adequate provided that the progress in computer technology arise as expected. The CMS Trigger table is initially optimized for the most inclusive selection allowing powerful event rejection while preserving signatures of possible discovery channels. The current trigger table, however, does not include some Standard Model physics selections relevant to CMS. An example is the B -physics selection for which highly efficient and low CPU cost HLT algorithms are under development [7].

REFERENCES

1. Jarlskog, C., and Rein, G., editors, *Proceedings of the Large Hadron Collider Workshop*, Aachen, 4-9 October 1990, CERN 90-10, ECFA 90-133.
2. Neal, H., "CMS Detector Update," these proceedings.
3. CMS Collaboration, *CMS, The Compact Muon Solenoid, Technical Proposal*, Tech. rep., CERN/LHCC 94-38 (1994).
4. CMS Collaboration, *CMS, The Trigger and Data Acquisition project, Volume I: The Level-1 Trigger*, Tech. rep., CERN/LHCC 2000-038 (2000).
5. CMS Collaboration, *CMS, The Trigger and Data Acquisition project, Volume II: Data Acquisition & High-Level Trigger*, Tech. rep., CERN/LHCC 2002-26 (2002).
6. Cucciarelli, S. *et al*, *Track-Parameter Evaluation and Primary-Vertex Finding with the Pixel Detector*, CMS NOTE-2003/026.
7. Marinelli, N., "*b* Physics at LHC with the CMS detector," these proceedings.

² The CPU time measurements with 1 GHz Pentium III machines were done in the year 2002. These machines are already obsolete.