Chapter 9

Muons

9.1 Muon reconstruction

The muon reconstruction software is able to perform reconstruction in the muon system and the silicon tracker. The software has been designed using the concept of *regional reconstruction* in order to allow its use in both the offline reconstruction and the High-Level Trigger [8] (online event selection).

Employing regional reconstruction results in very small parts of the detector actually needing to be reconstructed before a physics object is validated. As an example, the amount of information needed to reconstruct a muon track in the silicon tracker represents less than a few per cent of the total tracker data volume. The software does not reconstruct tracks in the entire tracker, but only in that part which can possibly be involved in the reconstruction of a charged particle track compatible with the hits in the muon chambers. This results in savings on the overall CPU power needed to process the events. The method depends strongly on the identification of a good "seed," providing initial values of the 5 trajectory parameters and their errors, that can start the reconstruction with high efficiency and reliability. In the online environment these seeds are provided by the Level-1 Trigger system [7].

The muon reconstruction algorithm used by the HLT is seeded by the muon candidates found by the Level-1 muon trigger, including those candidates that did not necessarily lead to a Level-1 trigger accept. These seeds define a region of interest in the muon system, in which local reconstruction is performed. For offline reconstruction a different seed-generation algorithm has been developed, which performs local reconstruction in the entire muon system and uses patterns of segments reconstructed in the CSC and/or DT chambers as initial seeds. Muon reconstruction is performed in 3 stages: local reconstruction (local-pattern recognition), standalone reconstruction and global reconstruction. Starting from a seed, the chambers compatible with the seed are identified and local reconstruction is performed only in these chambers. Standalone muon reconstruction uses only information from the muon system, while global-muon reconstruction uses also silicon tracker hits. The HLT standalone and global reconstruction are called Level-2 and Level-3 reconstruction, respectively.

9.1.1 Standalone muon reconstruction

The standalone/Level-2 muon reconstruction uses only data from the muon detectors—the silicon tracker (Chapter 6) is not used. Both tracking detectors (DT and CSC) and RPCs participate in the reconstruction. Despite the coarser spatial resolution, the RPCs complement the tracking chambers, especially where the geometrical coverage is problematic, mostly in the barrel-endcap overlap region.

The reconstruction starts with the track segments from the muon chambers obtained by the local reconstruction (Section 3.5). The state vectors (track position, momentum, and direction) associated with the segments found in the innermost chambers are used to seed the muon trajectories, working from inside out, using the Kalman-filter technique [232]. The predicted state vector at the next measurement surface is compared with existing measurements and updated accordingly. In the barrel DT chambers, reconstructed track segments are used as measurements in the Kalman-filter procedure. In the endcap CSC chambers, where the magnetic field is inhomogeneous, the individual reconstructed constituents (three-dimensional hits) of the segments are used instead. Reconstructed hits from the RPC chambers are also included. 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 state is propagated from one station to the next using the GEANE package [233], which takes into account the muon energy loss in the material, the effect of multiple scattering, and the nonuniform magnetic field in the muon system. The track parameters and the corresponding errors are updated at each step. The procedure is iterated until the outermost measurement surface of the muon system is reached. A backward Kalman filter is then applied, working from outside in, and the track parameters are defined at the innermost muon station. Finally, the track is extrapolated to the nominal interaction point (defined by the beam-spot size: $\sigma_{xy} = 15\,\mu$m and $\sigma_z = 5.3\,$cm) and a vertex-constrained fit to the track parameters is performed.

9.1.2 Global muon reconstruction

The global/Level-3 muon reconstruction consists in extending the muon trajectories to include hits in the silicon tracker (silicon strip and silicon pixel detectors). Starting from a standalone reconstructed muon, the muon trajectory is extrapolated from the innermost muon station to the outer tracker surface, taking into account the muon energy loss in the material and the effect of multiple scattering. The GEANE package is currently used for the propagation through the steel, the coil and the calorimeters. Silicon layers [6] compatible with the muon trajectory are then determined, and a region of interest within them is defined in which to perform regional track reconstruction. The determination of the region of interest is based on the track parameters and their corresponding uncertainties of the extrapolated muon trajectory, obtained with the assumption that the muon originates from the interaction point. The definition of the region of interest has a strong impact on the reconstruction efficiency, fake rate, and CPU reconstruction time: well measured muons are reconstructed faster and with higher efficiency than poorly measured ones.

Inside the region of interest, initial candidates for the muon trajectory (regional seeds) are built from pairs of reconstructed hits. The 2 hits forming a seed must come from 2 different tracker layers, and all combinations of compatible pixel and double-sided silicon strip layers are used in order to achieve high efficiency. In addition, a relaxed beam-spot constraint is applied to track candidates above a given transverse momentum threshold to obtain initial trajectory parameters. Starting from the regional seeds, a track-reconstruction algorithm, based on the Kalman-filter technique, is used to reconstruct tracks inside the selected region of interest. The track-reconstruction algorithm consists of the following steps: trajectory building (seeded pattern recognition), trajectory cleaning (resolution of ambiguities) and trajectory smoothing (final fit). In the first step, the trajectory builder transforms each seed into a set of trajectories. Starting from the innermost layer, the trajectory is propagated to the next

tracker reachable layer, and updated with compatible measurements found on that layer. In the second step, the trajectory cleaner resolves ambiguities between multiple trajectories that may result from a single seed on the basis of the number of hits and the χ^2 of the track fit. In the final step, all reconstructed tracks are fitted once again, without a beam-spot constraint, using the hits in the muon chambers from the original standalone reconstruction together with the hits in the silicon tracker. To resolve possible ambiguities a second cleaning step is performed which selects the final muon candidates on the basis of a χ^2 cut.

The selected trajectories are then refitted excluding measurements (hits or segments) with high χ^2 values in muon stations with high hit occupancy. In addition the trajectories are refitted using only silicon tracker hits plus hits in the innermost muon station (excluding RPC hits) and the χ^2 probability of the fit is compared with the χ^2 probability of the tracker-only trajectory in order to detect muon bremsstrahlung or any kind of significant energy loss of the muon before the first muon station. This procedure is important for the accurate momentum reconstruction of very high-p_{T} (TeV) muons. The precise reconstruction of these objects is very challenging because of catastrophic energy loss and severe electromagnetic showers in the muon system.

9.1.3 Performance

9.1.3.1 Ideal detector

This subsection documents the *ideal* offline reconstruction performance, i.e., without taking into account effects from misalignment, miscalibration, neutron background or pile-up. Simulated single-muon samples with both fixed and continuous values of p_{T} (between 10 GeV/c and 1 TeV/c), and with a flat distribution in η and ϕ are used for this study. For the dimuon reconstruction performance studies, the following event samples were used: $J/\psi \to \mu^+\mu^-$, $Z \to \mu^+\mu^-$, $Z' \to \mu^+\mu^-$ (with Z' masses of 1, 1.5, 3 and 5 TeV/c^2) and various Drell-Yan dimuon samples generated with mass cut-offs from 70 to 1000 GeV/c^2.

Figures 9.1a and 9.1b show the reconstruction efficiency as a function of pseudorapidity for various values of p_{T}. Simulated single-muon samples were used to obtain the results shown in Figs. 9.1–9.4. The reconstruction efficiency achieved is typically 95–99%, except in the pseudorapidity regions around $|\eta| = 0.25$ and $|\eta| = 0.8$ (the regions between 2 DT wheels) and around $|\eta| = 1.2$ (the transition region between the DT and CSC systems), where the efficiency drops. (Figure 1.6 displays the geometric structure of the muon system.)

Figure 9.2 shows the transverse momentum resolution as a function of pseudorapidity for muons reconstructed with the standalone reconstruction algorithm (a) and the global reconstruction algorithm (b). The transverse momentum resolution was obtained by a Gaussian to the distribution of the quantity

$$\frac{q^{\mathrm{rec}}/p_{\mathrm{T}}^{\mathrm{rec}} - q^{\mathrm{gen}}/p_{\mathrm{T}}^{\mathrm{gen}}}{q^{\mathrm{gen}}/p_{\mathrm{T}}^{\mathrm{gen}}}, \tag{9.1}$$

where q is the charge and $p_{\mathrm{T}}^{\mathrm{gen}}$ and $p_{\mathrm{T}}^{\mathrm{rec}}$ are the generated and reconstructed transverse momenta, respectively.

The standalone 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 interaction point as the origin of the muon. This measurement is dominated by multiple scattering

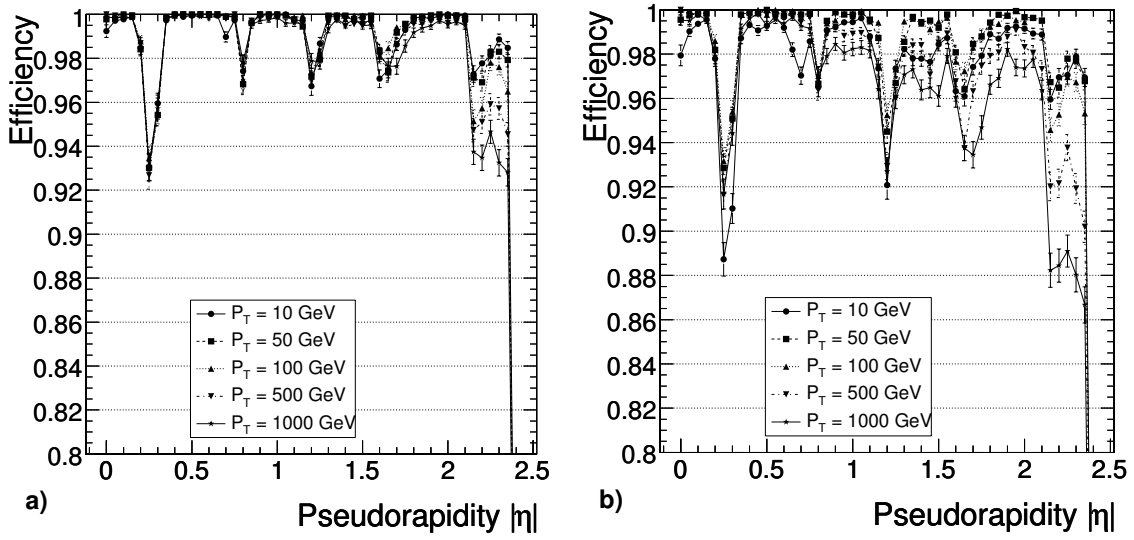

Figure 9.1: Muon reconstruction efficiency as a function of pseudorapidity for various values of p_{T}. a) Standalone reconstruction (using only hits from the muon system with a vertex constraint); b) Global reconstruction (using hits from the muon system and the tracker).

in the material before the first muon station up to p_{T} values of 200 GeV/c, when the chamber spatial resolution starts to dominate. For low-momentum muons, the best momentum resolution is given by the resolution obtained in the silicon tracker. Using measurements from the silicon tracker in addition to the muon system substantially improves the p_{T}-resolution (a factor of 10) compared to the resolution obtained by the standalone muon reconstruction.

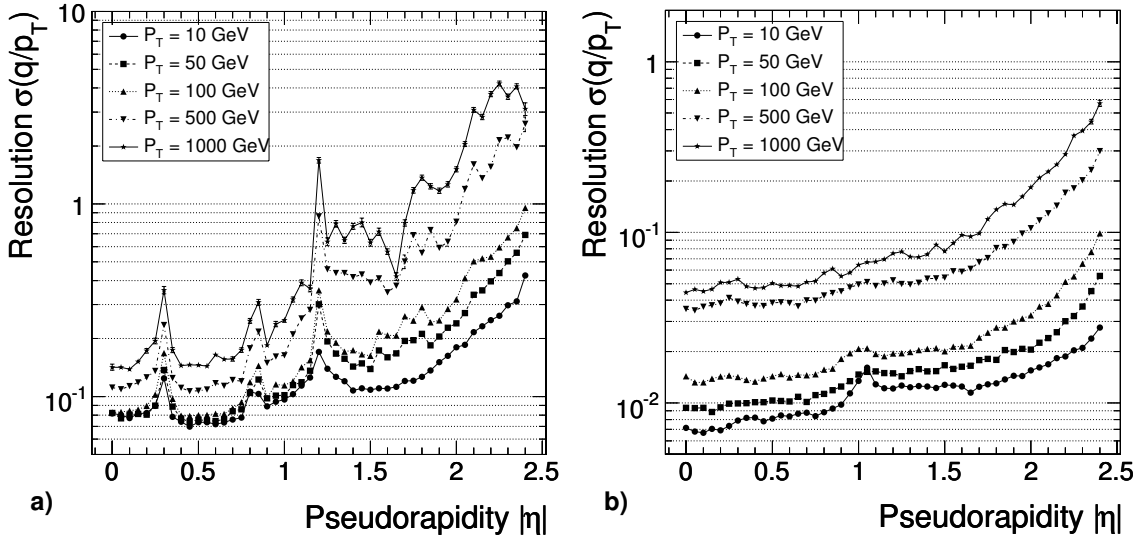

Figure 9.2: The q/p_{T} resolution for standalone (a) and globally (b) reconstructed muons (combined muon system and silicon tracker) as a function of pseudorapidity.

Figure 9.3 shows a comparison of the momentum resolution obtained from standalone and global reconstruction as a function of momentum. Two values of η have been chosen to illustrate the effect in the barrel and the endcap regions.

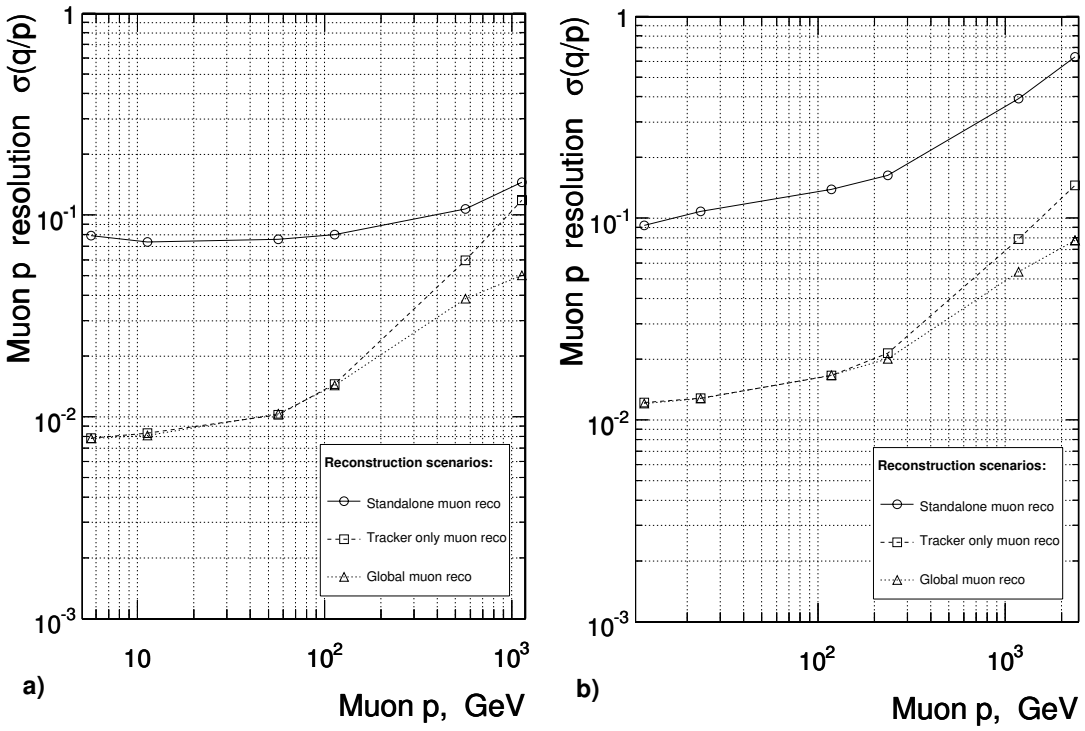

Figure 9.3: Resolution of $(1/p)$ versus p for standalone, global and tracker-only reconstruction. a) Barrel, $\eta = 0.5$; b) Endcap, $\eta = 1.5$.

In addition to an accurate measurement of its momentum, it is important that a reconstructed muon has the correct charge assignment. The probability to reconstruct muons with misassigned charge is shown in Fig. 9.4 as a function of pseudorapidity for various values of p_{T} for both standalone and global reconstruction. When the measurements from the silicon tracker are included, the percentage of misassigned charges for muons with $p_{\mathrm{T}} = 100$ GeV$/c$ is less than 0.1%.

9.1.3.2 Realistic detector

The muon reconstruction performance, taking into account effects from detector misalignment and event pile-up, is discussed in this Section. Events are produced using the full detector simulation package (OSCAR), which simulates the particle propagation and interactions through the detector (Section 2.5). A detailed description of the detector geometry, detector materials, and the magnetic field was used. The effect of pile-up has been taken into account assuming a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

Misalignment of muon chambers is introduced using a dedicated software package that provides the means to modify the simulated measurements after digitization. The package allows different misalignment displacements to be run, using the same set of digitized data. It also provides a description of 2 predefined misalignment scenarios, called the "First-Data-Taking Scenario" and the "Long-Term Scenario." The first one corresponds to the alignment at the beginning of data taking, while the second one describes the situation when all alignment procedures have enough data to obtain a full set of alignment constants. A detailed description of the alignment scenarios can be found in Section 3.2.2.

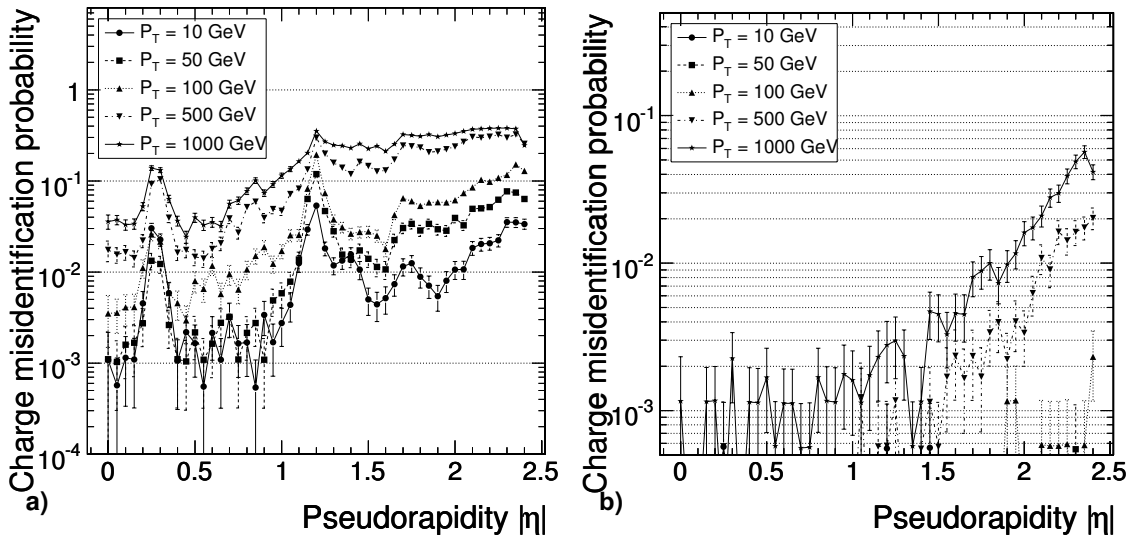

Figure 9.4: Muon charge-misidentification probability as a function of pseudorapidity for a) standalone and b) global reconstruction.

Figures 9.5, 9.6 and 9.7 show the effect of misalignment for standalone reconstruction. The effect of misalignment is demonstrated in terms of reconstruction efficiency, transverse momentum resolution and charge misassignment. The results were obtained using simulated single-muon samples. Figures 9.8, 9.9 and 9.10 show the same distributions for global muon reconstruction, using hits from the muon system and the silicon tracker.

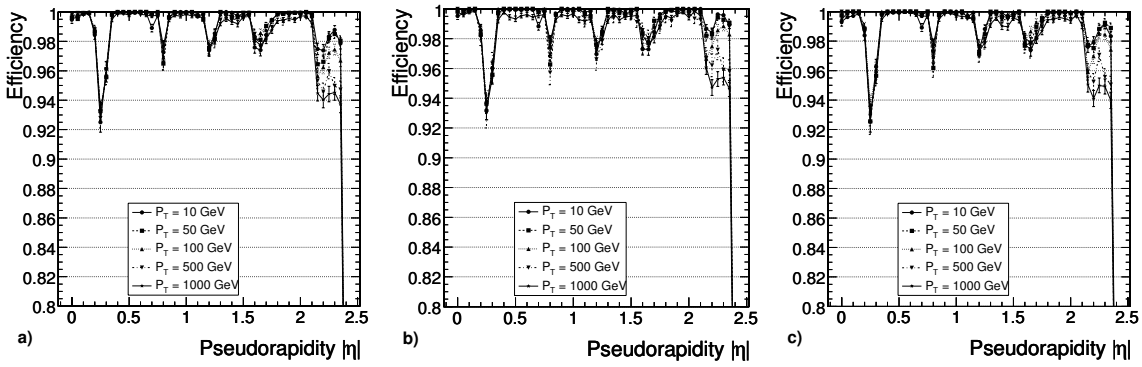

Figure 9.5: Standalone muon reconstruction efficiency as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the "First-Data-Taking Scenario;" and c) aligned according to the "Long-Term Scenario."

Figure 9.11 shows the dimuon reconstruction efficiency for various physics channels assuming different detector misalignment scenarios. The efficiencies shown are offline reconstruction efficiencies obtained from events having 2 muons within the muon detector acceptance. No requirement on the charges of the 2 muons was imposed in the reconstruction of these dimuon resonances. Figure 9.12 displays the corresponding dimuon invariant mass resolution. Samples of $J/\psi \rightarrow \mu^+\mu^-$ from $B_S^0 \rightarrow J/\psi\ \phi$, $Z^0 \rightarrow \mu^+\mu^-$, $Z' \rightarrow \mu^+\mu^-$ and Drell-Yan events were used to obtain these results.

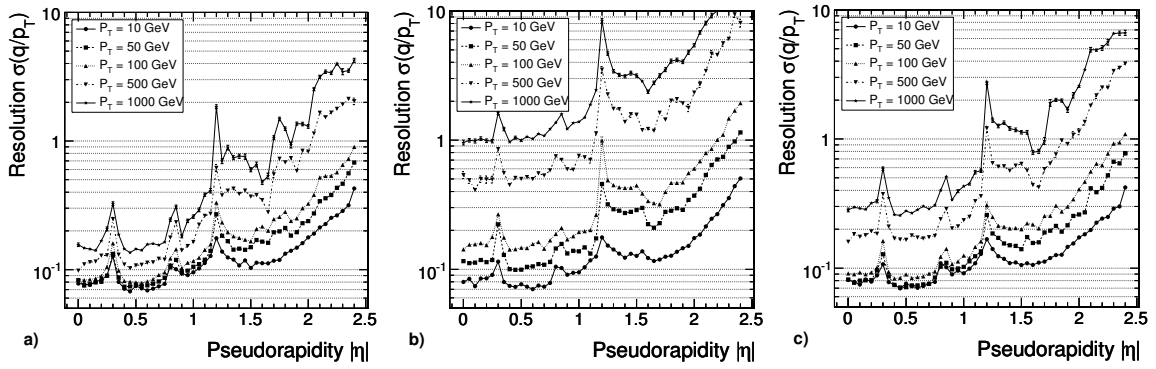

Figure 9.6: (q/p_{T}) resolution for standalone reconstructed muons (using only hits from the muon system with a vertex constraint) as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

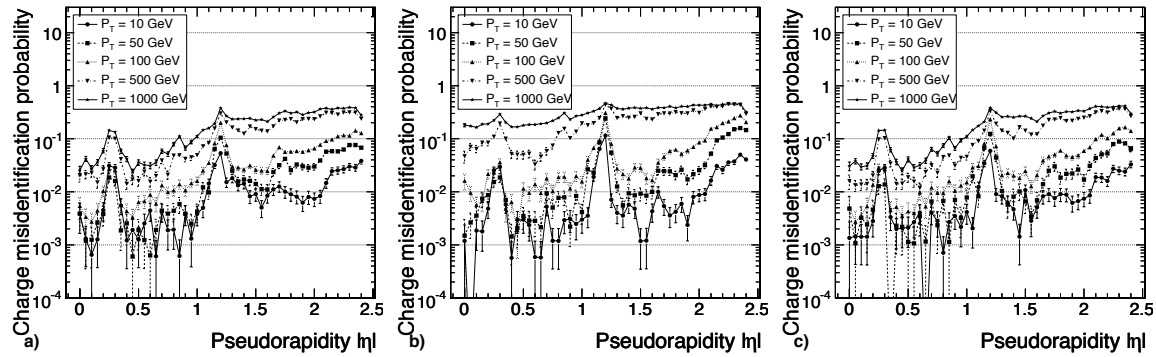

Figure 9.7: Muon charge misidentification probability for standalone reconstruction as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

Figure 9.13 shows the reconstructed invariant mass distributions for $J/\psi \rightarrow \mu^+\mu^-$ from $B_S^0 \rightarrow J/\psi\ \phi$ decays for different misalignment scenarios. Figure 9.14 displays the same distributions for $Z' \rightarrow \mu^+\mu^-$ decays, assuming a Z' mass of 1 TeV/c^2.

9.1.4 Cosmic-ray and beam-halo muon reconstruction

Beam-halo and cosmic-ray muons are 2 sources of particles that do not originate from LHC collisions but can nevertheless be found in the CMS detector. Whereas beam-halo muons move predominantly along the beam axis, cosmic-ray muons typically traverse the detector from top to bottom. While such muons have to be considered as sources of background, they can be employed for many purposes, for instance alignment, calibration, and the understanding of trigger and reconstruction efficiencies.

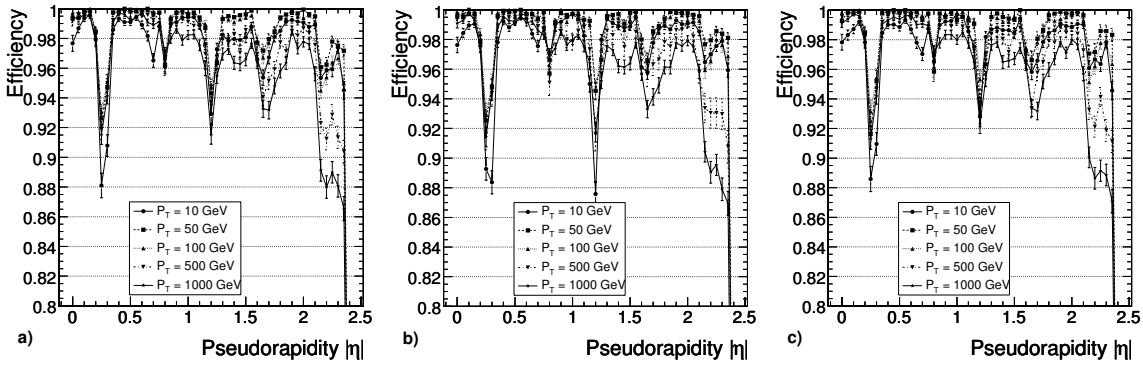

Figure 9.8: Global muon reconstruction efficiency as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

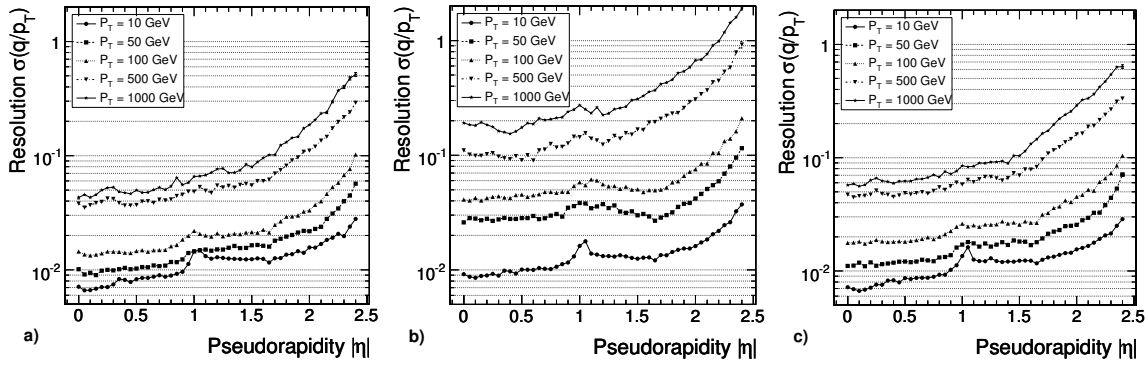

Figure 9.9: (q/p_{T}) resolution for global reconstructed muons (using only hits from the muon system with a vertex constraint) as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

9.1.4.1 Event generation and detector simulation

The simulation of beam-halo and cosmic-ray muons, described in detail in [199], is performed in two steps. The muons are first generated with $E(\mu^{\pm}) > 10$ GeV, and they are then passed through the CMS detector simulation.

Cosmic-ray muons are generated at the surface of the earth according to measurements described in [119], and are then propagated to the surface of the CMS detector, taking into account the energy loss in the material below the surface of the earth.

Beam-halo muons are generated according to kinematic distributions that have been provided by the LHC simulation group [234]. It should be stressed that the rates given in [234] are valid only for one specific set of LHC parameters, which is obsolete now. The absolute rates given here are thus to be understood as rough estimates. At the entrances of the LHC tunnel, the muons are passed to the detector simulation, which includes the simulation of the forward shielding.

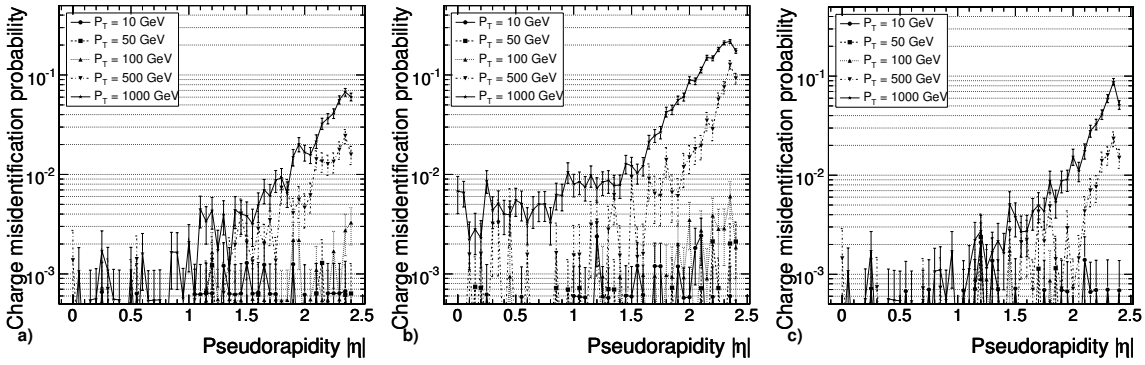

Figure 9.10: Muon charge misidentification probability in global reconstruction as a function of pseudorapidity for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long Term Scenario.

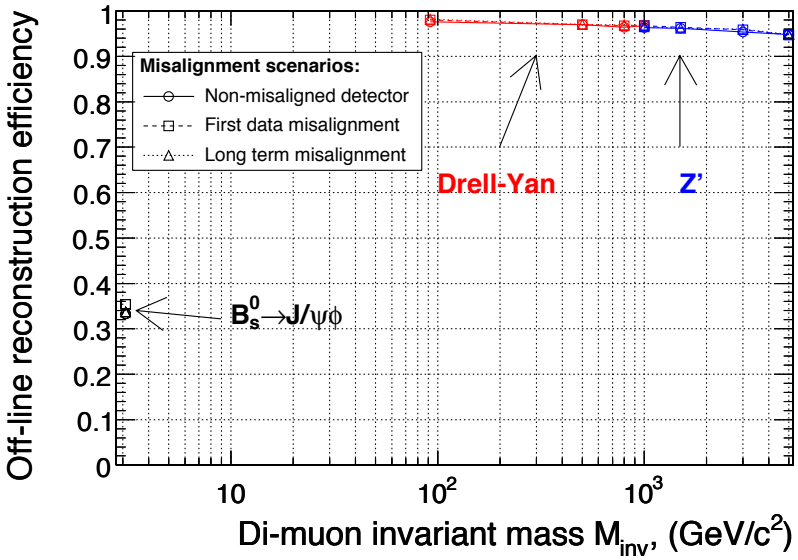

Figure 9.11: Dimuon reconstruction efficiency for different physics channels:: $J/\psi \rightarrow \mu^+\mu^-$ from $B_S^0 \rightarrow J/\psi\ \phi$, $Z^0 \rightarrow \mu^+\mu^-$ and $Z' \rightarrow \mu^+\mu^-$, and for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking scenario; c) aligned according to the Long-Term Scenario.

9.1.4.2 Level-1 trigger and standalone reconstruction

Two scenarios are discussed here, a "single-beam run" to study beam-halo muons and a "cosmic run" for cosmic-ray muons in the CMS cavern. Since an event consists of a single muon, a Level-1 trigger requirement is sufficient to control the event rate. In order to maximize the Level-1 rate, the look-up tables (LUT) for the CSC and DT detectors are widened as much as possible, which allows the triggering of muons with at least 2 hits and almost no angular constraint.

Because beam-halo and cosmic-ray muons enter the CMS detector from outside and do not

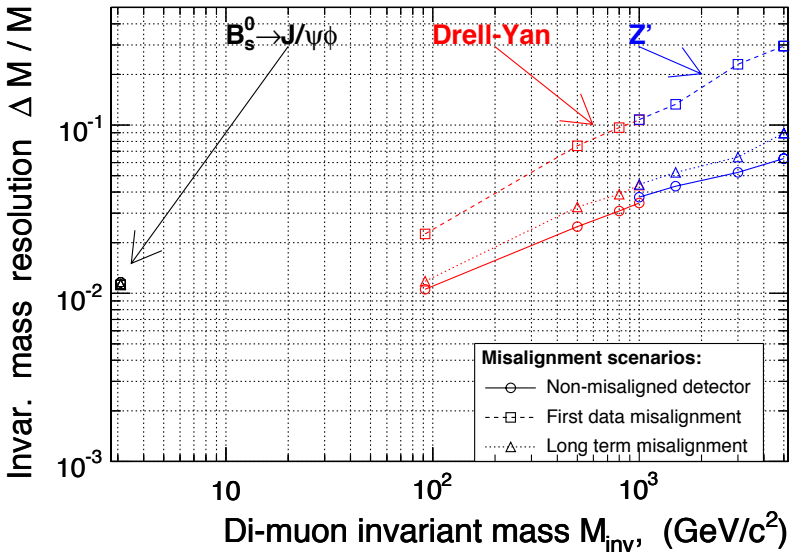

Figure 9.12: Dimuon invariant mass resolution for different physics channels: $J/\psi \to \mu^+\mu^-$ from $B_S^0 \to J/\psi\ \phi$, $Z^0 \to \mu^+\mu^-$ and $Z' \to \mu^+\mu^-$, and for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

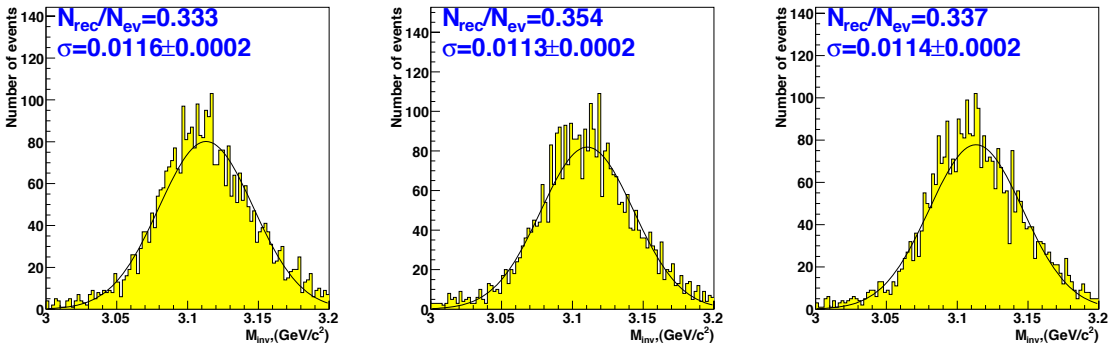

Figure 9.13: Reconstructed invariant mass distributions for $J/\psi \to \mu^+\mu^-$ from $B_S^0 \to J/\psi\ \phi$ decays for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

come from the interaction point, the time-of-flight patterns are substantially different from muons produced in pp collisions. The different timing requires a different synchronization of the detector. This is done separately for beam-halo and cosmic-ray muons by assigning an average, constant time-of-flight offset to each muon chamber.

Whereas most of the beam-halo muons are expected to have the same bunch structure as the proton beam itself, the bunch crossing assignment is rather artificial for the case of cosmic-ray muons. The Level-1 trigger rate increases by about 16% when the neighbouring bunches are added to the bunch with the highest rate.

The standalone (SA) muon reconstruction algorithm turns out to be flexible enough to per-

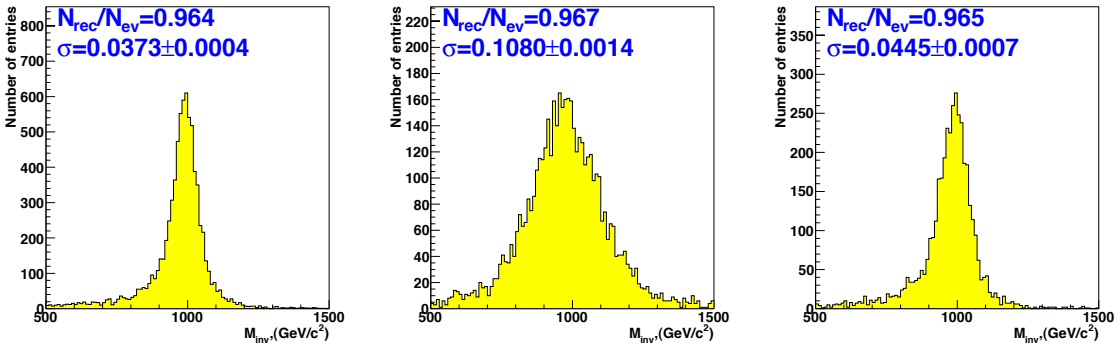

Figure 9.14: Reconstructed invariant mass for $Z' \rightarrow \mu^+\mu^-$ decays for different misalignment scenarios: detector a) perfectly aligned; b) aligned according to the First-Data-Taking Scenario; c) aligned according to the Long-Term Scenario.

form well on cosmic-ray muons. For beam-halo muons, the constraint on the pseudorapidity has to be ignored. At least 3 reconstructed hits are required for all SA muon tracks.

9.1.4.3 Kinematic properties and rate estimates

An important property of both beam-halo and cosmic-ray muons is that, if they are energetic enough, they can lead to 2 SA muon tracks. Beam-halo muons, as shown in Fig. 9.15, move almost parallel to the beam axis and lead to tracks in the forward muon detectors. When 2 SA muon tracks come from 1 beam-halo muon, one track is reconstructed on each side of the detector. This is useful for inter-alignment of the 2 halves of the forward muon chambers.

If only 1 SA muon track is found, it is most probably on the side from which the beam-halo muon enters the CMS detector. The opposite side is much less populated because many muons are stopped in the detector material, mainly the iron yoke, the magnet and the calorimeters, on their way through CMS.

The kinematic properties of cosmic-ray muons are again different, as illustrated in Fig. 9.16. A cosmic-ray muon typically traverses the CMS detector from top to bottom, and the majority of the tracks are reconstructed in the barrel region.

The majority of tracks are reconstructed in the barrel region. There are more single SA muon tracks reconstructed on the top than the bottom, due to the energy loss in the detector material. Cosmic-ray muons with 2 reconstructed tracks connect the upper half with the lower half of the CMS detector, including the inner tracker in cases where the muons are central enough.

In general, events with 2 SA muon tracks are useful for studying trigger and reconstruction efficiencies, energy loss in the detector material, and alignment on a large scale. Events with only 1 SA muon track are expected to add precision to the alignment of neighbouring muons chambers. The number of triggered events of each type are summarized in Table 9.1. It should be kept in mind that all rates depend on the trigger conditions and the track requirements.

In summary, beam-halo and cosmic-ray muons can be triggered and reconstructed with the CMS muon system. Both sources of muons are useful for studying and understanding the

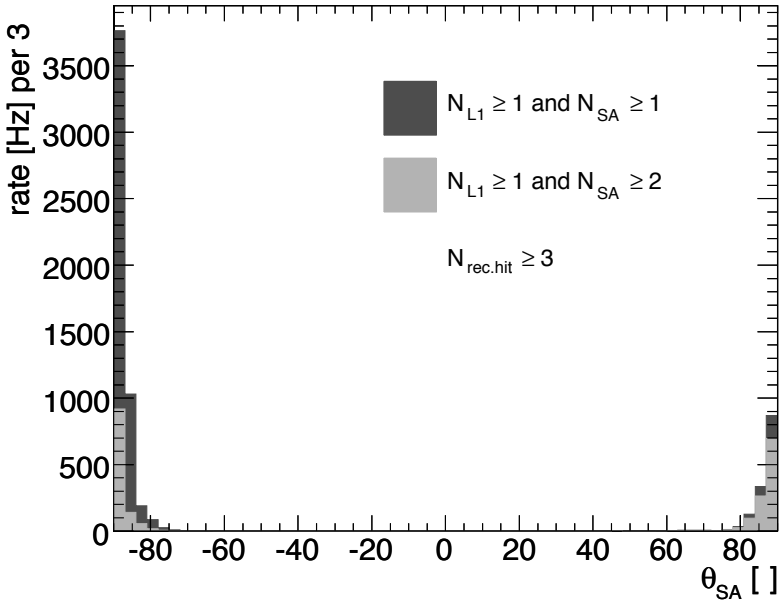

Figure 9.15: Rates of SA muon tracks as a function of θ_{SA} for beam-halo muons accepted by the Level-1 trigger. The beam-halo muons come from the $-z$ ($-\theta_{\mathrm{SA}}$) side and move to the $+z$ ($+\theta_{\mathrm{SA}}$) side.

Table 9.1: Event rates for beam-halo and cosmic-ray muons. All rates are for events accepted by the Level-1 trigger. Events with reconstructed SA muons are divided into 2 categories: events with exactly 1 SA muon track and events with exactly two SA muon tracks.

event rate	Level-1	Level-1 with 1 SA track	Level-1 with 2 SA tracks
beam-halo muons	6600 Hz	4300 Hz	1100 Hz
cosmic-ray muons	680 Hz	330 Hz	140 Hz

performance of the detector, especially the muon trigger system. Furthermore, these muons can help align and calibrate the CMS detector. A comprehensive study with many more details is given in [118].

9.1.5 Dimuon reconstruction in heavy-ion collisions

The reconstruction of muons in the heavy-ion environment requires specialized algorithms due to the high hit occupancy in the tracker [235]. In addition, it is important to extend the detector acceptance to the lowest-p_{T} muons. The relatively low expected luminosity in heavy-ion collisions allows a less stringent trigger criteria as compared to the pp case (Section 3.4.7).

The muon tracks in the muon chambers are used to identify hits in the outer tracker that can form the starting points (seeds) for the matching muon candidate tracks. The initial momentum value of the muon candidate is taken from the trigger momentum estimation. If

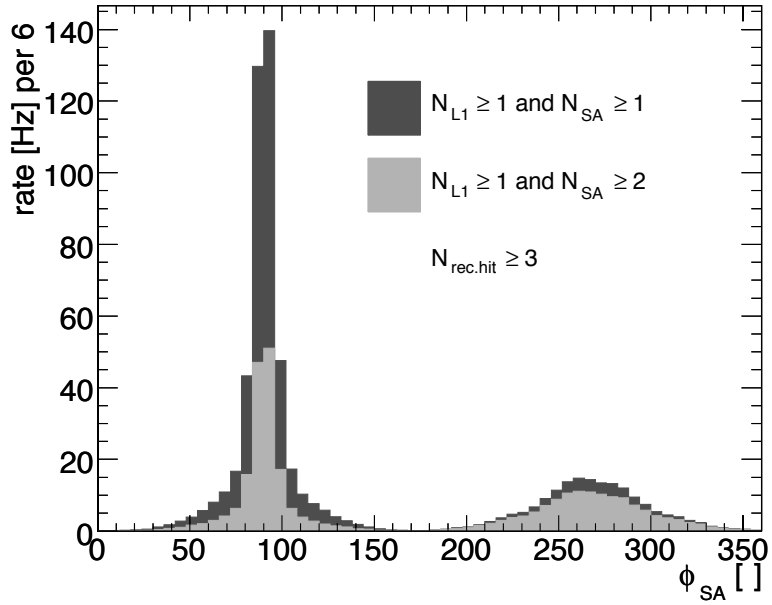

Figure 9.16: Rates of SA muon tracks as a function of ϕ_{SA} for cosmic-ray muons accepted by the Level-1 trigger. The maximum rate is at the top of the detector, which corresponds to $\phi_{\mathrm{SA}} = 90°$. The bump around $\phi_{\mathrm{SA}} = 270°$ is at the bottom of the CMS detector.

the candidate has both Level-1 and Level-2 momentum estimations, the momentum obtained by the Level-2 trigger is used. Only hits located within a search window are considered as seeds.

The track searching and fitting is done in several passes and is performed simultaneously for both muons forming the dimuon pair. During the first pass, both muon track candidates are required to have hits in all the outer silicon and pixel layers. For the successful candidates, a Kalman-fitting procedure is performed inside the tracker volume for both muons of the pair. Muon candidates that satisfy a χ^2 quality criterion are propagated back to the muon stations and are ordered according to their distance to the muon track segments. For the surviving pairs, a common vertex is calculated using a Principal Vertex Finder. In the final step, the pair that has the best quality (the smallest value of $\chi^2_{\mu_1} \times \chi^2_{\mu_2} \times \chi^2_{\mathrm{vertex}}$ and the smallest distance to the muon track segments) is chosen. If there are no pairs that satisfy the above criteria, the propagation pass is repeated, this time allowing one missing tracker layer for 1 of the muons. The selection procedure is terminated if either a pair is found or all possible combination of hits with 1 missing layer is considered. The detailed description of the track-finding procedure is given in [236].

9.2 Muon identification

Standard muon track reconstruction, as described in Section 9.1, starts from the muon system and combines standalone reconstructed muon tracks with tracks reconstructed in the silicon tracker. The muon identification algorithm provides an additional tool for selecting muon candidates using an inside-out algorithm that is by design complementary to the stan-

dard muon reconstruction. The algorithm starts with all reconstructed tracks from the silicon inner tracking detector and attempts to quantify a muon compatibility for each (including those tracks with no matching standalone tracks in the muon detectors). Muon identification also takes advantage of information from other detectors not used in reconstruction, such as associated energy deposits in the calorimeters for determining the compatibility of a given track with the muon hypothesis. In addition, the algorithm is able to make use of hit information from individual layers of the muon detectors, even in cases where those hits are not associated with a reconstructed, standalone muon track. Lower-p_{T} muons, for example, that range out within the iron yoke before depositing hits in the outer muon detector layers are difficult to reconstruct with the standalone muon track reconstruction. However, this class of muons can potentially be identified offline via this algorithm by matching the hits found in the inner muon detector layers with reconstructed silicon tracks and examining the associated calorimeter energy deposits to see if they are compatible with those from a minimum-ionizing particle.

The muon identification algorithm is useful for physics analyses involving event signatures that contain multiple muons (e.g., $H \rightarrow ZZ^* \rightarrow \mu^+\mu^-\mu^+\mu^-$), where the reconstruction conditions can be relaxed for some muons. If the muon identification algorithm is used to identify one or more of the muons in these channels rather than requiring all muons to be fully reconstructed, the overall selection efficiency for these events is increased. The muon identification algorithm can be particularly useful in cases where one has a soft p_{T} spectrum for the lowest-energy muon in the event since these muons are more likely to range out in the detector and not be fully reconstructed. Even for event topologies in which all of the final state muons are fully reconstructed, the information provided by the identification algorithm helps in understanding muon fake rates and solving potential ambiguities at the analysis level. In addition, the algorithm will be useful during commissioning for picking out a class of potential, unreconstructed muon candidates for studying and improving the reconstruction software.

9.2.1 Algorithm

The basic idea of the muon identification algorithm is to extrapolate each reconstructed silicon track outward to its most probable location within each detector of interest (ECAL, HCAL, HO, muon system). The track extrapolation takes into account the magnetic field and energy loss via the GEANE software package [233]. The algorithm then searches in a cone around that most probable location for any signals that might be associated with the track of interest. The cone sizes used for making these searches are different for each detector and are parameterized as functions of the p_{T} and η of the silicon track. Optimization of cone sizes is based on studies of simulated event samples. After collecting the associated signals from each detector, the algorithm determines a combined compatibility value corresponding to how well the observed signals fit with the hypothesis that the silicon track used to seed the algorithm is produced by a muon. The assignment of the compatibility value is somewhat arbitrary and based primarily on studies of various simulated event samples. The compatibility variable output by the algorithm is constructed so that different physics analyses can make different cuts on the minimum value required to identify muons. In choosing this threshold, each analysis can independently optimize the trade off between muon selection efficiency and the background fake rate for muons.

9.2.1.1 Identification in the calorimeters

The minimum-ionizing signature of muons in the calorimeters can be used to identify reconstructed tracks produced by muons. The identification algorithm looks for energy deposits associated with candidate tracks in both the electromagnetic and hadron calorimeters and compares the observed energies with those expected for a minimum-ionizing particle. In the central region of the detector ($|\eta| < 1.2$), the algorithm also checks if a signal is observed in the additional outer hadron calorimeter (HO), a scintillator layer located just outside of the superconducting solenoid in front of the inner muon chambers. Particles that enter this layer, having passed through the material in the calorimeter modules and magnet, are more likely to be minimum-ionizing. In addition, any excess energy deposit in cells from sources such as pile-up events should have less of an effect in the outer layers. There should, therefore, be a relatively lower probability of misidentification in the outer layers of the calorimeters, especially in the HO.

The muon identification algorithm extrapolates the seed track to its most likely location within each calorimeter volume. In the case of the HO layer, the location of the extrapolated track at the innermost muon chambers is used. The algorithm determines the total energy associated with the track candidate by summing the energies in towers whose centers are within some distance $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$ of the extrapolated track position. The value of ΔR is optimized separately for each calorimeter to account for differences in tower granularities and to ensure that a big enough search cone is used to collect all of the energy associated with muons that either share their energy in adjacent cells or undergo significant multiple scattering. Optimization of cone sizes is also done independently for 3 separate bins in track η: barrel ($|\eta| < 0.8$), overlap ($0.8 < |\eta| < 1.2$), and endcap ($|\eta| > 1.2$) to account for differences in the calorimeter modules used in these regions. In order to avoid collecting energies not directly associated with the candidate track, the ΔR values used by the algorithm in each region are chosen to be as small as possible. However, they are not so small that a high collection efficiency for the energies associated with real muon tracks is not maintained. For the finely segmented electromagnetic calorimeter modules, studies of simulated single-muon event samples advocate optimal ΔR values between 0.02 and 0.03 (all energy deposited in 1 crystal). In the case of the more coarsely segmented hadron calorimeter modules and the HO scintillator layer, bigger ΔR values between 0.08 and 0.13 are found to be the optimal choice.

The energy thresholds for observing the minimum-ionizing calorimeter signals are also important. For the offline algorithm, the single-tower thresholds are made as low as possible so that the small minimum-ionizing signal can be observed, especially in the case of the signals from the electromagnetic calorimeter. The minimum-ionizing energy expected to be deposited in the crystals is on the order of 300 MeV in the barrel region and 400 MeV in the endcap. Due to the fact that a high threshold on the order of 500 MeV is applied during the online readout of the endcap crystals, it is not always possible to observe the minimum-ionizing signal in the forward-electromagnetic calorimeter modules. A much lower threshold of about 60 MeV is applied during the online readout of the barrel crystals, which means that there is a much higher efficiency for observing the minimum-ionizing signal in the central electromagnetic calorimeter modules. The difference is illustrated in Fig. 9.17. The top and bottom figures show the electromagnetic energy returned from the muon identification algorithm for single $p_{\mathrm{T}} = 10\ \mathrm{GeV}/c$ muon tracks in the barrel and endcap regions of the detector. In the endcap region, there are a substantial number of cases where the algorithm finds no electromagnetic energy associated with the track. For the hadron calorimeter modules, the online

readout thresholds are well below the threshold of the expected minimum-ionizing signal. Therefore, an offline single-tower threshold on the order of 700 MeV is applied within the identification algorithm for forming the hadron calorimeter energy sum. This threshold is applied to remove potentially noisy towers.

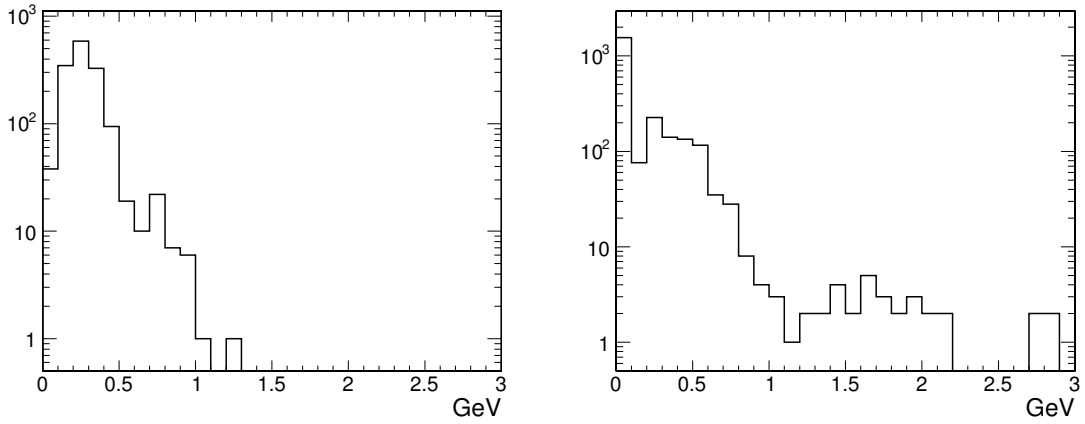

Figure 9.17: Electromagnetic calorimeter energy associated with single $p_{\mathrm{T}} = 10\ \mathrm{GeV}/c$ muon tracks as returned by the muon identification algorithm for the barrel (left) and endcap (right) regions of the detector. The significant number of zeroes returned for muons in the endcap region is due to the high online thresholds applied in the readout of the endcap crystals.

Based on the measured energies associated with candidate tracks, the muon identification algorithm calculates a compatibility value between 0 and 1 that attempts to describe how consistent these energies are with respect to what one expects for a muon. The compatibility value is obtained from a three-dimensional likelihood function of the form

$$\frac{P_S(x) \cdot P_S(y) \cdot P_S(z)}{P_S(x) \cdot P_S(y) \cdot P_S(z) + P_B(x) \cdot P_B(y) \cdot P_B(z)}\ , \tag{9.2}$$

where P_S and P_B are the signal and background probabilities as functions of the measured energies in the electromagnetic calorimeter (x), the hadron calorimeter (y), and the HO scintillator layer (z). The signal and background probability distributions are obtained from simulated samples of single-muons and pions, respectively. Independent distributions are obtained for the barrel, overlap, and endcap regions of the detector and for different track p_{T} ranges. Because there is no HO scintillator layer in the endcap region of the detector, the $P_S(z)$ and $P_B(z)$ functions for this region are set to one. The resulting compatibility values for tracks associated with simulated $p_{\mathrm{T}} = 10\ \mathrm{GeV}/c$ single muons and pions are shown in Fig. 9.18 for the barrel and endcap regions of the detector.

9.2.1.2 Muon detector identification

Reconstructed hit information from the muon detectors can be used to identify muon candidates even in cases where the standalone muon reconstruction fails. Low-p_{T} muons, in particular, can be bent back in the magnetic return field toward the interaction region or range out in the magnet yoke. In these cases, no hits are produced in the outermost layers of the muon detectors, and consequently standalone reconstruction is less efficient. Conversely,

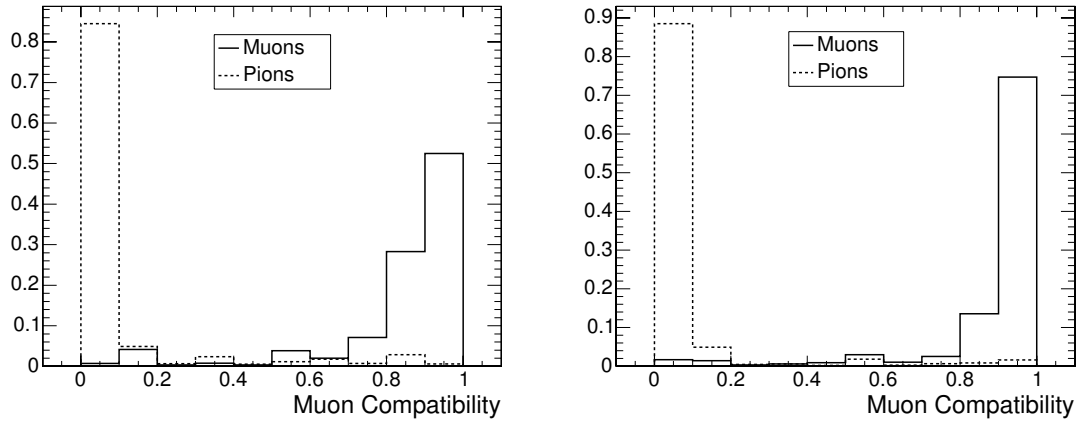

Figure 9.18: Muon compatibility values for tracks in the silicon tracker associated with simulated $p_{\mathrm{T}} = 10\ \mathrm{GeV}/c$ single muons (solid lines) and pions (dashed lines) in the barrel (left) and endcap (right) regions of the detector. The compatibility values shown here are based solely on observed energy deposits in the calorimeters matched with the candidate track.

muon detector hits coming from even a single inner layer that can be associated with a track reconstructed in the silicon tracker (thereafter simply called a "track") are potentially helpful in determining the muon compatibility of that track.

To incorporate muon detector hit information into the muon identification algorithm, tracks are extrapolated sequentially into each layer of the muon detector. A search road is defined to look for hits and reconstructed segments within the layer that may be compatible with the extrapolated trajectory. For each potential matching hit or segment, the algorithm calculates a χ^2 measuring the compatibility of the position and direction (for segments) with those of the extrapolated track at the corresponding muon detector layer. The algorithm defines hits and segments to be associated with the track if the corresponding χ^2 is below a programmable threshold. Using this definition, the algorithm outputs a list of matching hits and segments for each reconstructed track.

Using this list, the identification algorithm calculates a second muon compatibility value for each track based solely on muon detector hit information. The weight assigned to a single matching segment or hit depends on the dimensional content. Drift tube hits, for example, provide matching information in only 1 dimension. Drift tube chamber segments, on the other hand, contain matching information in 2 or even 4 dimensions if hits from all superlayers are incorporated. CSC and RPC hits provide slightly better matching in 2 dimensions, while CSC segments also provide ideal matching in all 4 dimensions. Matching hits and segments are also weighted by layer. Outer-layer hits and segments are more likely to be associated with a real muon since more material must be traversed to reach these layers, while inner-layer hits are more likely to be associated with punch-through of hadrons.

For each track, the algorithm assigns a compatibility value between 0 and 1 based on the number of matching muon detector segments and hits. The calculation is designed so that tracks matching to the best possible combination of hits and/or segments in each of the muon detector layers have a corresponding compatibility value of one. Of course, based on this definition, tracks with an associated compatibility value near 1 are almost sure to be globally reconstructed as muons. In the case of muon identification, therefore, cuts on the

compatibility value designed to recover muon reconstruction inefficiencies will need to be set well below this maximum value. Details of the algorithm are described in Ref. [237].

Based on the algorithm described above, the muon compatibility value distributions for tracks associated with simulated $p_{\mathrm{T}} = 10\,\mathrm{GeV}/c$ single muons and pions are shown in Fig. 9.19 for the barrel and endcap regions of the detector.

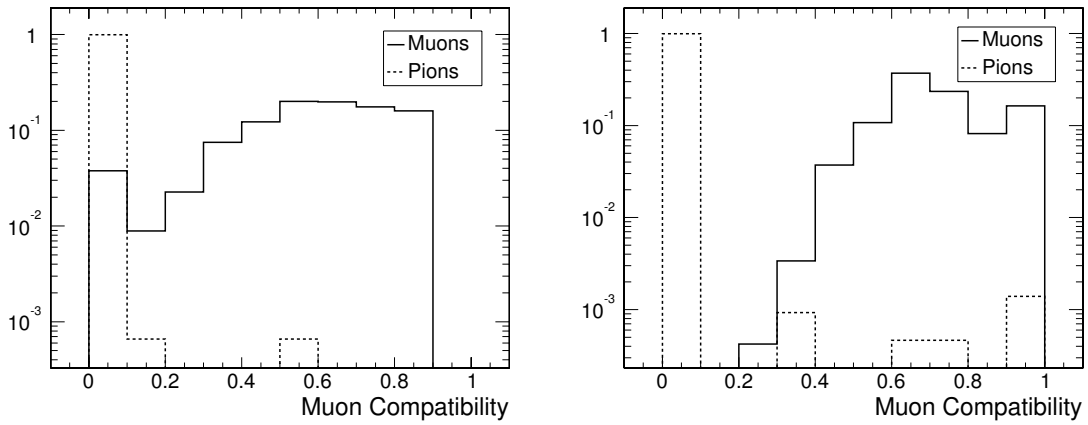

Figure 9.19: Muon compatibility values for tracks associated with simulated $p_{\mathrm{T}} = 10\,\mathrm{GeV}/c$ single muons (solid lines) and pions (dashed lines) in the barrel (left) and endcap (right) regions of the detector. The compatibility values shown here are based solely on matched hits and segments in the muon detectors.

9.2.2 Performance

The performance of the muon identification algorithm was studied on several simulated event samples. For these studies, a track is considered to be identified as a muon if its associated compatibility value based on energy deposits in the calorimeters is above 0.8 and its associated compatibility value based on hit and segment information in the muon detectors is greater than 0.4. These cut values are programmable and in general require optimization for specific algorithm applications. Here, the same cut values are used on several different types of event samples to provide a baseline overview of algorithm performance. Low-p_{T} muons are of particular interest since in a significant fraction of cases these muons do not produce hits in the outer muon detector layers and are therefore more difficult to reconstruct using the standard standalone algorithm. The algorithm was run on a sample of $50\,000$ simulated single-muon events. The muons were generated flat in ϕ and η (from -2.5 to 2.5) with a single-p_{T} value of 5 GeV/c. The single muons were generated on top of simulated pile-up events corresponding to a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The efficiency distributions for both global muon reconstruction and global muon reconstruction plus muon identification are shown in Fig. 9.20 (left), as a function of the simulated muon pseudorapidity. For the events in this sample, the efficiency for reconstructing single, low-p_{T} muons is 68.7%. Adding muon identification as an additional selection mechanism increases the overall selection efficiency for the muons in this sample by roughly 15% (for a combined efficiency of 78.6%). The muon identification algorithm was also run on all of the non-muon tracks in these events to investigate the rate at which the algorithm misidentified these tracks as muon candidates. The fake rate as a function of track η is shown in Fig. 9.20 (right). Only

30 out of over $250\,000$ non-muon tracks in these events are identified as muon candidates by the algorithm.

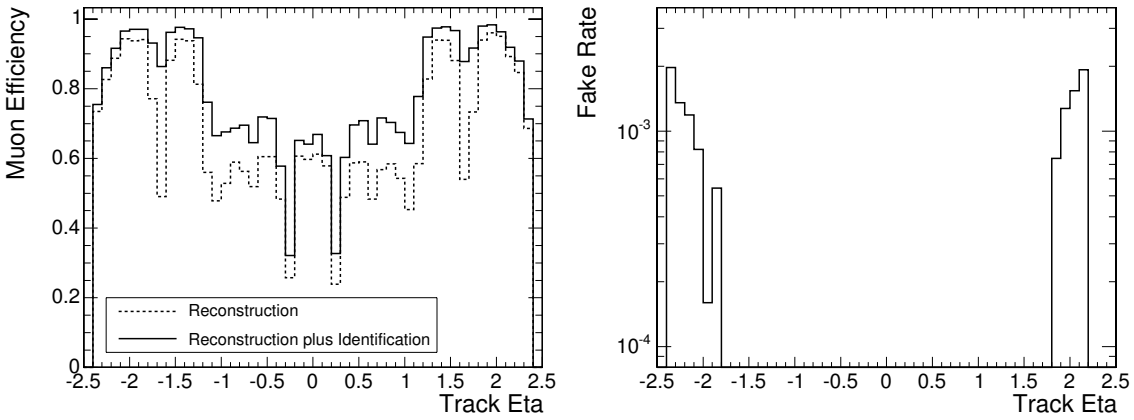

Figure 9.20: Efficiencies for global muon reconstruction (dashed lines) and global muon reconstruction plus muon identification (solid lines), as a function of pseudorapidity for $p_{\mathrm{T}} = 5\ \mathrm{GeV}/c$ single muons (left). Fake rate at which non-muon tracks in the same events are tagged as muon candidates using the muon identification algorithm also as a function of pseudorapidity (right). The observation of fake rates equal to 0 in the central part of the detector ($|\eta| < 1.8$) is due to insufficient statistics.

The muon identification algorithm was also run on a sample of simulated $H \rightarrow WW \rightarrow \mu^+\mu^-\nu\overline{\nu}$ events generated using a Higgs boson mass of $200\ \mathrm{GeV}/c^2$. The events include pile-up events corresponding to a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Due to the low overall production cross section for this process, observing a signal in the data requires maintaining high event selection efficiency. Since the transverse momentum of the lower-p_{T} muon is observed to be less than $10\ \mathrm{GeV}/c$ in roughly 20% of these events, muon identification can potentially be a useful tool for increasing the overall selection efficiency in this channel. Because the algorithm has not yet been optimized for selecting very low-p_{T} muons, this study is restricted to the subset of events where the p_{T} of both muon candidates is greater than $5\ \mathrm{GeV}/c$. The distributions in Fig. 9.21 (left) show the selection efficiency for events in this channel as a function of the simulated p_{T} of the lower-p_{T} muon, using both global reconstruction and global reconstruction plus muon identification. From these distributions, a net gain on the order of 5% is observed in the overall selection efficiency incorporating the muon identification algorithm. The rate at which non-muon tracks in these events are tagged as muons by the identification algorithm was also measured. The fake rate as a function of track p_{T} is shown in Fig. 9.21 (right). Out of over $750\,000$ non-muon tracks in these events, less than 125 are identified as muon candidates using the algorithm.

Low-p_{T} muons are also produced in the semi-leptonic decays of heavy quarks. The tagging of muons within jets (soft-lepton tagging) is therefore a useful tool in selecting jets originating from bottom quark decays (an important requirement for many analyses, and discussed further in Section 12.2.4). In order to see if the muon identification tool might be useful in improving the efficiency for soft-lepton tagging in bottom quark jets, the algorithm was run on a simulated sample of bottom quark jets with generated p_{T} values between 50 and $80\ \mathrm{GeV}/c^2$. The events were generated on top of simulated pile-up events corresponding to a

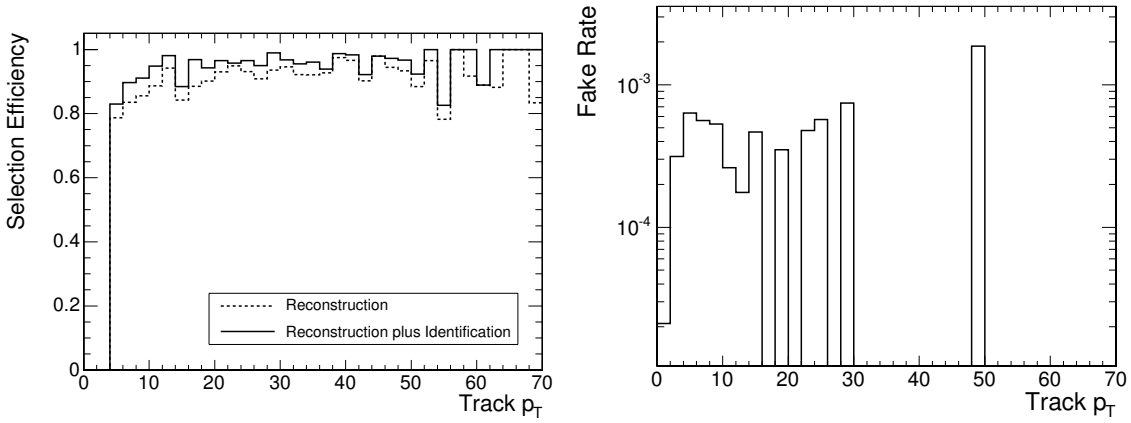

Figure 9.21: $H \to WW \to \mu^+\mu^-\nu\overline{\nu}$ event selection efficiencies using global muon reconstruction (dashed lines) and global muon reconstruction plus muon identification (solid lines), as a function of the p_{T} of the lower-p_{T} muon in the events (left). Fake rate at which non-muon tracks in the same events are tagged as muon candidates using the muon identification algorithm as a function of track p_{T} (right). The observation of fake rates equal to 0 for the majority of p_{T} values above 15 GeV/c is due to insufficient statistics.

luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The distributions shown in Fig. 9.22 are efficiency curves for both global muon reconstruction and global muon reconstruction plus muon identification for the muons produced in the semi-leptonic decays of the bottom quarks in these events, as functions of the simulated muon p_{T} and pseudorapidity. The efficiencies shown here are for muons with a minimum p_{T} of 5 GeV/c. The study is restricted to this range since the muon identification algorithm has not yet been optimized for selecting very low-p_{T} muons. The global reconstruction efficiency for these muons is 71% and incorporating the muon identification algorithm increases the overall selection efficiency to 84%. To be useful in soft-lepton tagging, the associated fake rate of the algorithm (the rate at which non-muon tracks within jets are identified as muons) must also be low since jets originating from both light and heavy quarks, as well as those from gluons, contain many non-muon tracks. The algorithm fake rate is studied using a sample of reconstructed pion tracks found within the subset of bottom quark jets from our simulated sample that contain no prompt muons. The efficiencies for tagging these pion tracks as muons using global muon reconstruction and global muon reconstruction plus muon identification are shown in Fig. 9.23 as functions of the simulated pion p_{T} and pseudorapidity. The measured fake rate at which the pion tracks within bottom quark jets are globally reconstructed as muons is 0.17%. Adding muon identification as an additional selection tool, increases the muon fake rate for pion tracks to roughly 0.20%. The small increase in the observed fake rate indicates that the muon identification tool can be potentially useful in soft lepton bottom quark tagging.

9.3 Muon isolation

9.3.1 Principles of the muon-isolation algorithms

Muons selected by a transverse-momentum cut in the muon trigger are mostly real muons. In the momentum range relevant for triggering (p_{T} threshold in the range 10–30 GeV/c), the

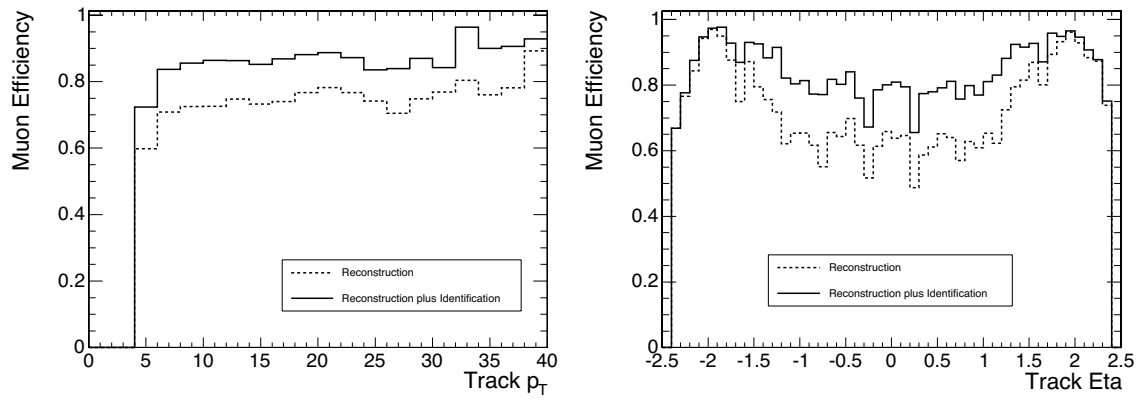

Figure 9.22: Efficiencies for global muon reconstruction (dashed lines) and global muon reconstruction plus muon identification (solid lines) for prompt muons contained within bottom quark jets, as a function of muon p_{T} (left) and pseudorapidity (right).

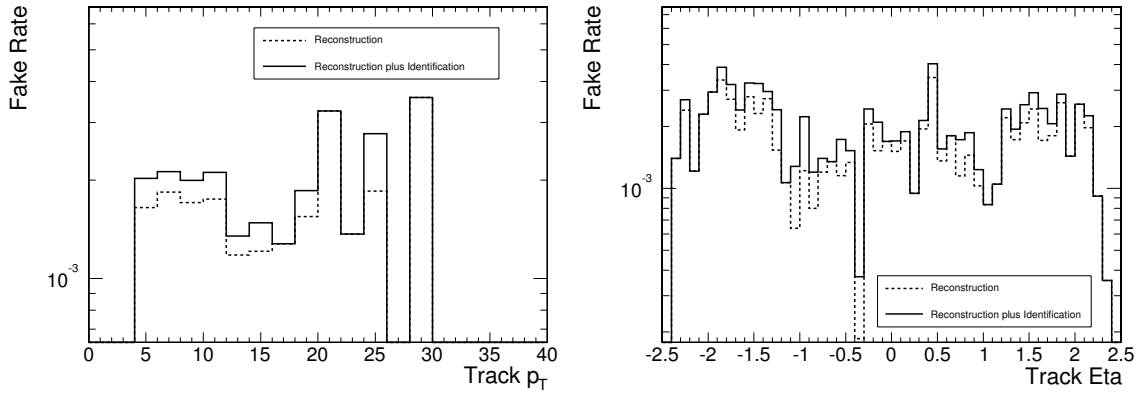

Figure 9.23: Fake rates for global muon reconstruction (dashed lines) and global muon reconstruction plus muon identification (solid lines) for pions contained within bottom quark jets, as a function of pion p_{T} (left) and pseudorapidity (right).

main sources of muons are from b and c decays. Another important contribution, mostly for low-p_{T} muons, is given by muonic K and π decays. All of these muons are produced in (usually soft) jets and are thus accompanied by nearby particles. Only for p_{T} above approximately $30\ \mathrm{GeV}/c$ do muons from W and Z decays become dominant. Muons from W, Z, and other heavy objects are accompanied only by particles from pile-up and by uncorrelated particles from the underlying event. Muon isolation is a tool to distinguish between the muons produced in jets and those coming from the decays of heavy objects.

The isolation algorithms that have been developed rely on the comparison of the total energy deposited in a cone around the muon with a predefined threshold. The deposit can be transverse energy in a calorimeter or the sum of transverse momenta of reconstructed charged-particle tracks. The cone axis is chosen according to the muon direction with a procedure that is tailored to the specific properties of each algorithm. The geometrical definition of the cone is given by the condition $\Delta R \leq \Delta R_{\mathrm{MAX}}$, where $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$, with $\Delta\eta$ and $\Delta\phi$ being the distances in pseudorapidity and azimuthal angle between the deposit and the cone axis, respectively. The muon itself contributes to the energy measurement inside the cone. This contribution (called the veto value) can be subtracted to improve the discriminating power of the isolation algorithm. A schematic illustration of the isolation cone is shown in Fig. 9.24.

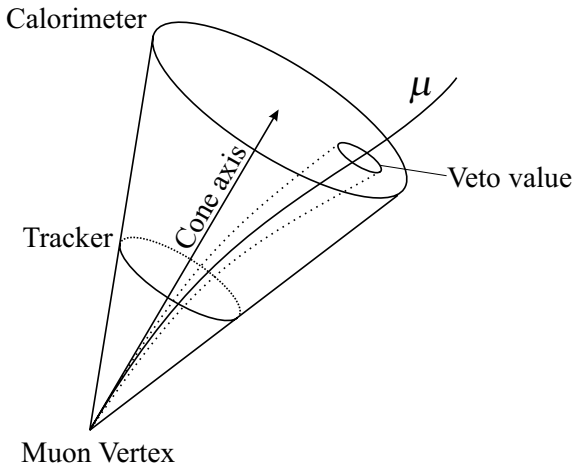

Figure 9.24: Schematic illustration of the isolation cone. The muon direction at the vertex defines the cone axis. The energy deposit ($\sum p_{\mathrm{T}}, \sum E_T^{\mathrm{weighted}}$) in the cone is computed, and the muon contribution is removed by excluding the small area around the muon (the "veto value") from the cone. Comparison of the deposit in the cone with a predefined threshold determines the muon isolation.

9.3.2 Online event selection

For the purposes of the High-Level Trigger, 3 isolation techniques were developed. Their detailed description can be found in [8, 238].

9.3.2.1 Calorimeter isolation

The calorimeter isolation is based on a standard technique of summing the calorimeter transverse energy in a cone around the muon direction. This method can be used with the stan-

dalone reconstruction at Level-2 and with fully-reconstructed Level-3 muons. The energy deposit in the cone is defined as the weighted sum of the transverse electromagnetic ($E_{\mathrm{T}}^{\mathrm{ECAL}}$) and hadron calorimeter energies ($E_{\mathrm{T}}^{\mathrm{HCAL}}$), using $E_{\mathrm{T}} = \alpha E_{\mathrm{T}}^{\mathrm{ECAL}} + E_{\mathrm{T}}^{\mathrm{HCAL}}$, with $\alpha = 1.5$, which reflects the better discrimination performance of the electromagnetic calorimeter. The veto cone is defined by the extrapolation of the muon to the ECAL/HCAL boundary (Fig. 9.25). Additional cuts against noise are used at the level of contributing crystals and towers. Typical energy-deposit distributions are shown in the left-most plots of Fig. 9.26.

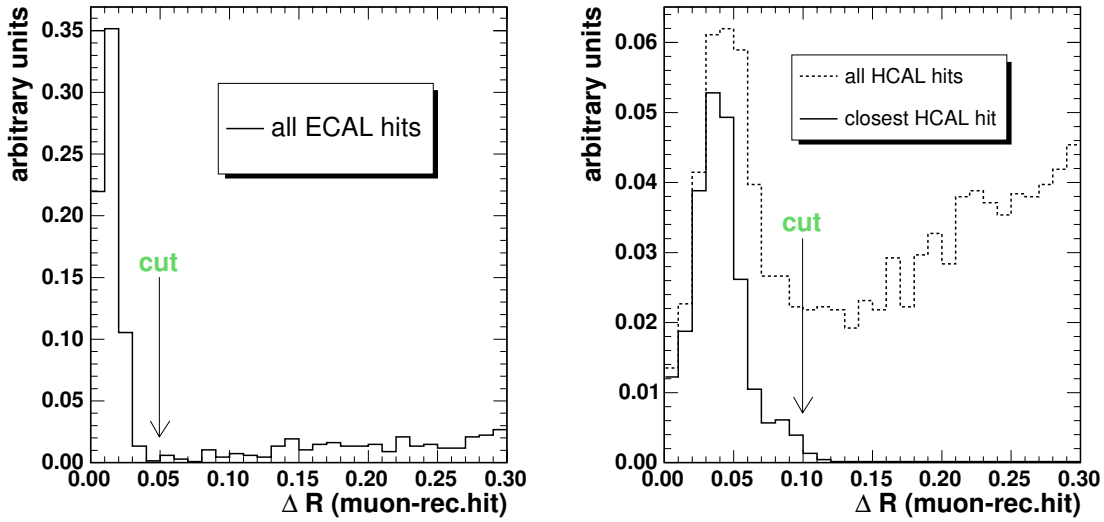

Figure 9.25: Distribution of the distance between a reconstructed muon and its energy deposit in the ECAL (left) and HCAL (right). The muon position is taken at the boundary between the ECAL and HCAL. The arrows indicate the size of the veto cone.

9.3.2.2 Pixel and full-tracker isolation

The pixel and full-tracker isolation algorithms are based on the scalar sum of p_{T} from all tracks reconstructed in a cone around the direction of the Level-3 muon at the vertex, neglecting the contribution from the muon itself. In both algorithms, the tracks are reconstructed using regional tracking.

In the case of the pixel isolation, the track candidates are searched for only in the pixel detector by dedicated pixel-reconstruction software. Such a reconstruction is based on 3-hit tracks. It is optimized for speed but, due to the small number of hits, has very poor momentum resolution and limited track-finding efficiency ($\approx 90\%$). For the full-tracker reconstruction, a Kalman-Filter fitting procedure is used at the price of longer CPU execution time. For both methods, the isolation cone defines the tracking region, with the vertex constraint coming from the isolated muon. Thus, the influence of tracks from pile-up is minimized. Since the direction resolution is good for both algorithms, the veto cone is very small ($\Delta R \approx 0.015$). A comparison of the energy-deposit distributions for the full-tracker and the pixel isolation is presented in the centre and right-most plots of Fig. 9.26, respectively.

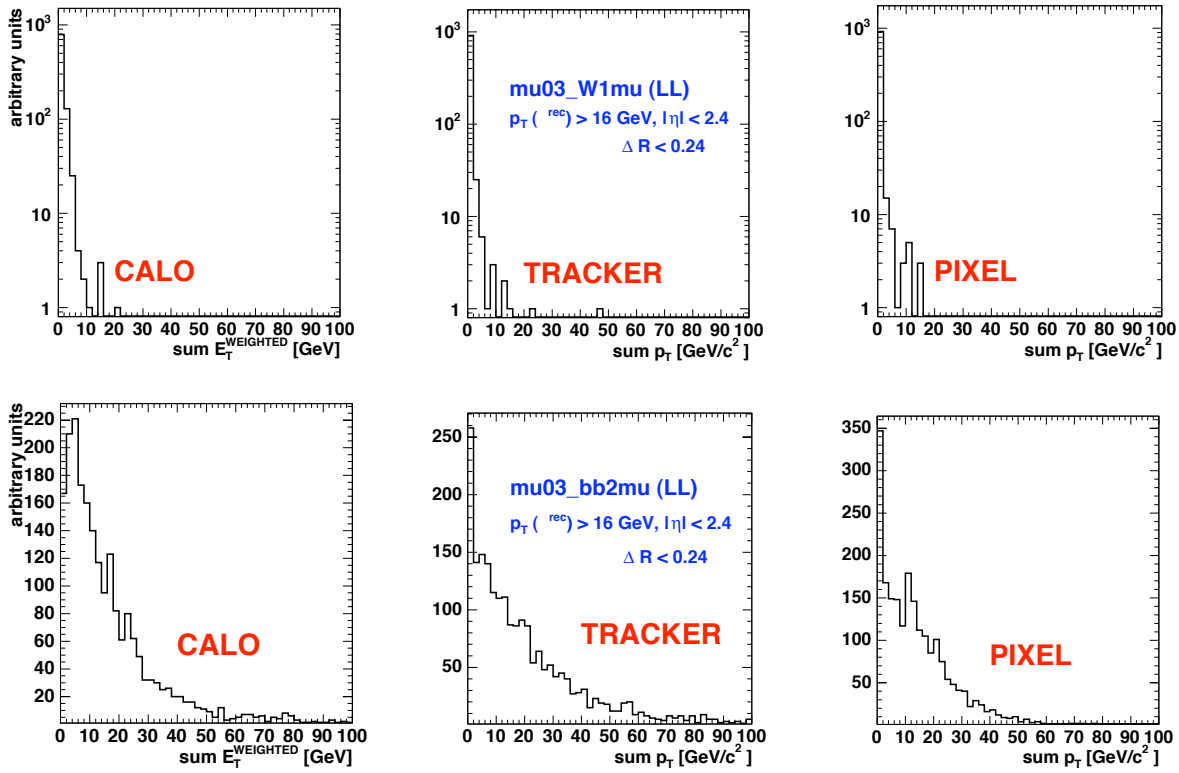

Figure 9.26: Distributions of the energy deposited in a $\Delta R = 0.24$ isolation cone for the calorimeter (left), the full tracker (centre), and the pixel tracker only (right). The difference in energy deposits between a typical source of isolated muons ($W \to \mu X$, top) and non-isolated muons ($b\bar{b} \to 2\mu X$, bottom) is well visible. The distributions are for the low-luminosity case.

9.3.2.3 Optimization and performance

The first step in the optimization of the isolation algorithm is the determination of the transverse energy and momentum thresholds as a function of cone size. Since some of the detector effects (noise, efficiency) may depend on pseudorapidity, the thresholds are evaluated in bins of η. The values of the thresholds are obtained using muons from $W \to \mu\nu$ decays (the "reference signal"), a typical source of isolated muons at the HLT level. For each combination of a cone size and an η bin, the threshold on the transverse energy or momentum in the cone is determined such that the efficiency of the reference signal is equal to some target efficiency (the "nominal efficiency"); however, for some choices of cone sizes, this is not always possible and the reference signal efficiency is somewhat larger than the chosen nominal efficiency.

The next step is the choice of a particular cone size (sizes considered are listed in Table 9.2). It depends on the background process that has to be rejected. For purposes of the HLT, $b\bar{b} \to \mu X$ events ("reference background") were used. The dependence of the background rejection on cone size is presented in Fig. 9.27 for several choices of the nominal efficiency. The efficiency of the HLT algorithms after the cone-size optimization is shown in Fig. 9.28 for the reference signal (left) and background (right) events, as a function of the nominal efficiency.

Table 9.2: The relationship between the cone-size indices and the ΔR values used in the optimization procedure.

Cone Index	ΔR_{MAX}	Cone Index	ΔR_{MAX}
1	0.02	8	0.28
2	0.045	9	0.32
3	0.09	10	0.38
4	0.13	11	0.45
5	0.17	12	0.5
6	0.2	13	0.6
7	0.24	14	0.7

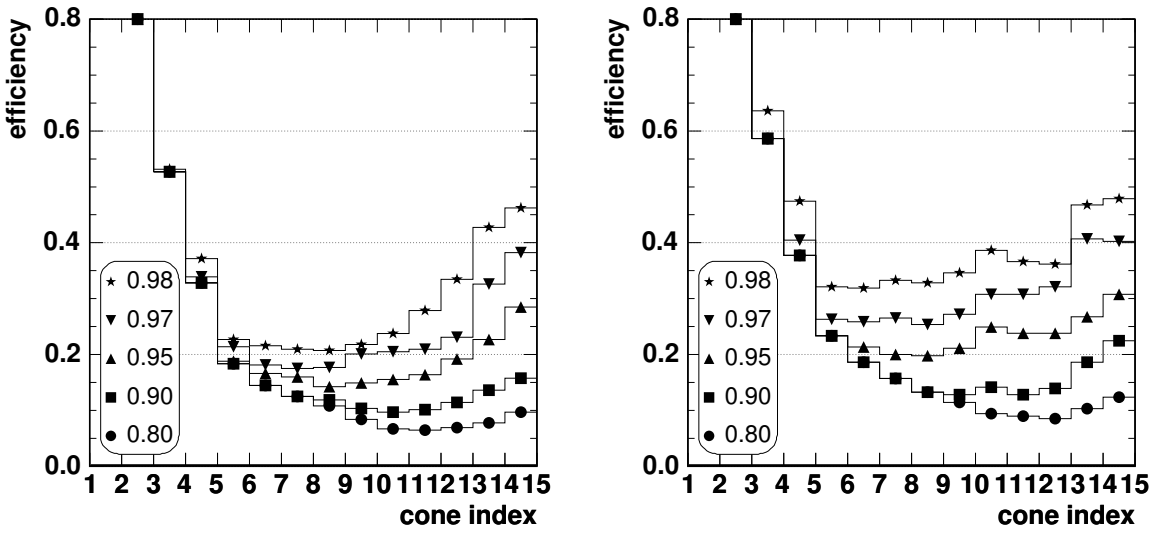

Figure 9.27: An example of an optimization of cone size: the efficiency for a background sample is shown as a function of the cone-size index for different predefined nominal efficiencies. The ΔR values corresponding to the different cone-size indices (horizontal axis) are given in Table 9.2. The left plot is for the full-tracker isolation in the low-luminosity case, the right plot is done with the pixel-based isolation for the high-luminosity environment.

9.3.3 Example of the isolation procedure applied offline

An offline algorithm should be optimized for a given signal and its most difficult background. For example, a calorimeter-based isolation requirement applied to high-p_{T} muons is bound to be inefficient because such muons are likely to radiate photons. Conversely, the isolation algorithm based on the distance between the muon and nearby jets does not efficiently reject background muons in soft jets that are likely to pass the HLT isolation. The aim of the study presented here is to show the possibility of further suppressing minimum-bias events that pass the HLT isolation chain by using additional offline isolation requirements. The HLT reference signal ($W \rightarrow \mu\nu X$) and background ($b\bar{b} \rightarrow \mu X$) processes are used, as before, for optimization.

The goal of the implemented offline algorithm is to combine precisely reconstructed charged particles measured in the tracker with neutral particles measured by the calorimeters. Since CPU processing time is less of an issue, the algorithms can benefit from global reconstruction in order to reduce pile-up contributions to the isolation cone.

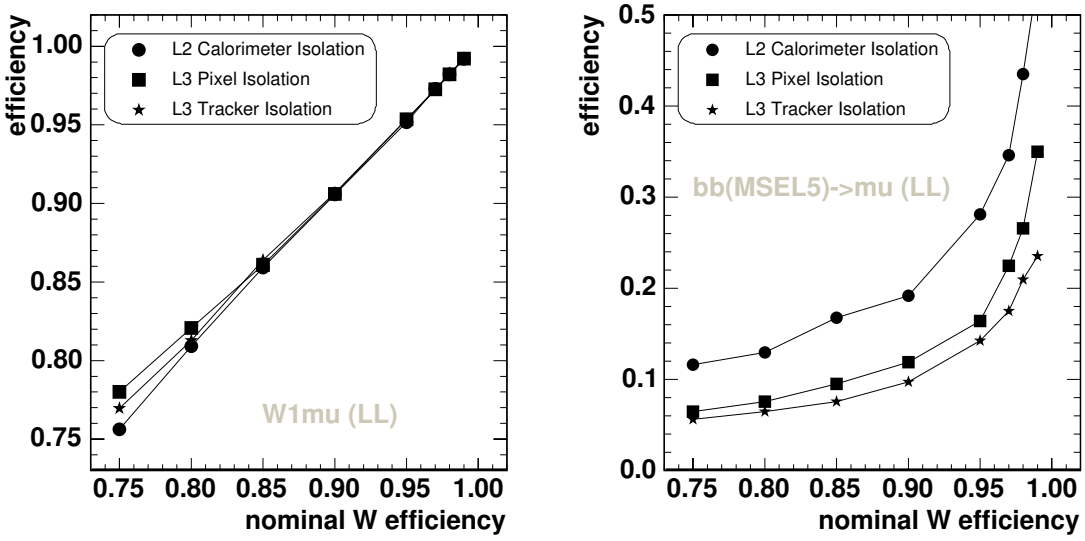

Figure 9.28: Comparison of the performance of 3 types of isolation algorithms at the HLT level. The isolation efficiency is plotted as a function of the nominal efficiency for signal events (left) and background events (right).

Particle trajectories from the global track reconstruction are propagated to the calorimeters. The muon energy deposit is identified and removed (by setting a veto cone) and is not used further in the computation. The deposit in the isolation cone is taken as the sum of the transverse momenta of reconstructed charged particles and the calorimeter transverse energy, both inside the isolation cone, from which the charged-particle deposit from the global reconstruction is subtracted. The distribution of deposits is shown in Fig. 9.29. The subtraction of the charged particle contribution from the calorimeter energy deposit is illustrated in Fig. 9.30.

The procedure to find the optimal threshold and cone size is the same as for the HLT. The results are shown in Fig. 9.31. The best background rejection that preserves 97% of the signal events (after already passing the HLT selection) is for a cone size of 0.28. A rejection factor of approximately 2 is obtained.

9.3.4 Sensitivity of the muon isolation cut efficiency to the underlying event uncertainties

The uncertainties in predicting the muon isolation cut efficiency were studied in the context of the $H \rightarrow ZZ \rightarrow 4\mu$ analysis by varying the PYTHIA parameters responsible for the simulation of the underlying event (UE) [239]. The isolation cut efficiency per muon can vary as much as $\pm 5\%$ due to uncertainties in the UE model (depending strongly on the energy deposit threshold), as shown in Fig. 9.32 (left). This results in a 4-muon isolation cut efficiency for the $ZZ \rightarrow 4\mu$ background of $(78 \pm 6)\%$. To eliminate these large uncertainties, the isolation cut efficiency can be calibrated from data using Z-inclusive events ($Z \rightarrow 2\mu$) and a random-cone technique, as displayed in Fig. 9.32 (right). This procedure eliminates the uncertainties associated with a poor understanding of the UE physics. There might be at most a $\approx 2\%$ systematic shift in the 4-muon isolation cut efficiencies using this method. This shift could be corrected for in principle, but it does not appear to be necessary since

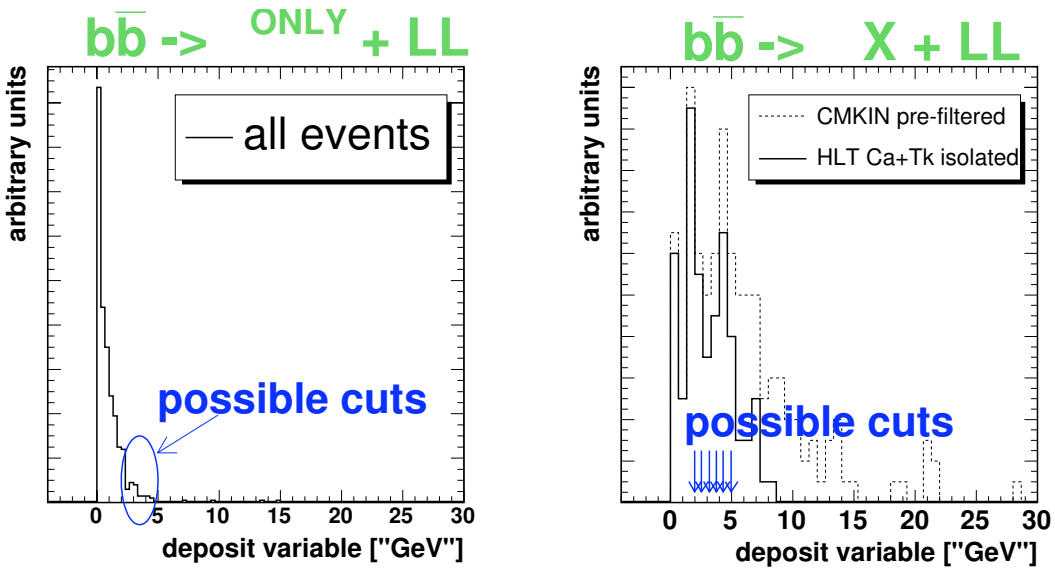

Figure 9.29: Energy deposits found offline in a $\Delta R = 0.24$ cone for isolated muons (left) and muons in $b\bar{b}$ jets (right). Only events passing the online isolation cut are used. Possible values for the energy deposit cut are marked on both plots.

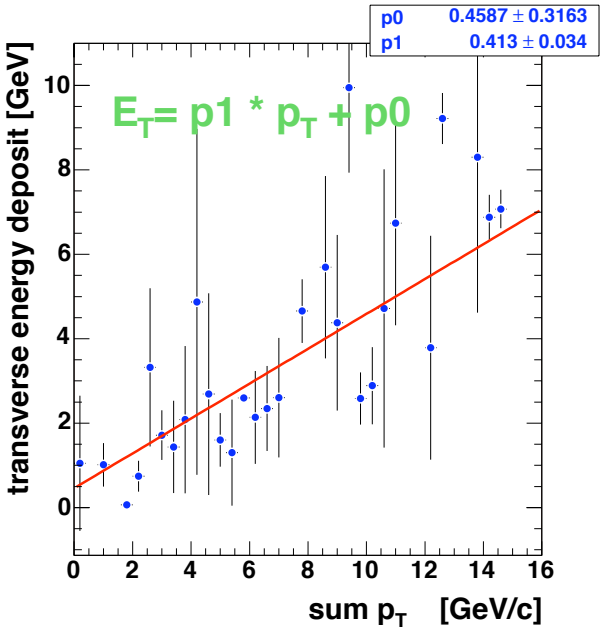

Figure 9.30: The correspondence between the measured transverse energy (vertical axis) and the total transverse momenta of all charged particles (horizontal axis) contributing to a deposit in the calorimeters. For the purpose of this figure, only charged particles were simulated. The cone size used was 0.24.

it is already small compared to the expected statistical uncertainties. These results and the analysis technique may be of interest to other analyses relying on lepton isolation cuts.

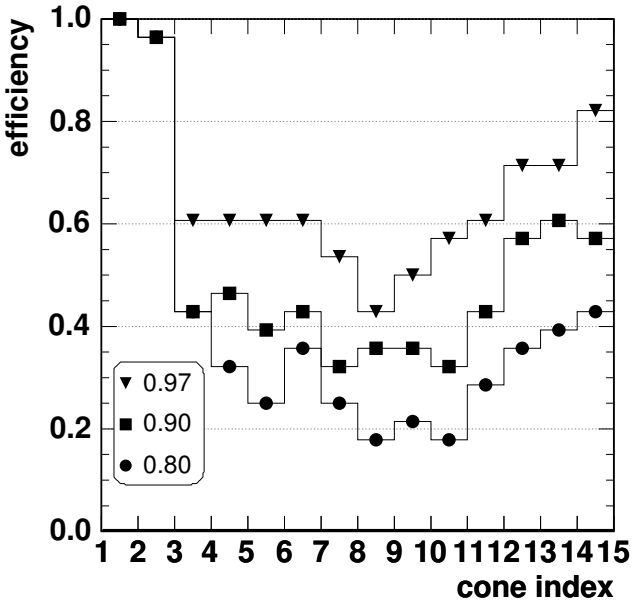

Figure 9.31: Performance of an offline isolation algorithm, in addition to the online selection. The efficiency for the background sample is shown as a function of the cone index for predefined nominal efficiencies in the low-luminosity case. The ΔR value for each cone index (horizontal axis) is given in Table 9.2.

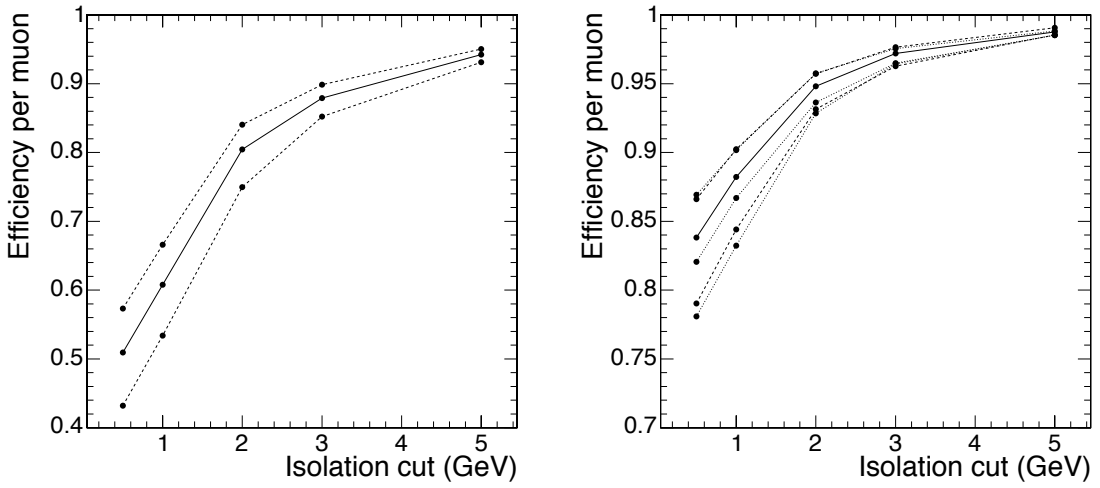

Figure 9.32: Left: the muon isolation cut efficiency for the least-isolated muon from $H \to 4\mu$ events. The middle solid line is for the default Multiple Interaction (MI) p_{T} "cut-off" value, the upper dashed line is for a -3σ variation in the p_{T} cut-off value, and the lower dashed line is for a $+3\sigma$ variation. Right: the muon isolation cut efficiency found using a random-cone technique in ZZ (solid and dashed lines) and Z-inclusive (dotted lines) events. The meaning of the middle, upper and lower lines is the same as for the left plot.

9.4 High-Level trigger

A large fraction of the expected LHC physics is characterized by the presence of high-p_{T} leptons, and CMS has been designed and optimized, in particular, to reconstruct muons. Muons provide a very clean signature; very high efficiency can be reached with inclusive-muon triggers based on the selection of 1 or more muons above a given p_{T} threshold. While this threshold should be low enough not to lose efficiency, it must be high enough to reduce the trigger rate to an acceptable level.

The online selection of muons within the High-Level Trigger (HLT) system is performed in 2 steps. First, to confirm the Level-1 decision, a Level-2 muon reconstruction is performed, i.e., muons are reconstructed using only information from the muon system and with the Level-1 candidates as a seed. The details of this reconstruction algorithm are described in Section 9.1.1. Secondly, the Level-2 p_{T} threshold reduces the rate enough to then allow a Level-3 reconstruction using the full tracker information and seeded by the Level-2 candidates. The Level-3 muon reconstruction algorithm is described in Section 9.1.2. Additional rate reductions can be achieved, if needed, by requiring muons to be isolated (Section 9.3). Already after the Level-2 selection, the calorimeter isolation can be required (Section 9.3.2.1), and the tracker and pixel isolation can be required after the Level-3 selection.

After each reconstruction step, a set of selection cuts is applied. These selection cuts are different for single and dimuon triggers. The details of the selection cuts are described in [8]. The p_{T} thresholds are set to achieve a sufficient rate reduction. As discussed in [8], for a luminosity of $2\times10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ an inclusive single-muon trigger rate of 25 Hz can be achieved with a p_{T} threshold of $19\ \mathrm{GeV}/c$ (including isolation). For dimuon triggers, a typical p_{T} threshold is $7\ \mathrm{GeV}/c$, resulting in a rate of 4 Hz.

9.4.1 Performance

A detailed performance study of single and dimuon high-level triggers can be found in [8]. However, since the time this study was performed, the CMS simulation and reconstruction software has been significantly changed. Hence, some of the findings presented in [8] were cross-checked with the results obtained using the most recent version of the CMS simulation and reconstruction software. The overall efficiency for muons to pass the Level-1 through Level-3 single-muon trigger criteria cumulatively as a function of η is shown in Fig. 9.33. Muons were generated without pile-up and with flat distributions in the intervals $5 < p_{\mathrm{T}} < 100\ \mathrm{GeV}/c$ and $|\eta| < 2.1$. Figure 9.34 shows the single-muon HLT rates without applying muons isolation cuts for a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. A sample of 10^6 minimum-bias events was used to calculate these rates. The single-muon rate at Level-2 is still dominated by feed-through of low-p_{T} muons. The improved Level-3 momentum resolution allows these muons to be discarded. Despite the fact that completely different event samples and a different version of the CMS simulation and reconstruction software were used, the result are in agreement with those in [8].

9.4.2 High-Level trigger for muons in heavy-ion collisions

The High-Level Trigger will be used to select dimuons from quarkonia produced in heavy-ion collisions. The J/ψs and Υs produced with low p_{T} are expected to provide information on the state of the hot nuclear matter produced in these collision. The relatively low luminosity of $L = 10^{27}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ allows any low quality muon in the CSC trigger system (a muon

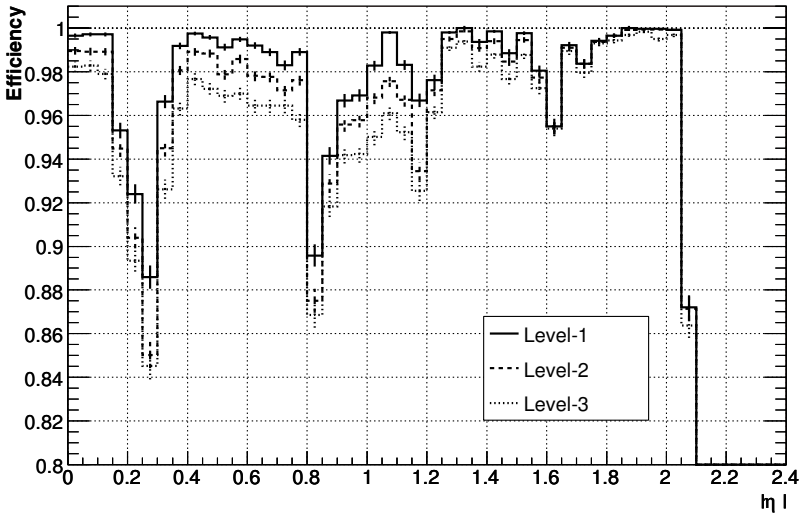

Figure 9.33: Cumulative efficiency for single muons to pass Level-1 (solid), Level-2 (dashed), and Level-3 (dot-dashed) trigger levels as a function of the generated muon pseudorapidity. No thresholds on p_{T} are applied. The vertical axis is zero-suppressed.

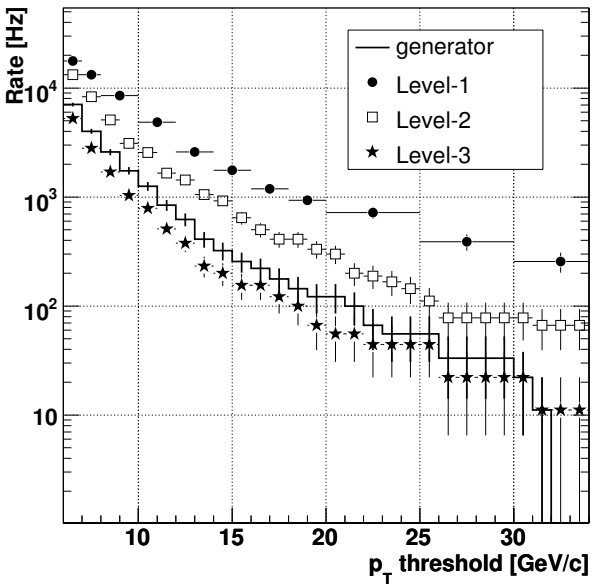

Figure 9.34: HLT single-muon rates after Level-1, Level-2, and Level-3 triggers, as a function of the p_{T} threshold for a luminosity of $2 \times 10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The rates are obtained without applying isolation cuts.

without a segment in ME1, as described in Section 3.4.3) to be sufficient to generate a Level-1 accept. Such muons passing the looser requirements at Level-1 are denoted hereafter "Open Level-1" muons (OL1), and are considered in the heavy-ion HLT trigger decision. The final trigger decision is delayed until the detailed reconstruction is performed at Level-2 and Level-3 of the HLT.

Different configurations of the Level-1 and Level-2 trigger flags were considered to maximize the selection efficiency for J/ψ and Υ. No cuts on any kinematic variables (p_{T}, η, ϕ) and no isolation cuts in the calorimeter are applied, apart from those caused by the detector acceptance.

The information obtained from Level-1 and Level-2 processing is used to estimate the muon momentum and its direction based on muon chamber information only. The Level-2 processing is applied on the output from Level-1 (including both "high quality" Level-1 muons and OL1 muons). If Level-2 fails, the information from Level-1 is used to estimate the muon momentum. For the case when Level-2 fails in the *barrel* and the momentum of the OL1 muon is undetermined, then the p_{T} of the OL1 muon is set to $3.5\ \mathrm{GeV}/c$. If Level-2 fails in the *endcap* and the momentum of the OL1 muon is undetermined, then the p_{T} of the OL1 is fixed to $1\ \mathrm{GeV}/c$. In both cases, the momentum components are recalculated according to the Level-1 muon direction.

The chosen strategy requires either 2 opposite-sign Level-1 or 2 opposite-sign Level-2 muon candidates.

The use of OL1 muons as the Level-2 input favors the survival of low-p_{T} J/ψs; the selection efficiency is $\approx 1\%$, as compared to 0.41% with the more stringent Level-1 usage without such low quality muons. For Υs, the acceptances are 21% for the OL1/Level-2 combination and 16% for the more stringent Level-1/Level-2 combination. The η and p_{T} distributions for accepted J/ψs and Υs are shown in Figs. 9.35, 9.36, 9.37 and 9.38. The results obtained using OL1 as input to Level-2 are shown with a solid-line histogram and those using Level-1 as input to Level-2 are shown with the shaded histogram.

While the trigger criteria are optimized for muons from decaying quarkonia, they will also be used to select same-sign dimuon pairs needed to estimate the uncorrelated background [240].

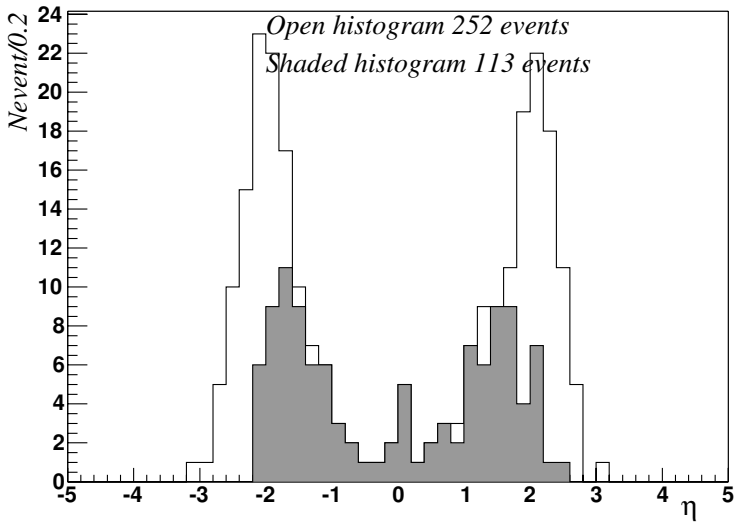

Figure 9.35: The η distribution of J/ψs with both muons triggered by the "Open Level-1" and Level-2 trigger combination (solid-line histogram) as compared with the more stringent Level-1 and Level-2 trigger (shaded histogram). Two opposite-sign Level-1 or 2 opposite-sign Level-2 candidates are required. See text for details.

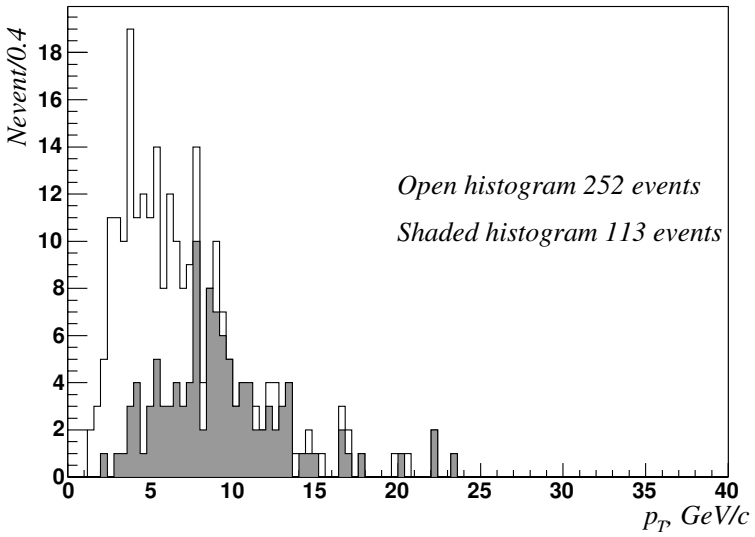

Figure 9.36: The p_{T} distribution of J/ψs with both muons triggered by the Open Level-1 and Level-2 trigger (solid-line histogram) as compared with the more stringent Level-1 and Level-2 trigger (shaded histogram). Two opposite-sign Level-1 or 2 opposite-sign Level-2 candidates are required.

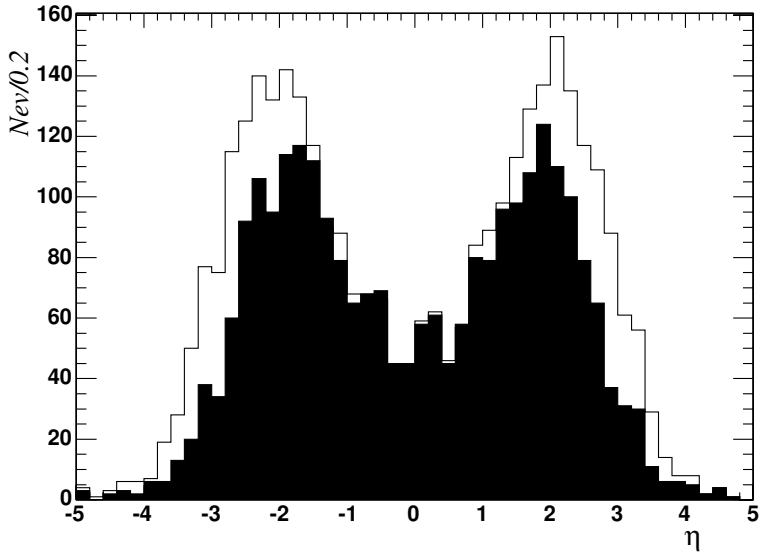

Figure 9.37: The η distribution of Υs with both muons triggered by the Open Level-1 and Level-2 trigger (solid-line histogram) as compared with the Level-1 and Level-2 trigger (shaded histogram). Two opposite-sign Level-1 or 2 opposite-sign Level-2 candidates are required.

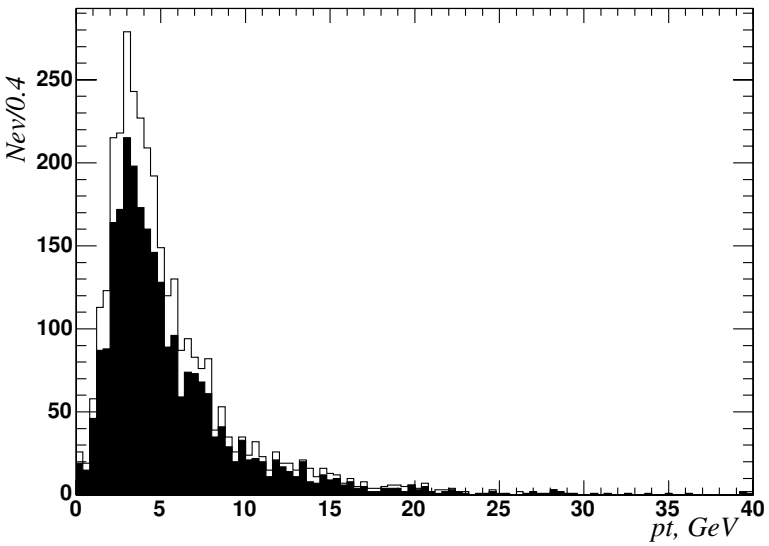

Figure 9.38: The p_{T} distribution of Υs with both muons triggered by the Open Level-1 and Level-2 trigger (solid-line histogram) as compared with the Level-1 and Level-2 trigger (shaded histogram). Two opposite-sign Level-1 or 2 opposite-sign Level-2 candidates are required.