Eur Phys J C **34**, s01, s75–s83 (2004)
Digital Object Identifier (DOI) 10.1140/epjcd/s2004-04-009-7

EPJ C direct
electronic only

The CMS muon system: Physics performance

S. Lacaprara, on behalf of the CMS collaboration

Università di Padova and INFN, via Marzolo, 8, 35131 Padova, Italy, e-mail: Stefano.Lacaprara@pd.infn.it

Received: 4 August 2003 / Accepted: 18 February 2004 /
Published Online: 13 July 2004 –

Abstract. The CMS experiment will take data at the CERN LHC starting from 2007. The CMS muon system has been designed to identify, reconstruct and measure muons with high efficiency and accuracy. In this paper the layout of the system and the key features of the detectors are presented. The reconstruction algorithms, in the context of the High Level Trigger, and their performance in terms of resolution, rate reduction and trigger efficiency are also discussed.

PACS. 2 5.70.Ef – 21.60.Gx – 27.30.+t

1 Introduction

The CMS (Compact Muon Solenoid) experiment is one of the two general purpose experiments that will study proton-proton and heavy ion collisions at the Large Hadron Collider (LHC). The accelerator is currently under construction at CERN, and is scheduled to start delivering collisions from mid 2007.

Muons with high transverse momenta are one of the probes for discovery of new phenomena as well as for precision measurements; the CMS apparatus has a robust, flexible and redundant system to identify and reconstruct such particles [1].

The main feature of the CMS detector is a big superconducting solenoid, which provides a strong 4 T magnetic field in the central part, ensuring high momentum resolution for charged particles. A tracking system and both electromagnetic and hadron calorimeters are placed inside the magnet (bore radius $\sim$ 3 m), which leads to a very compact design. The tracking system is based on silicon strip and silicon pixel detectors for the vertex reconstruction. The electromagnetic calorimeter is homogeneous, made of $PbWO_4$ crystals. The hadron calorimeter is made of copper absorber plates interleaved with 4 mm thick plastic scintilator tiles. Finally, the muon system consists of tracking and triggering chambers and is embedded in the iron of the magnet return yoke. The CMS detector has a length of 21.6 m, a diameter 14.6 m and a total weight of $\sim$ 12500 tons.

This article describes the design and performance of the CMS muon system, i.e. the detectors, the reconstruction algorithms and the physics performance. The muon reconstruction algorithms have been developed mainly in the context of the High Level Trigger (HLT) [2]. The CMS HLT system is based on trigger algorithms implemented in software running on a farm of commercial processors. Thus, even if the off-line reconstruction uses, as much as possible, the same algorithms (with full use of calibration and alignment information), the main focus has been devoted to rate reduction and trigger selection.

2 Muon system layout

The choice of the detector technology has been driven by the very large surface to be covered and by the different background radiation environments. A space resolution of the order of 100 μm is required for the muon chambers to improve the p_t measurement done by the inner tracker for very high momenta ($p_t \gtrsim 200$ GeV/c). For lower momenta the momentum resolution is given by the tracker reconstruction alone. Other crucial features of the detectors chosen are the triggering and the bunch crossing identification capabilities.

Three types of gaseous detectors are used to identify and measure muons. In the barrel region ($|\eta| < 1.2$), where the neutron induced background is negligible, the muon rate is low ($R(\mu) \lesssim 1$ Hz/cm^2) and the residual magnetic field is low, Drift Tube chambers (DT) are used. In the endcap region, where the muon rate as well as the neutron background is high ($R(\mu) \lesssim 10$ kHz/cm^2, $R(n,\gamma) \sim 10$ kHz/cm^2), and the magnetic field is high as well, Cathode Strip Chambers (CSC) are placed and cover up to $|\eta| < 2.4$. Finally, a third type of detector, Resistive Plate Chambers (RPC), dedicated to the trigger, is used in both barrel and endcaps, up to $|\eta| < 2.1$. The layout of one quarter of the CMS muon system is shown in Fig. 1.

In the barrel region, four stations of detectors are arranged in cylinders interleaved with the iron yoke. Each muon station has a DT and a RPC, plus one additional layer of RPC in each of the first two stations, to improve the low p_t trigger. The segmentation along the beam direction follows the five wheels of the yoke.

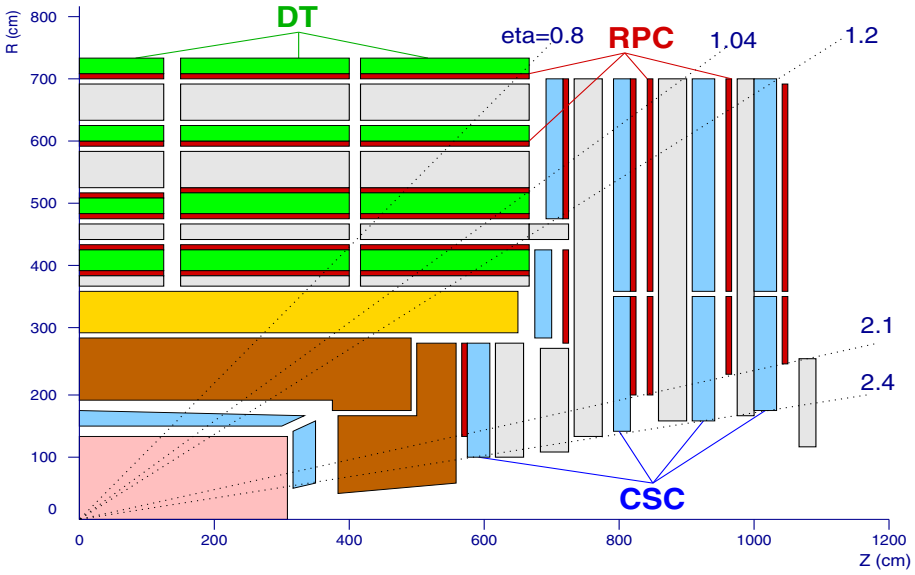

Fig. 1. Layout of the CMS muon system

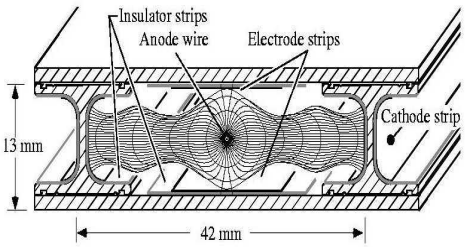

Fig. 2. Schematic view of a drift cell with electric field lines

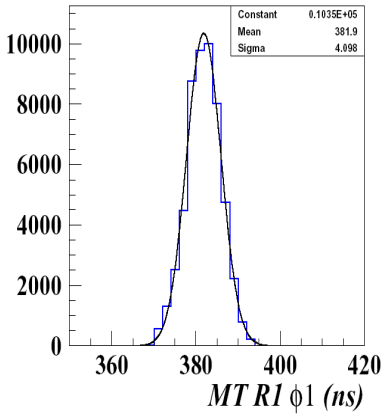

Fig. 3. Time resolution of BTI

In each of the endcaps, the CSCs and RPCs are arranged in four disks perpendicular to the beam, and in concentric rings around the disks, three in the innermost, and two in the others.

In the initial phase of LHC, a staged scenario is foreseen. In the endcaps, the outer ring of the last disk will be missing and the CSC electronics for the Level 1 trigger will be not present in the innermost chambers. The Level 1 trigger acceptance will therefore be limited to $|\eta| < 2.1$. Moreover the RPC will be installed only up to $|\eta| < 1.6$.

2.1 Drift Tube chamber

The basic component of a Drift Tube chamber is the drift cell, shown in Fig. 2. The cell dimensions are 42×13 mm^2, and the length varies from chamber to chamber (from ~ 2 to ~ 3 m). The design of the cell has been optimized in order to have excellent linearity in the time-to-distance relationship, needed for the trigger. The gas mixture is Ar/CO_2 (85/15%); for nominal high voltage setting, the maximum drift time is $\sim$ 380 ns and the resolution is ~ 180 μm.

The drift cells are assembled in layers. Four layers are glued together to form a *superlayer* (SL), with consecutive layers staggered by half a cell. Each chamber is made of three SLs, two measuring the bending coordinate (Rϕ), and one for the coordinate along the beam (Rz). The first Rϕ SL is followed by an honeycomb spacer, then by the Rz SL and finally by the second Rϕ SL. The spacer, ~ 13 cm thick, increases the lever arm between the two Rϕ SLs and ensures the needed mechanical rigidity with low Z material. In the outermost station, the Rz SL is missing.

The important features of these chambers are the self-triggering and the bunch crossing identification capabilities. Both are achieved by a staggering of the layers in a superlayer. The Bunch and Trigger Identification (BTI) electronics combines the drift time measured in three consecutive layers, applying a mean timer technique. The excellent linearity of the time-to-distance relationship of the cell allows a time resolution of about ~ 4 ns (Fig. 3) [3]. At the same time, track segments of aligned hits are built. The multi-layer structure of the chambers improves the spatial resolution of the single cell, which reaches the needed $\sigma_{R\phi} \sim 100$ μm. It also reduces the impact of uncorrelated

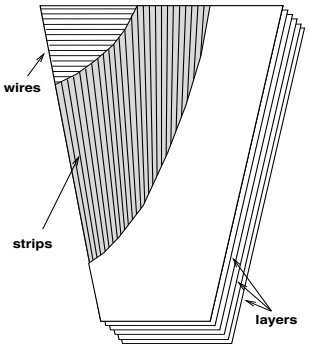

Fig. 4. Schematic layout of a CSC

hits in the layers due to electronic noise or neutron and γ induced activity.

2.2 Cathode strip chambers

In the endcap region, the background coming from thermal neutrons and halo muons is expected to be sizeable, and the magnetic field is strong and highly inhomogeneous. A different detector technology, Cathode Strip Chambers, has therefore to be used. The CSCs are multiwire chambers able to provide precise time and position measurements in this difficult environment. The basic detector element is a layer, defined by two cathode planes, one segmented into radial strips, defining a gas gap with an array of parallel wires (Fig. 4).

The avalanche produced by a crossing charged particle is collected by the wires, and the image charge is collected in several adjacent strips. An interpolation of the strip signal gives a very precise spatial measurement. In addition, the wire signal measurement allows the other coordinate to be determined, although with a coarse precision. The resolution is $\sigma_{R\phi} \sim 100 \div 240$ μm for strip measurements, depending on the chamber, and $\sigma_R \sim 0.5$ cm for wire measurements. Finally, the small wire spacing of about 3.4 mm gives a fast chamber response. The gas mixture is Ar-CO_2-CF_4 (30%-50%-20%). No hydrogen is present to reduce the neutron-induced noise.

The CSC detectors use the multi-layer structure for self triggering. A Local Charged Track (LCT) is built from the spatial and time correlation of the hits in the two projections. The anode (wires) signals are used to identify the bunch crossing with a multi-layer time coincidence in the LCT pattern. Typically two layers are used to set the bunch crossing, while four are required to establish the existence of a muon track.

2.3 Resistive Plate Chambers

The third detector technology used in the CMS muon system is the Resistive Plate Chambers. A RPC (Fig. 5) is formed by two planes of material with high resistivity (Bakelite) separated by a gas gap of few mm. The planes are externally coated by graphite, which forms the

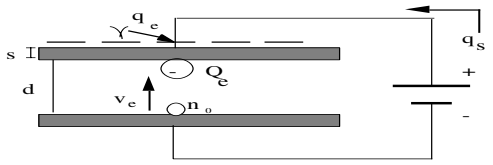

Fig. 5. Operating schema of the RPC

cathode for the high voltage (9.5 kV). The avalanche generated by a crossing particle induces a signal in aluminum strips placed outside the graphite cathodes, which is then read out. The RPC operates in "avalanche mode" in order to improve the rate capability. This mode requires a reduced electric field (and consequently a reduced gas multiplication) with respect to the streamer mode and therefore a robust electronic amplification is needed. In CMS double-gap RPC are used, with two gas gaps read out by a single set of strips in the middle. This design increases the signal on the read-out strip, which sees the sum of the single gap signals. The gas mixture consists of 95% $C_2H_2F_4$, 5% i-C_4H_{10}.

The time resolution of the RPC is $\sigma_t \sim 1$ ns. The spatial resolution, in the bending coordinate, is determined by the pitch of the strips and is of the order of $\sigma_{R\phi} \sim 1$ cm. No measurement is provided in the other coordinate.

The excellent time resolution of the RPC provides an unambiguous bunch crossing assignment. The RPC are used also to identify and measure muon tracks at the Level 1 trigger with a pattern comparator logic. Muons are reconstructed from correlated patterns of hits in the different stations. The parameters (p_t, η, ϕ) are assigned via look up tables.

3 Muon reconstruction

The muon reconstruction algorithms presented in this paper have been developed in the context of the High Level Trigger. The reconstruction proceeds in consecutive steps, each of which uses the previous result as starting point to improve the definition of the muon. The basic feature of the HLT reconstruction is the "regional" behavior. The reconstruction is performed only in a limited region of the detector, where the previous step has already found a muon candidate. This technique allows the time needed for the algorithm to be substantially reduced, a crucial issue for the HLT.

The algorithm developed for the HLT starts from a Level 1 trigger accept decision, and uses the Level 1 candidates as initial seeds to define the region of the detector where the so-called "Level 2" reconstruction is performed. Because the HLT is fully software based, there is no strict separation between the different levels of the High Level Trigger. Two different levels (Level 2 and Level 3) however are defined. The Level 2 uses only information from the muon system and the calorimeters, while the Level 3 uses also inner-tracker hits. In contrast to the Level 1, the Level 2 has access to event data with full granularity and therefore can exploit the full detector resolution.

3.1 Local reconstruction

The first step is local reconstruction in the multi-layer chambers (DT and CSC), which associates aligned hits to build track segments, subsequently used by the regional reconstruction [4].

3.1.1 Local reconstruction for DT

In the DT chambers, hits in the three superlayers are reconstructed in the two different projections, "bending" (Rϕ) and "along the beam axis" (Rz). First, the hits in the cells are reconstructed, with a linear time-space conversion. A single cell has an intrinsic left-right ambiguity which can be solved only by associating hits from different layers. In order to get the best cell resolution, the drift velocity in the gas used for the reconstruction of the hits is not a constant. It depends on magnetic field value (along and perpendicular to the wire) and on the impact angle in the cell. These three effects are particularly important for chambers in the outer wheels, near the endcaps, where the residual magnetic field in the chambers can be sizeable. Moreover, for the Rz superlayer, the impact angle is systematically big. To apply a precise correction, the coordinate along the wire, as well as the reconstructed impact angle are needed. As they are not available at the first step of the hit reconstruction, an iterative process is applied. Initially the magnetic field value at the centre of the wire is used. The impact angle is set to zero for the Rϕ SL, while it is assumed that the muon comes from the nominal interaction point for the Rz SL. The uncertainties associated to the hits depend also on magnetic field and impact angle, and are iteratively improved as well.

For each projection, two-dimensional segments are built, with a linear fit. The left/right ambiguities are solved with a best χ^2 criterion. In the Rϕ projection, hits in the two SLs are used. A correction for the impact angle is applied on the hits forming the segments, and the fit is repeated. Finally, the two projections are associated and a three-dimensional segment is built. At this point, the coordinate along the wire is known for both projections. Further corrections to the hit reconstruction from a more refined knowledge of the magnetic field are applied, and are followed by a last fit. The resolution of the local reconstruction amounts to $\sigma_x \sim 100$ μm for position and $\sigma_\theta \sim 1$ mrad for direction.

3.1.2 Local reconstruction for CSC

Each layer of the cathodes strip chambers is able to provide both coordinates. The measurement in the bending coordinate is obtained by clustering together adjacent strips with signals, and by fitting the charge distribution with a proper function (Gatti function [5]) to get the cluster centroid and width. The resolution of this measurement is ~ 100 μm for the innermost CSCs, where the strip pitch is smaller, while it is ~ 200 μm for the other chambers. The dominant causes for signal degradation are the electronic noise and cross talk between nearby strips.

The other coordinate is obtained from the wire signal, discriminated and read out after grouping several wires together. The resolution is therefore $w/\sqrt{12}$, where w is the width of wire group, and is of the order of $\sim 0.5 \div 1$ cm. The two projections are associated using a time coincidence of the hits.

Finally, the three-dimensional hits are used to create three-dimensional track segments with a linear fit. The pattern recognition starts with two hits in the first and last layers. Hits in other layers are searched for and chosen according to a χ^2 compatibility. No hit sharing is allowed between different segments.

3.2 Level 2 reconstruction

The Level 2 reconstruction is based only on the muon detectors, without usage of the inner-tracker. Both tracking detectors (DT and CSC) and RPCs participate in the reconstruction. In spite of the poor spatial resolution, the latter complement the tracking chambers, especially where the geometrical coverage is problematic, mostly in the barrel-endcap overlap region.

The reconstruction starts from a seed, given by a muon candidate from the Level 1 muon trigger. A different seed generation has been developed, which uses patterns of segments reconstructed in CSC and/or DT chambers, and assigns an initial transverse momentum with a parametrization for the barrel part, and with an helix hypothesis for the muon trajectory in the endcaps. This seed generation requires the segments to be reconstructed, and therefore lacks the regional behavior. In principle, however, any other seed can be used, including one from the inner-tracker reconstruction, in case the need is to identify as a muon a high-p_t track seen in the tracker.

A seed is a state vector, with position, direction and p_t estimate, defined at the innermost station of the muon system. Starting from a seed, the chambers compatible with the state are identified, and the local reconstruction, as described above, is performed only in these chambers. This is the key point of the regional behavior of the reconstruction. The output of the local reconstruction is used to grow the muon trajectory, initially inside-out, with a Kalman filter technique [6]. It uses directly the three-dimensional segments for the DT chambers, with either position and direction information. For the CSC, where the magnetic field is stronger, three-dimensional hits constituents of the segments are used. Finally, hits from the RPC are used. A suitable χ^2 cut is applied in order to reject bad hits, mostly due to showering, delta rays and pair production. In case no matching hits (or segments) are found, e.g. due to detector inefficiencies, geometrical cracks or hard showering, the search is continued in the next station.

The improved state obtained by including the hits and segments is used to refine the state vector, and to further reduce the list of detectors where the local reconstruction is performed. The state is propagated from one station to the other using the `GEANE` [7] package, which takes into

account both the magnetic field and the material effects, such as multiple scattering and energy loss. The uncertainties on the parameters (position, direction and momentum) are updated together with the state. This process is very CPU-time consuming, as shown in Sect. 4.3, and dominates the total algorithm CPU-time.

This procedure is iterated until the outermost chamber is reached, where the state is now well defined. A backward Kalman filter is then applied, from outside in, and the track parameters are defined at the innermost muon station. Finally, the track is extrapolated to the nominal interaction point and a vertex constraint is applied with the nominal vertex smearing ($\sigma_{xy} = 15$ μm, $\sigma_z = 5.3$ cm). A cut on the χ^2 of the track is applied to reject badly reconstructed tracks, mostly muons coming from π/K decays in flight, and eventual ambiguities are solved according to the number of hits and χ^2 criterion.

Figure 6 shows the transverse momentum resolution for muons coming from $W \to \mu\nu$ decays, defined as

$$\frac{\sigma_{1/p_t}}{1/p_t} = \frac{1/p_t^{\rm rec} - 1/p_t^{\rm gen}}{1/p_t^{\rm gen}} \ . \qquad (1)$$

The resolution amounts to 10% in the barrel region ($|\eta| < 0.8$), 15% in the overlap region ($0.8 < |\eta| < 1.2$) and 16% in the endcaps ($1.2 < |\eta| < 2.1$). The Level 2 muon momentum precision is essentially determined by the measurement in the transverse plane of the muon bending angle at the exit of the 4 T coil, taking the beam spot as the origin of the muon. This measurement is dominated by multiple scattering in the material before the first muon station up to p_t values of 200 GeV/c, when the chamber space resolution starts to dominate. The efficiency of the reconstruction algorithm is everywhere larger than 99% with respect to the Level 1 input.

3.3 Level 3 reconstruction

The following step in the muon reconstruction, called "Level 3" reconstruction, is the inclusion of the inner-tracker hits. As in the previous section, a seed is defined as the output of the previous step, the Level 2 reconstruction in this case. Muon tracks are extrapolated to the outermost tracker surface and to the nominal interaction point. Silicon layers compatible with the muon trajectory are identified. A region of interest inside these layers is then determined, according to the muon track parameters and uncertainties, including those coming from the extrapolation. The region is defined with the assumption that the muon origins from the nominal interaction point. This assumption has the side effect of strongly reducing the efficiency for muons which do not originate from the main vertex, namely those coming from π/K decays in flight.

Inside the region of interest, a set of local track seeds are built, as pairs of hits from different layers, with a (loose) beam spot constraint. Each seed is transformed into a trajectory with a pattern recognition based on a forward Kalman filter, similar to that used in the muon system, even if applied to detectors with much smaller dimensions. If more than one hit is found in a given layer, the trajectory is duplicated and grown in parallel. If no hit is found in a layer, the next layer is used in turn. Appropriate selections are applied in order to limit the combinatorics. A trajectory is stopped if no compatible hits are found in two consecutive layers or if the transverse momentum is too small (0.9 GeV/c).

The surviving tracks are checked for ambiguities, solved on the basis of the number of hits and the χ^2 of the fit. Finally, all reconstructed tracks are fit once again with the tracker hits and the muon hits found in the Level 2 reconstruction, and selected according to a χ^2 cut.

The high resolution of the silicon detectors and the strong magnetic field present inside the solenoid improve transverse momentum resolution by about a factor ten with respect to the muon-system stand-alone reconstruction. The momentum resolution for muons coming from $W \to \mu\nu$ decays is shown in Fig. 7 and amounts to 1.0% in the barrel region ($|\eta| < 0.8$), 1.4% in the overlap region ($0.8 < |\eta| < 1.2$) and 1.7% in the endcaps ($1.2 < |\eta| < 2.1$).

The algorithmic efficiency of the Level 3 reconstruction, relative to Level 2, is shown in Fig. 8 as a function of $|\eta|$, for a sample of single muons generated in the p_t range $10 < p_t < 100$ GeV/c, without the inclusion of pile-up. The efficiency is typically 99%, except for the difficult region around $|\eta| \sim 1$, where DT and CSC systems overlap, where the efficiency drops to about 97%.

3.4 Isolation

The isolation algorithms are important for the muon HLT selection [8]. Indeed, they help in suppressing the muons coming from b, c, π and K decays, generally accompanied by other nearby particles. The general strategy is to count the total energy deposited in the calorimeters or the sum of the momenta of charged particle tracks reconstructed in the tracker in a cone around the reconstructed muon, and to apply a threshold. The cone is actually a circle in the (η, ϕ) plane centered around the muon direction (extrapolated to the proper surface), defined by $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} \leqslant \Delta R_{\rm max}$.

Three different algorithms have been developed. The first uses the energy deposited in the electromagnetic and hadron calorimeters, the second uses tracks reconstructed in the pixel detector and the third uses the full tracker reconstruction. In all cases, the cones size and thresholds are optimized in order to get maximum rejection for minimum-bias events above a given trigger p_t threshold for a fixed efficiency in a reference signal sample ($W \to \mu\nu$ decays). Thresholds are $|\eta|$ dependent so that the efficiency for the reference sample is flat with respect to $|\eta|$.

The calorimeter algorithm uses a combination of electromagnetic and hadronic energy deposit in a weighted sum $E_t = \alpha E_t^{\rm ecal} + E_t^{\rm hcal}$, with $\alpha = 1.5$, a value found to give optimal results. The energy deposited by the muon itself is subtracted. Thresholds on single calorimeter towers

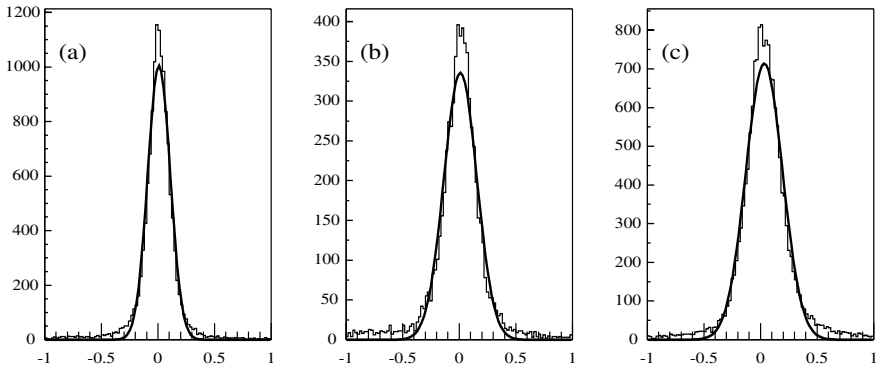

Fig. 6. $\sigma_{1/p_t}/(1/p_t)$ for muons from W $\to \mu\nu$ decays for Level 2 reconstruction in barrel (**a**), overlap (**b**) and endcaps (**c**) region

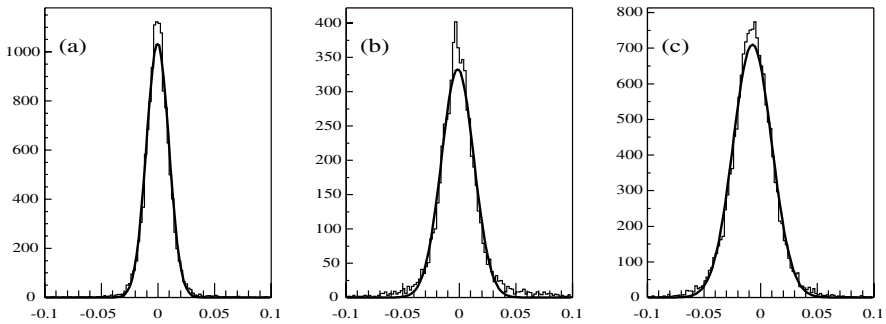

Fig. 7. $\sigma_{1/p_t}/(1/p_t)$ for muons from W $\to \mu\nu$ decays for Level 3 reconstruction in the barrel (**a**), the overlap (**b**) and the endcaps (**c**) region. The x-axis ranges are different from Fig. 6

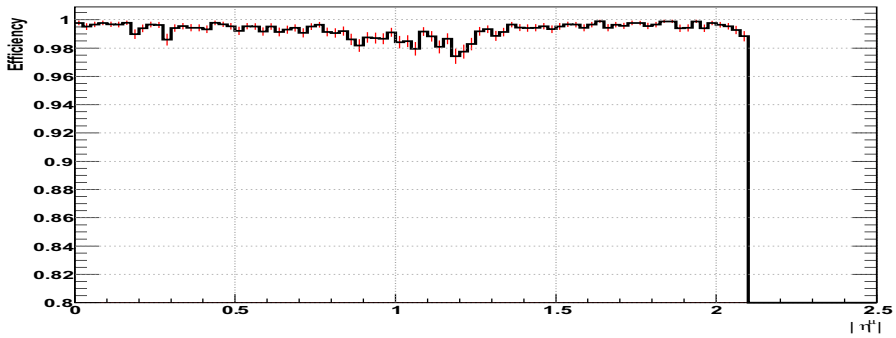

Fig. 8. Level 3 algorithmic efficiency with respect to Level 2 as a function of $|\eta|$ for single muon without pile-up

are applied to reject noise and pile-up deposits. The typical thresholds on $\sum E_t$ are from 6.5 to 9 GeV, for typical cone sizes of $\Delta R_{\max} = 0.2$. This algorithm is sensitive to the pile-up conditions, since it is not able to distinguish between signal and pile-up energy deposit. It uses only information from the calorimeters and the muon system, and therefore it can be applied at Level 2.

The pixel algorithm uses track reconstruction with the three pixel layers of the vertex detector. A primary vertex is reconstructed, defined as that linked to the pixel track closest to the reconstructed muon. Only the tracks which originate from the the same vertex as the muon are used to compute the sum of transverse momenta, to reduce the effect of pile-up. This algorithm can be applied either after Level 2 reconstruction or after Level 3 reconstruction. In the former case, the uncertainties of the parameters of the muon track at the interaction point are sizeable, given the resolution of the muon system stand-alone and the multiple scattering which occurs between the innermost muon station and the vertex. In this case, the vertex requirement cannot be used, and the algorithm is less effective.

The last algorithm uses the full tracker reconstruction, performed only in a region around the muon trajectory, and sums the transverse momenta of all the tracks inside a cone, except the muon. As in the pixel isolation algorithm, a vertex constraint is applied.

The performance of the three algorithms depend strongly on the threshold applied to the p_t of the muon, as shown in Fig. 9. This fact has important implications on the effectiveness of the selection when applied at Level

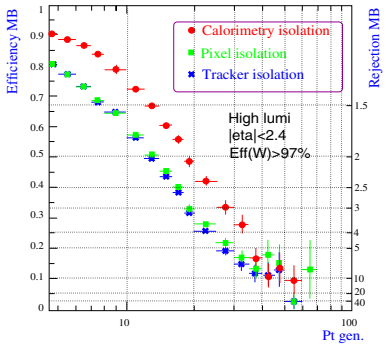

Fig. 9. Efficiency of the three isolation algorithms as a function of the generated muon p_t for minimum-bias events, for a fixed efficiency ($\epsilon > 97\%$) for a $W \to \mu\nu$ sample, at high luminosity

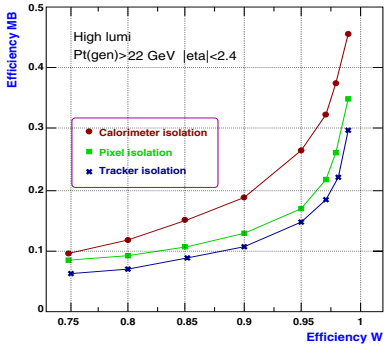

Fig. 10. Efficiency of the three isolation algorithms for minimum-bias events as a function of the analogous efficiency for the reference sample ($W \to \mu\nu$), at high luminosity. A threshold of 22 GeV/c on the generated p_t is applied

2, as shown in Sect. 4.1. At Level 2 a large fraction of the events passing even a rather high p_t threshold have only low p_t muons. Due to problems in the reconstruction (e.g. for hard showering, multiple scattering, etc.), these low p_t muons are reconstructed with a p_t above the threshold. Because for these events the isolation algorithms are applied on low p_t muons, they are not effective.

Figure 10 shows the rejection of minimum-bias event as a function of the efficiency for the reference sample ($W \to \mu\nu$), when a threshold of 22 GeV/c is applied on the generated muon p_t. In this configuration it is possible to achieve a factor five to seven rejection on minimum-bias events, while $W \to \mu\nu$ events are accepted with $\sim 97\%$ efficiency. Such a sharp p_t threshold cannot be applied on real events, but the excellent resolution available at Level 3 is very close to this ideal scenario.

4 Performance

The muon reconstruction has up to now been developed mainly for the HLT selection. The performance of the muon system are therefore shown here in that context. The most important two features of the muon reconstruction in the context of HLT are rate reduction and signal efficiency, shown in the next two sections.

4.1 Rate reduction

Figure 11 shows the integral rate after each trigger level for single muon events, both for low ($\mathcal{L} = 2 \times 10^{33}$ cm^{-2}s^{-1}) and high ($\mathcal{L} = 10^{34}$ cm^{-2}s^{-1}) LHC luminosity, as a function of the p_t trigger threshold. All relevant physics processes have been included [9], i.e. minimum-bias events starting from a very low p_t^{μ}, including heavy flavour generation, W, Z and $t\bar{t}$ production.

It is possible to see the effect of the so-called "feed-through", namely the effect of promoting badly reconstructed low p_t muons to high p_t.

The feed-through is the reason why, at Level 1, the measured rate is substantially higher than the generated rate. It is the convolution of the tails of the p_t resolution distribution with the very steep p_t spectrum of muons, especially at low momentum. A similar effect, although reduced, is present also at Level 2, in spite of the improved resolution. Nevertheless, a rate reduction of about a factor of ten is achieved, for a threshold around 20 GeV/c. As mentioned above, the isolation algorithm is not very effective at Level 2, due to the feed-through contamination. At Level 3, due to the improved resolution, but mainly to the absence of significant tails, the reconstructed spectrum follows nicely the generated spectrum. Finally, the rate reduction achieved applying the isolation algorithms at Level 3 becomes visible. Above $\sim$ 25 GeV/c, the total rate is dominated by muons coming from $W \to \mu\nu$ decays, not rejected by the isolation criteria.

4.2 Signal efficiency

The selection efficiency for signal events depends on the p_t threshold. The threshold in turn is set to reduce the output rate after each level so as to match the allowed bandwidth. It is beyond the scope of this paper to describe the full Level 1 and HLT trigger Tables [10]. Only some informations relevant to the muon system performance illustration are given.

The proposed thresholds at HLT for the low and high luminosity scenarios together with the output rate and the physics content are summarized in Table 1.

The flexibility of the software allows for more complex and exclusive trigger selections. As an example, in order to enhance the b content of the final rate it is possible to have a dedicated trigger for the $B_s \to J/\psi\ \phi$ signal. This trigger is based on the invariant mass of two muons passing at Level 1 with a fast tracker reconstruction [11]. Since the Level 1 thresholds are lower than the corresponding HLT thresholds (at low luminosity they amount to 14 and 3 GeV/c for single and double muon events, respectively, at high luminosity to 20 and 5 GeV/c) it is possible to select many events below the HLT p_t cut. The estimated efficiency can be increased up to $\sim$ 5%.

The efficiency for two reference signals $W \to \mu\nu$ and $t\bar{t} \to \mu + X$ is shown in Fig. 12 as a function of the single muon p_t threshold after the various trigger levels.

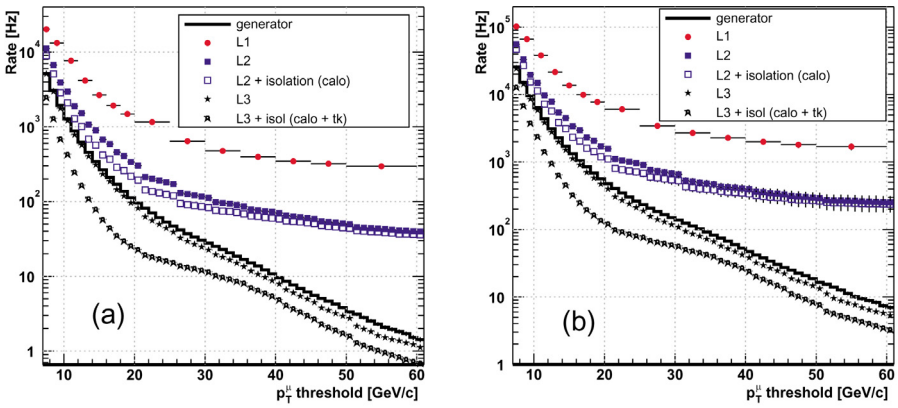

Fig. 11. Integral rate as a function of p_t^μ threshold for events with one muon in the final state, for successive trigger levels, for low luminosity (**a**) and high luminosity (**b**) scenarios

Table 1. Proposed thresholds and total and detailed output rate for single and double muon triggers at HLT for low and high luminosity

$\mathcal{L}$ ($\mathrm{cm^{-2}s^{-1}}$)	single μ thr. (GeV/c)	di μ thr. (GeV/c)	tot. rate (Hz)	π K (Hz)	b c τ (Hz)	W (Hz)	Z (Hz)
2×10^{33}	19	7	29	3.4	8.7	14.5	2.4
10^{34}	31	10	55	0.8	2	42	7.6

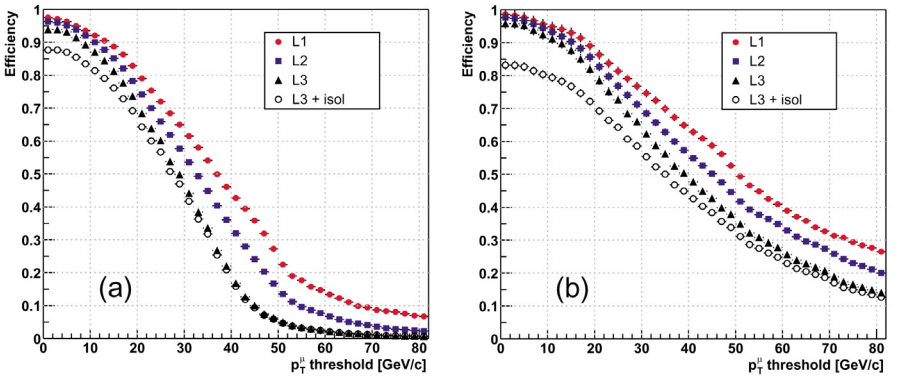

Fig. 12. Efficiency at different levels of the trigger selection for $\mathrm{W} \to \mu\nu$ events (**a**) and $\mathrm{t\bar{t}} \to \mu + \mathrm{X}$ events (**b**) as a function of the single muon p_t threshold for low luminosity

Table 2 shows the inclusive trigger efficiency using single and double muon selections for different signal samples, including Higgs bosons at different masses.

4.3 CPU time

Because the HLT algorithms are executed on an on-line farm, the CPU time needed to reconstruct muons is a crucial issue. Table 3 shows the time per event needed on average to process the muon reconstruction and isolation algorithms Each algorithm is applied on the events passing the previous level selection. At Level 1, the threshold is $p_t > 10$ GeV/c, (18 GeV/c) for low (high) luminosity. The time given in ms was computed on a benchmark `INTEL PIII`, 1 GHz CPU. Results are presented for both low and high luminosity. The time excluding the fraction spent in propagation (using `GEANE` packages) is also shown.

These results show that most of the time spent in the Level 2 reconstruction is spent just in propagating states, mainly across the iron of the magnet yoke. Work is underway to replace `GEANE` with an optimized propagation package customized for the CMS geometry. A significant speed-up is therefore expected.

5 Conclusion

The layout of the CMS muon system, and the main features of the three muon detectors have been summarized. The reconstruction of muons in the framework of the High Level Trigger, both stand-alone and with the

Table 3. Average CPU time needed to reconstruct an event by the HTL algorithms. The *Total/L1 event* time represents the average time spent per Level 1 event, taking into account the rate reduction obtained by applying the various trigger levels on a realistic data spectrum

	mean CPU time (ms/event)			
	Low Lumi ($p_t > 10$ GeV/c)		High Lumi ($p_t > 18$ GeV/c)	
algorithm	total	excluding `GEANE`	total	excluding `GEANE`
Level 2	640	100	580	100
Calo iso	100	25	90	40
Level 3	420	200	590	420
Pixel iso	65	65	320	320
Tk iso	190	190	370	370
Total/L1 event	710	125	660	150

Table 2. Inclusive trigger efficiency for different signal samples using single and double muon selections. Higgs boson masses are in GeV/c^2, detector acceptance is not included. Results for low (LL) and high luminosity (HL) scenario are presented

Signal	ϵ_{LL}	ϵ_{HL}
$W \to \mu\nu$	69%	42%
$Z \to \mu\mu$	92%	86%
$t\bar{t} \to \mu + X$	72%	58%
$H_{m=120} \to WW \to \mu\nu\mu\nu$	87%	64%
$H_{m=160} \to WW \to \mu\nu\mu\nu$	92%	77%
$H_{m=150} \to ZZ^* \to \mu\mu\mu\mu$	98%	97%
$H_{m=200} \to ZZ \to \mu\mu\mu\mu$	99%	99%

inner-tracker, have been discussed. The performance of the muon reconstruction, isolation and selection algorithms has been shown. The background rate reduction and the efficiency at selecting some reference signal samples have been demonstrated, together with the actual status of the CPU time needed to perform the reconstruction algorithms.

The CMS muon system has been designed to be robust, flexible and redundant, and the actual implementation of the reconstruction algorithms already exploits much of these features. More sophisticated algorithms can, and most probably will, be developed in the future. They are already now adequate for the rich physics program expected at CMS.

Acknowledgements. I would like to thank Darin Acosta, Norbert Neumeister and Patrick Janot for their suggestions and corrections.

References

1. The CMS Collaboration: The Muon Project, Technical Design Report, CERN/LHCC 97-32, CMS TDR 3 (1997)
2. The CMS Collaboration: The Trigger and Data Acquisition Project, Volume II, Data Acquisition & High-Level Trigger Technical Design Report, CERN/LHCC 2002-26, CMS TDR 6.2 (2002)
3. M. Cerrada et al.: CMS-NOTE 2003/007 (2003)
4. G. Bruno et al.: CMS-NOTE 2003/043 (2003)
5. E. Gatti et al.: Nucl. Instrum. and Methods, **188**, 327-346 (1981)
6. R. Frühwirth: Nucl. Instrum. and Methods, A **262**, 444 (1987), W5013 (1991)
7. V. Innocente, M. Maire, and E. Nagy: CERN program library, IT-ASD W5013 (1991)
8. N. Amapane, M. Fierro, and M. Konecki: CMS-NOTE 2003/040 (2003)
9. N. Amapane et al: CMS-NOTE 2003/041 (2003)
10. C. Seez: *"The CMS trigger system"*, these proceedings (2003)
11. V. Ciulli: *"CMS B-Physics Reach"*, these proceedings (2003)