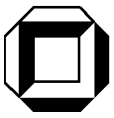

Universität Karlsruhe (TH)
Forschungsuniversität • gegründet 1825

VIRTUALISATION OF GRID RESOURCES AND PROSPECTS OF THE MEASUREMENT OF Z BOSON PRODUCTION IN ASSOCIATION WITH JETS AT THE LHC

Zur Erlangung des akademischen Grades eines
DOKTORS DER NATURWISSENSCHAFTEN
von der Fakultät für Physik der
Universität Karlsruhe (TH) genehmigte

DISSERTATION

von

Dipl. Phys. Volker Martin Büge
aus Karlsruhe

Tag der mündlichen Prüfung: 18. Juli 2008

Referent: Prof. Dr. G. Quast
Institut für Experimentelle Kernphysik

Korreferent: Prof. Dr. M. Feindt
Institut für Experimentelle Kernphysik

meiner Familie

Virtualisation of Grid Resources and Prospects of the Measurement of Z Boson Production in Association with Jets at the LHC

At the Large Hadron Collider (LHC), a large number of events containing Z bosons will be available. These events will be used for precise electroweak measurements but also for the commissioning of different detector components. Furthermore, the calibration of the absolute jet energy scale using Z boson events becomes feasible for the first time. As the uncertainty on this observable will be the dominant systematical uncertainty on most LHC analyses, its precise calibration is indispensable. For such purpose, events in which the Z boson decays into two muons are well suited as the kinematics of the boson can be reconstructed very precisely and without relying on calorimetric information. The absolute jet energy scale can be calibrated through a comparison of the measured momentum of the jet with the momentum of the Z boson. But this equivalence only holds in events in which the Z boson momentum is balanced by exactly one parton of the hard process. In this thesis, the feasibility of such a calibration within the CMS experiment is demonstrated. Taking purely statistical uncertainties into account, the determination of a calibration from Z boson balancing is possible up to transverse momenta of 350 GeV, even for an expected number of events corresponding to an integrated luminosity of only $100\,\mathrm{pb}^{-1}$. The Z plus jet based calibration of the absolute jet energy scale is deduced for all algorithms currently used in CMS, including the investigation of effects from the underlying event, the choice of the jet size parameter and the event selection.

The four LHC experiments will produce about 15 petabytes of data annually. The safe storage and processing of this enormous amount of data are quite challenging and the experiments cope with this using grid technologies. This, however, requires additional services which do not fully load server machines at smaller grid sites. The virtualisation of operating systems enables the consolidation of different services on a few server machines. After a brief introduction to the concept of operating system virtualisation, different techniques are discussed with respect to their suitability for grid service consolidation followed by a presentation of the final consolidation of the services of the institute's computing and grid infrastructure. Furthermore, the concept of server consolidation is extended to the field of batch queuing systems to overcome intrinsic limitations of shared computing infrastructures. Based on virtualisation techniques, multiple operating systems can be offered to different user groups without losing the advantage of load balancing. Experiences with a prototype implementation of such functionality is presented and discussed in detail. Applying virtualisation techniques helps to increase the load, stability and maintainability of the local computing infrastructure.

Contents

Introduction

The aim of physics is to obtain a detailed and consistent description of the basic laws of nature. For this, theoretical models are deduced from experimental observations and in return, the predictions of these models have to be tested by future experiments to confirm or falsify the assumptions. The Standard Model of particle physics is the theory which has been developed and successfully been tested in numerous experiments over the last decades. It describes the fundamental particles and their interactions to an unprecedented precision. Nevertheless some fundamental questions like the origin of mass remain unsolved.

One of the major successes was the discovery of the mediators of the weak interaction, the charged W bosons and the neutral Z boson with the Super Proton Antiproton Synchrotron at CERN[1] in 1983. Up to now, the properties of these particles have been studied in various experiments. Particularly, the experiments at the Large Electron Positron Collider have measured the properties of the Z boson with highest accuracy.

However, in current and planned experiments the focus on the Z boson changes and besides precise electroweak measurements, it becomes a promising candidate for various purposes. The large number of events containing Z bosons can be used for a precise monitoring of the luminosity of the collider or for the commissioning and calibration of several detector components.

At the Large Hadron Collider (LHC) which will start operation this year, processes of quantum chromodynamics have by far the largest cross-section. The large number of such events is predestined for early analyses leading to a precise understanding of the fundamental processes of the Standard Model at the LHC energy scale and therewith of the backgrounds for most physics analyses. In addition, they are well suited for the commissioning of different detector components.

Since coloured partons of the hard process hadronise, only a collimated stream of colour-neutral objects can be observed with detectors. The concept of jets is introduced to infer on the properties

[1] Conseil Européen pour la Recherche Nucléaire

like momentum and energy of the parton. A jet is defined as the cluster of all particles which are supposed to originate from the same initiator. However, the reconstructed energy of the jet in the detector does not correspond exactly to the energy of the initial parton due to various effects and a precise calibration of its energy is indispensable as the uncertainty on the jet energy scale will dominate the systematical uncertainty for almost all analyses at the LHC.

Several methods have been deployed at the different experiments to calibrate jets. With the large number of Z bosons which will be produced at the LHC, the precise calibration of the absolute jets energy scale with these events becomes feasible for the first time and is investigated in detail in this thesis within the CMS[2] experiment.

In chapter one, a brief introduction to the Standard Model of particle physics is given including the basic concepts of local gauge invariance which leads to an interacting field theory and the electroweak unification. As the composite structure of the proton is important for hadron colliders, its measurement is also discussed in this chapter which closes with the basic concepts of the strong interaction and the production of Z bosons. This theoretical introduction is followed by a general overview of the CMS detector in chapter two.

At the LHC, data management and the access to computing resources are realised using grid technologies. In chapter three, the LHC computing concept is presented focussing on the requirements of the CMS experiment. This includes an introduction to the basic concept of grid computing with a discussion of the required services, their usage, the administration and the particularities of the computing infrastructure of the Institut für Experimentelle Kernphysik (IEKP) with respect to its connection to the grid.

However, such an environment of distributed storage and computing power is challenging and will benefit from optimisations. To share temporarily unused resources with the collaboration, several grid services have to be deployed at each participating site which requires additional hardware. Especially at the large number of smaller grid sites like the IEKP, these services do not fully load the server machines. This hardware overhead can be reduced using virtualisation techniques. In addition, the concept of virtualisation can be extended to the field of batch queuing systems to combine the support of multiple operating systems adapted for the usage of different user groups with the advantages of opportunistic use. Both approaches are presented in chapter four including a brief introduction to the concept of virtualisation.

The CMS software framework is introduced in chapter five which comprises a description of the full simulation and reconstruction chain focussing on the reconstruction of muons and jets. In this context the definition of the jet clustering algorithms which are applied in the analysis is presented.

The jet energy correction procedure envisaged by CMS is outlined in chapter six concentrating on the calibration of the absolute jet energy scale using Z boson events. For this, the required event topology and the applied selection criteria are discussed. In a first step, the self-consistency

[2] Compact Muon Solenoid

of the method is demonstrated using Monte Carlo simulations. This is followed by an evaluation with Z boson events of existing calibrations in CMS. After a discussion of the influence of the underlying event and the jet size parameter on this method, the deduction of a calibration of the absolute jet energy scale with Z boson events is presented.

The Standard Model of Particle Physics

1.1 Introduction to the Standard Model

The Standard Model of particle physics, formulated in the 1970s, is a theory which describes the fundamental particles and their interactions. According to this model, all matter is built from a small number of fundamental spin 1/2 particles, namely six leptons and six quarks. For each of these particles, which are called fermions according to their half-integer spin, an antiparticle exists which has the same properties as the corresponding particle but the signs of its quantum numbers which are listed in the tables below are reversed.

The four fundamental interactions between particles are described by the exchange of integer spin mediators which are called bosons: the photon for the electromagnetic force, two W bosons and the Z boson for the weak interaction as well as eight gluons for the strong interaction. Gravity takes a special position in this context as it is not included in the Standard Model at present and its predicted mediator, the graviton, has not been observed to date.

Half of the six leptons carry a charge of $-e$ and can be paired with a neutral lepton, the neutrinos, to form three families consisting of the electron, the muon and the tau with their corresponding neutrino. Characteristic for each family is a quantum number called the electron, muon and tau number, which is conserved by all interactions. However, the mass eigenstates of the neutrinos differ from their energy eigenstates leading to experimentally observed oscillations between different flavours. Obviously, these oscillations do not conserve the family specific lepton numbers but only their sum. An overview of the three lepton families is given in Table 1.1.

The six quarks carry a fraction of 2/3 or -1/3 of the elementary charge and can also be grouped into three families by requiring a combined absolute charge of one and an ordering by mass. Each quark flavour has an own quantum number which is conserved by all interactions except the weak force. This violation is a result of the difference between the mass eigenstates of the quarks and

	Particle	Q	L_e	L_μ	L_τ
First	e	-1	1	0	0
generation	ν_e	0	1	0	0
Second	μ	-1	0	1	0
generation	ν_μ	0	0	1	0
Third	τ	-1	0	0	1
generation	ν_τ	0	0	0	1

Table 1.1: *The three lepton families with the corresponding quantum numbers of charge (Q), electron number (L_e), muon number (L_μ) and tau number (L_τ).*

the eigenstates of the weak interaction. The two representations are connected by the Cabibbo-Kobayashi-Maskawa (CKM) matrix and flavour changes without conservation of the dedicated quantum number become possible. A detailed overview of the different quark flavours is presented in Table 1.2.

With the Δ^{++} resonance, a spin 3/2 particle consisting of three up quarks has been observed. As this particle is assumed to have no angular momentum, the three quarks are in the same state and an additional quantum number is required to preserve the Pauli principle. This quantum number is called the colour-charge and can adopt the three values of red, green and blue. All particles observed to date are colour-neutral, which indicates that quarks do not exist as free particles. According to chromatics, the colour disappears if a colour and its anti-colour are combined. This is possible for a bound state of a quark and an anti-quark, which is called meson. Baryons are bound states of three quarks where each colour is present exactly once. These two strongly interacting bound quark states are called hadrons.

	Particle	Q	D	U	S	C	B	T
First	d	$-\frac{1}{3}$	-1	0	0	0	0	0
generation	u	$\frac{2}{3}$	0	1	0	0	0	0
Second	s	$-\frac{1}{3}$	0	0	-1	0	0	0
generation	c	$\frac{2}{3}$	0	0	0	1	0	0
Third	b	$-\frac{1}{3}$	0	0	0	0	-1	0
generation	t	$\frac{2}{3}$	0	0	0	0	0	1

Table 1.2: *Overview of the three quark families with the corresponding quantum numbers D, U, S, C, B and T. Each of these quantum numbers is characteristic for one quark flavour. For example, up-quarks carry the quantum number U, down-quarks the quantum number D and so on. Q represents the charge and additional quantum numbers like the spin are not listed.*

The mathematical concept of the Standard Model is based on quantum field theory, which will be introduced briefly. This includes a discussion of local gauge invariance by the example of the electromagnetic interaction with its mediator, the photon as well as the principles of electroweak unification. This is followed by an introduction to quantum chromodynamics which is the theory of the strong interaction.

1.2 Lagrangians in Relativistic Field Theories

In classical mechanics, the position of an object as a function of time can be calculated either by applying Newton's second law of motion or following the alternative formulation which starts with the Lagrangian

$$L(q_i, \dot{q}_i, t) = T - U \tag{1.1}$$

where the T denotes the kinetic energy of a particle in a scalar potential U. This Lagrangian is a function of the coordinates q_i and their time derivatives $\dot{q}_i$, for example the three Cartesian coordinates and the corresponding velocities. The equation of motion can be obtained from the Euler-Lagrange equation

$$\frac{d}{dt}\left(\frac{\partial L}{\partial \dot{q}_i}\right) = \frac{\partial L}{\partial q_i} \qquad (i = 1, 2, 3). \tag{1.2}$$

However, this formalism is so far only valid for localized, discrete particles and not for fields, occupying some region in space. Instead of dealing with the particle's position as a function of time, the concern of relativistic field theories is to calculate one or more functions ϕ_i of the space-time coordinates x_μ. Therefore, the Lagrangian L of classical mechanics has to be replaced by the Lagrangian density $\mathcal{L}$, which is a function of the fields ϕ_i and their derivatives $\partial_\mu \phi_i$:

$$L(q_i, \dot{q}_i, t) \rightarrow \mathcal{L}(\phi_i, \partial_\mu \phi_i, x_\mu) \quad \text{with} \quad \partial_\mu \phi_i \equiv \frac{\partial \phi_i}{\partial x^\mu}. \tag{1.3}$$

The left side of the classical Euler-Lagrange Equation 1.2 involves only time derivatives. However, a relativistic theory does not distinguish between time and space coordinates and the Euler-Lagrange equations generalise to

$$\partial_\mu \left(\frac{\partial \mathcal{L}}{\partial(\partial_\mu \phi_i)}\right) = \frac{\partial \mathcal{L}}{\partial \phi_i} \qquad (i = 1, 2, 3, \dots). \tag{1.4}$$

It has to be emphasized that in contrast to the Lagrangian of classical mechanics which could be derived analytically, the Lagrange densities in relativistic field theories are usually taken as axiomatic. They are constructed in such a way that they reproduce the desired field equations.

In the following, the three Lagrangians for spin-0, spin-1/2 and spin-1 particles are introduced [1] which, however, describe only free fields, without sources or interactions[1]:

- **Klein-Gordon Lagrangian** for a scalar (spin-0) field ψ:

$$\mathcal{L} = \frac{1}{2}(\partial_\mu \psi)(\partial^\mu \psi) - \frac{1}{2}m^2\psi^2 \tag{1.5}$$

[1] From here on, the convention $c = \hbar = 1$ is used.

which reproduces the Klein-Gordon equation

$$\partial^\mu \partial_\mu \psi - m^2 \psi = 0 \tag{1.6}$$

for a particle with spin 0 and a mass m.

- **Dirac Lagrangian** for a spinor (spin-1/2) field ψ:

$$\mathcal{L} = i\overline{\psi}\gamma^\mu \partial_\mu \psi - m\overline{\psi}\psi \tag{1.7}$$

Each of the four components of the complex spinor ψ and its adjoint are treated as independent field variables. Applying Equation 1.4 to $\overline{\psi}$ leads to

$$\left(i\gamma^\mu \partial_\mu - m\right)\psi = 0 \tag{1.8}$$

which is the Dirac equation describing a particle with spin 1/2 and mass m.

- **Proca Lagrangian** for a vector (spin-1) field A:

$$\mathcal{L} = -\frac{1}{4}(\partial^\nu A^\mu - \partial^\mu A^\nu)(\partial_\nu A_\mu - \partial_\mu A_\nu) + \frac{1}{2}m^2 A^\nu A_\nu. \tag{1.9}$$

Since the expression $(\partial^\nu A^\mu - \partial^\mu A^\nu)$ is used repeatedly in this context the following abbreviation is introduced:

$$F^{\mu\nu} = -F^{\nu\mu} = \partial^\nu A^\mu - \partial^\mu A^\nu. \tag{1.10}$$

Now, the Lagrangian can be expressed more clearly as

$$\mathcal{L} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + \frac{1}{2}m^2 A^\nu A_\nu. \tag{1.11}$$

and the corresponding field equation for a field with mass m becomes

$$\partial_\mu F^{\mu\nu} + m^2 A^\nu = 0. \tag{1.12}$$

1.3 Local Gauge Invariance

The principle of gauge invariance is well known from classical electrodynamics. The observables in this theory, the electric field $\vec{E}$ and the magnetic field $\vec{B}$ can be related to potentials according to Maxwell's equations

$$\vec{E} = -\nabla\phi - \frac{\partial \vec{A}}{\partial t} \quad \text{and} \quad \vec{B} = \nabla \times \vec{A}. \tag{1.13}$$

These potentials form the four-vector $A^\mu = (\phi, \vec{A})$ which is related to the electromagnetic field-strength tensor through Equation 1.10. A transformation of the vector field

$$A^\mu \rightarrow A'^\mu - \partial^\mu \chi \tag{1.14}$$

where χ is an arbitrary scalar function is called global gauge transformation. Such a transformation leaves the observables and thus the field strength tensor invariant. Different potentials can therefore describe the same physical quantities.

The same argumentation is valid for the Lagrangian of a free Dirac particle as defined in Equation 1.7 which is invariant under the transformation

$$\psi \rightarrow \psi' = e^{iq\chi}\psi, \tag{1.15}$$

as the additional exponential factors cancel out in the combination of the spinors $\overline{\psi}\psi$. But this changes for local gauge transformations where the phase factor is a function of x_μ. Now, the spinor transforms like

$$\begin{aligned} \psi(x) &\rightarrow e^{i\alpha(x)}\psi(x) \ , \\ \overline{\psi}(x) &\rightarrow e^{-i\alpha(x)}\overline{\psi}(x) \ . \end{aligned}$$

Such a phase transformation does not affect the second term in the Dirac Lagrangian, see Equation 1.7. The first term, however, contains a derivative which transforms like

$$\partial_\mu \psi \rightarrow e^{i\alpha(x)}\partial_\mu \psi + ie^{i\alpha(x)}\psi\partial_\mu \alpha(x) \tag{1.16}$$

and local gauge invariance is broken by the additional term $\partial_\mu \alpha(x)$.

To achieve local phase invariance nevertheless, a modified derivative D_μ fulfilling the condition

$$D_\mu \psi \rightarrow e^{i\alpha(x)} D_\mu \psi \tag{1.17}$$

is required. The introduction of an additional vector field A_μ which transforms like

$$A_\mu \rightarrow A_\mu - \frac{1}{q}\partial_\mu \alpha \tag{1.18}$$

leads to such a derivative cancelling out the additional term in Equation 1.16 which breaks the gauge invariance. With this new vector field, the covariant derivative which keeps the Lagrangian invariant can be derived to

$$D_\mu \equiv \partial_\mu + iqA_\mu. \tag{1.19}$$

By requiring local gauge invariance, a new local vector field, the gauge field, has to be introduced. This vector field can be regarded as the photon field and an additional term corresponding to its

energy has to be added. As the photon is a spin-1 particle, the Proca-Lagrangian (Equation 1.11) is a reasonable choice. However, the mass term $\frac{1}{2}\mathrm{m}^2 A^\nu A_\nu$ of this Lagrangian would spoil local gauge invariance and the photon field is therefore required to be massless.

This leads to the final Lagrangian of Quantum Electrodynamics (QED):

$$\mathcal{L}_{QED} = \overline{\psi}\left(i\gamma^\mu\partial_\mu - m\right)\psi - q\overline{\psi}\gamma^\mu\psi A_\mu - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}\,. \tag{1.20}$$

In the previous example, local gauge invariance has been investigated for the group of the transformation

$$\psi \to U\psi, \quad \text{where} \quad UU^\dagger = 1. \tag{1.21}$$

Here, $U = e^{iq\chi}$ is a scalar and the group of all these 1×1 matrices is called $U(1)$. Insisting, that a global gauge invariance holds locally for this group leads to the Lagrangian for the QED with a massless vector field, the photon field.

The same strategy was extended to the $SU(2)$ group[2] by Yang and Mills. For this group, the 2×2 unitary transformation matrix can be expressed as

$$U = e^{\frac{i}{2}\vec{\tau}\cdot\vec{a}} \tag{1.22}$$

where the three components of $\vec{\tau}$ are the Pauli matrices

$$\tau_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \tau_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \quad \text{and} \quad \tau_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{1.23}$$

To get a Lagrangian which is again invariant under local gauge transformations, three new vector fields W_μ have to be introduced, each requiring its own free Lagrangian. Again, the vector fields are required to be massless to hold local gauge invariance and the covariant derivative for this transformation is

$$D_\mu \equiv \partial_\mu + i\frac{g}{2}\vec{\tau}\vec{W}_\mu. \tag{1.24}$$

Historically, the Yang-Mills theory was designed to be applied to the strong interaction which lead to difficulties. However, it is very useful in the context of the weak interaction which is discussed in more detail later.

To derive the Lagragian describing Quantum Chromodynamics (QCD), the principle of local gauge invariance is applied to the colour group $SU(3)$, a subgroup of $U(3)$, [3]. The transformation of this group can be expressed as

$$U = e^{\frac{i}{2}\sum\alpha_j\lambda_j} \quad \alpha_i \in \mathbb{R} \tag{1.25}$$

[2] Any element of the $U(2)$ group can be expressed as an element of the $SU(2)$ group times a phase factor which is an element of the $U(1)$ group. This is denoted shortly as $U(2) = SU(2) \otimes U(1)$.

[3] $U(3) = SU(3) \otimes U(1)$.

with the 3×3 matrices λ_i. Requiring U to be hermetic and to have a determinant of 1 leads to exactly eight linearly independent and hermetic matrices, the Gell-Mann matrices [2]. Therefore, the derivative in the Lagrangian has to be replaced again by the covariant derivative

$$D_\mu \equiv \partial_\mu + i\frac{g}{2}\tau^i A^i_\mu. \tag{1.26}$$

where A_μ stands for eight massless gauge fields corresponding to the gluon octet that can be represented by

$$r\bar{b}, r\bar{g}, b\bar{g}, b\bar{r}, g\bar{r}, g\bar{b}, \frac{1}{\sqrt{2}}(r\bar{r} - b\bar{b}), \frac{1}{\sqrt{6}}(r\bar{r} + b\bar{b} - 2g\bar{g}). \tag{1.27}$$

In contrast to the electrically neutral photon as mediator of the electromagnetic interaction, gluons themselves are colour-charged and are able to couple directly to other gluons. Hence, in addition to the normal quark-gluon vertex, also three gluon and four gluon vertices exist.

However, combining the three colours with the three anti-colours suggests $3^2 = 9$ combinations which leads to an additional colour singlet state for the gluon

$$\frac{1}{\sqrt{3}}(r\bar{r} + g\bar{g} + b\bar{b}). \tag{1.28}$$

However, this state has to be excluded as it carries no net colour-charge which would lead to long range couplings between hadrons.

1.4 V-A Theory of the Weak Interaction

Experiments have shown, that parity is not conserved in the β decay and other processes which are mediated by the exchange of a W boson. This has been demonstrated for the first time in the Wu experiment by observing the beta decay of radioactive cobalt 60 nuclei. In this experiment, the spins of the nuclei have been aligned in the same direction by a magnetic field and the direction of the emitted electron was recorded. It was found, that most of the electrons are emitted in the direction of the spin, which implies that parity is not conserved in this decay.

One of the consequences of this violation is that spin and momentum are correlated for weakly interacting particles. States in which the spin of a particle is parallel to the direction of its velocity are called left-handed. If they are oriented in opposite directions, it is called right-handed. Massless particles created in a weak interaction are left-handed whereas antiparticles are right-handed. However, for massive particles, the degree of polarisation follows $\beta = -v/c$. This behaviour is described by the V-A theory[4].

[4] V-A stands for vector minus axial-vector.

In this theory, each Dirac-spinor ψ can be divided into two orthogonal chirality components

$$\psi = \frac{1}{2}(1+\gamma^5)\psi + \frac{1}{2}(1-\gamma^5)\psi = \psi_R + \psi_L \ , \tag{1.29}$$

where $P_R = \frac{1}{2}(1+\gamma^5)$ and $P_L = \frac{1}{2}(1-\gamma^5)$ are the corresponding chirality projection operators. The spinor ψ_R denote the right-handed and ψ_L the left-handed component of ψ. In the Feynman rules for the calculation of the invariant amplitude, the corresponding chirality components have to be used. For example, the $\mu\nu$ vertex in the W boson decay reads as:

$$\overline{\psi}(\mu)\gamma_\mu\psi_L(\nu) = \overline{\psi}(\mu)\gamma_\mu\frac{1-\gamma^5}{2}\psi(\nu) \ . \tag{1.30}$$

In this expression, two different terms can be identified:

- The vector-current
$$V^\mu = \overline{\psi}(\mu)\gamma^\mu\psi(\nu)$$
transforming like a four-vector and
- the axial-vector-current
$$A^\mu = \overline{\psi}(\mu)\gamma^\mu\gamma^5\psi(\nu)$$
which transforms like an axial-vector.

Axial-vectors behave like four-vectors under Lorentz-transformations, but keep the sign under parity transformations. In this construct, the vector and axial-vector component have the same weight. Therefore, parity conservation is maximally violated in this example. The fact that only left-handed particles and right-handed anti-particles participate in the weak interaction mediated by W bosons can be motivated as follows: The two properties of the projection operator P_L

$$P_L = P_L^2 \tag{1.31}$$

$$\gamma^\mu P_L = P_R\gamma^\mu \tag{1.32}$$

applied on Equation 1.30 lead to

$$\overline{\psi}_\mu\gamma_\mu\frac{1-\gamma^5}{2}\psi_\nu = \overline{\psi}_\mu\gamma_\mu\left(\frac{1-\gamma^5}{2}\right)^2\psi_\nu = \overline{(\psi_\mu)}_L\gamma_\mu(\psi_\nu)_L \ .$$

However, the coupling of the neutral mediator of the weak interaction, the Z boson, is more complicated. Instead of the purely V-A vertex factor of the W boson, a mixed form has to be used:

$$\gamma^\mu(c_V^f - c_A^f\gamma^5) \ .$$

The coefficients c_V^f and c_A^f depend on the particular quark or lepton (f) involved. They are determined by a single fundamental parameter θ_W, the weak mixing angle which is discussed in the next section.

1.5 The Glashow-Salam-Weinberg Model

With the introduction of the weak isospin I whose third component I_3 can take the values $\pm 1/2$, all fermions can be arranged in multiplets of this new quantum number.

The left-handed fermions form the isospin doublet ψ_L

$$\begin{array}{lccc} & & & I_3 \\ \text{Leptons:} & \begin{pmatrix} \nu_e \\ e^- \end{pmatrix}_L & \begin{pmatrix} \nu_\mu \\ \mu^- \end{pmatrix}_L & \begin{pmatrix} \nu_\tau \\ \tau^- \end{pmatrix}_L & 1/2 \\ & & & & -1/2 \end{array}$$

$$\begin{array}{lccc} & & & I_3 \\ \text{Quarks:} & \begin{pmatrix} u \\ d' \end{pmatrix}_L & \begin{pmatrix} c \\ s' \end{pmatrix}_L & \begin{pmatrix} t \\ b' \end{pmatrix}_L & 1/2 \\ & & & & -1/2 \end{array} \quad .$$

The right-handed leptons and quarks do not couple to the charged weak currents and are therefore arranged in singlets ψ_R:

$$I = 0: \;\; e_R^-,\; \mu_R^-,\; \tau_R^-,\; u_R,\; d_R,\; s_R,\; c_R,\; b_R,\; t_R \; .$$

In this context, the weak hypercharge is defined as $Y = 2(Q - I_3)$ where Q is the charge of the particle. In this notation, the left-handed leptons have a weak hypercharge $Y_L = -1$ whereas the right-handed leptons have $Y_R = -2$.

Combining the left- and right-handed components of ψ with the electromagnetic interaction into a single Lagrangian leads to the electroweak unification which is discussed in the following exemplarily for the leptons. The transformation of the left-handed leptons is described by the group $SU(2)_L$ which brings three additional vector fields $\vec{W}^\mu$ as discussed in Section 1.3. The transformation for the electromagnetic interaction is represented by the $U(1)$ group and requires one additional vector field, from here on denoted as B^μ.

To get a single Lagrangian for both interactions which is invariant under local gauge transformation, the derivative has to be replaced by this covariant derivative

$$D^\mu = \partial^\mu + ig\vec{T} \cdot \vec{W}^\mu + i\frac{g'}{2} Y B^\mu \; . \tag{1.33}$$

which reads for left-handed leptons ($\vec{T} = \vec{\tau}/2$ and $Y_L = -1$) as

$$D^\mu = \partial^\mu + i\frac{g}{2}\vec{\tau}\cdot\vec{W}^\mu - i\frac{g'}{2}YB^\mu \ . \tag{1.34}$$

The corresponding covariant derivative for right-handed leptons ($\vec{T} = 0$ and $Y_L = -2$) is then

$$D^\mu = \partial^\mu - ig'YB^\mu \ . \tag{1.35}$$

According to the coupling of the vector-fields to left-handed leptons, the fields of the two charged vector-bosons $W^\pm$ can be derived to be

$$W^{(\pm)\mu} = \frac{1}{\sqrt{2}}(W_1^\mu \pm W_2^\mu) \ . \tag{1.36}$$

However, the electromagnetic field A^μ cannot be identified with one of the remaining fields B^μ or W_3^μ as both couple to the neutrino. A proper linear combination of these two fields leads to a solution with a massless particle not coupling to neutrinos, the photon. The other orthogonal combination represents the gauge field of the neutral weak current with the Z boson as mediator:

$$A^\mu = B^\mu \cos\theta_W + W_3^\mu \sin\theta_W,$$

$$Z^\mu = -B^\mu \sin\theta_W + W_3^\mu \cos\theta_W \ .$$

The value of the weak-mixing angle θ_W is constraint by requiring a vanishing coupling of the photon field to the neutrino. The current value of the weak-mixing angle from measurements is $\sin^2\theta_W = 0.232$ [3].

As beautiful as the unification of electromagnetic and weak interaction appears, one big problem remains. Local gauge invariance has been achieved for the weak interaction by using the Proca Lagrangian for massless fields. Whereas this approach is valid for the photon and the gluon fields as their mediators are massless, the vector bosons are certainly massive. An approach to formulate a gauge theory which is able to handle massive fields as required for the $W^\pm$ and the Z boson is the Higgs mechanism, which will not be discussed here. A detailed introduction to this theory can be found in [2].

1.6 Cross-Section for Hadron Collisions

Since the discovery of the substructure of atoms by Rutherford in 1903, particle physics experiments have been based on the same method. High energetic probe particles are used to get information on the structure and the interactions of the target material. An important parameter of such

investigations is the resolution ΔR which can be achieved with a probe particle. Here, ΔR is a measure for the minimal distance between two objects which still can be resolved. For particle beams, this resolution is in the order of magnitude of their de Broglie wavelength

$$\lambda = \frac{h}{p} \tag{1.37}$$

where h is Planck's constant and p is the momentum of the beam particle.

To get deeper and deeper insights of the substructure of particles requires therefore increasing beam energies. In addition, to create particles with high masses in experiments and to study their properties, the centre-of-mass energy $\sqrt{s}$ of the particle collisions has to exceed the energy of the desired particle, related by Einstein's equivalence of mass and energy $E = mc^2$.

For fixed target experiments, $\sqrt{s}$ increases only with the square root of the energy. Therefore, collider experiments in which the particles of the two beams have equal energy E are favoured because of the linear increase of the centre-of-mass energy with E, following

$$\sqrt{s} = 2E. \tag{1.38}$$

1.6.1 Cross-Section

A measure for the probability for a certain process to take place in a collision of two particle beams is the cross-section σ which connects the reaction rate per target particle W with the incoming flux ϕ by

$$\sigma = \frac{W}{\phi}. \tag{1.39}$$

The unit of the cross-section is the barn [b] [5]. According to this equation, σ can be determined experimentally by measuring the reaction rate. From a theory point of view, this rate depends on the Hamiltonian of the interaction $H_{\text{interaction}}$ and can be derived from quantum mechanics. According to Fermi's Second Golden Rule, the reaction rate W reads as

$$W = \frac{2\pi}{\hbar}|M_{if}|^2\rho_f. \tag{1.40}$$

Here, ρ_f denotes the phase space which is available for the final states and M_{if} is the matrix element between the initial and final-state wave function ψ which reads as

$$M_{if} = < \psi_f|H_{\text{interaction}}|\psi_i > . \tag{1.41}$$

[5] The unit barn is defined by 1 b = 10^{-28} m^2.

1.6.2 Elastic Electron-Proton Scattering

To investigate the substructure of the proton, electrons are an ideal candidate for scattering experiments as they are to the current standard of knowledge point-like particles, which means that no substructure has to be taken into account. To resolve the substructure of the proton with an approximate diameter of about 1 fm, the de Broglie wavelength of the accelerated electrons must be at least in the same order of magnitude, which corresponds to an energy of 200 MeV. Increasing this energy leads to the emerging of finer substructures. In the following, the basics of elastic scattering are briefly discussed.

For the description of an elastic scattering, the transferred four-momentum is one of the key variables and can be expressed as

$$Q^2 = -q^2 = -(k - k')^2 \approx 4EE'\sin^2\left(\frac{\theta}{2}\right). \tag{1.42}$$

In this equation, k and k' denote the four-vectors of the electron respectively before and after the interaction. As both, the electron and the proton are charged spin-1/2 particles, the interaction of the charge of the electron with the spin of the proton has to be considered as well as the interaction of both spins. The latter is already included in the Mott cross-section [4]. The final differential cross-section for elastic electron-proton scattering is described by the Rosenbluth formula

$$\left(\frac{d\sigma}{d\Omega}\right) = \left(\frac{d\sigma}{d\Omega}\right)_{\text{Mott}} \cdot \frac{E'}{E}\left[\frac{G_E^2(Q^2) + \tau G_M^2(Q^2)}{1+\tau} + 2\tau G_M^2(Q^2)\tan^2\frac{\Theta}{2}\right] \tag{1.43}$$

This formula takes the dimension of the proton into account by introducing two form factors which depend on the transferred four-momentum. The electric form factor $G_E^2(Q^2)$ is the Fourier transformation of the charge distribution inside the proton whereas the magnetic form factor $G_M^2(Q^2)$ represents the distribution of the magnetic moment of the proton. The abbreviation τ is defined as

$$\tau = \frac{Q^2}{4M_P^2}. \tag{1.44}$$

and M_P denotes the mass of the proton.

1.6.3 Deep-Inelastic Scattering and PDFs

A further increase of the energy of the electron leads to inelastic scattering which can be observed for electron-proton collisions with a centre-of-mass energy larger than ≈5 GeV. In such interactions, the proton can be excited, which has to be taken into account for the energy conservation when deriving a formula for the differential cross-section.

The invariant mass W of the excited proton state after the interaction is calculated from the four-vector of the exchanged photon q and proton P before the interaction according to

$$W^2 = P'^2 = (P+q)^2 = M^2 + 2Pq + q^2 = M^2 + 2M\nu - Q^2 \tag{1.45}$$

with ν defined as

$$\nu = \frac{Pq}{M}. \tag{1.46}$$

In the case of a fixed target experiment, the protons are at rest in the laboratory frame and the last equation simplifies to $\nu = E - E'$. As a result of the excitation of the proton, which implies that $W > M$, the following constraint can be deduced from Equation 1.45 for inelastic processes:

$$2M\nu - Q^2 > 0 \tag{1.47}$$

For elastic processes, $M = W$ and the last expression reads as $2M\nu - Q^2 = 0$. With the new constraint for the inelastic scattering, the cross-section can be expressed to be

$$\frac{d^2\sigma}{d\Omega dE'} = \left(\frac{d\sigma}{d\Omega}\right)_{\text{Mott}} \cdot \left[W_2(Q^2,\nu) + 2W_1(Q^2,\nu)\tan^2\frac{\Theta}{2}\right]. \tag{1.48}$$

As a result of the additional degree of freedom due to the inelasticity of the process, the cross-section is presented in double differential form. Again, the two functions $W_{1,2}(Q^2,\nu)$ describe the distribution of magnetic momentum and electric charge of the proton. In this context, a new variable is introduced, the Bjorken-x which is defined as

$$x = \frac{Q^2}{2M\nu} \tag{1.49}$$

and corresponds to the fraction of the proton momentum carried by the scattered parton in the infinite momentum frame in which the proton is regarded as a collinear stream of fast moving partons and masses are negligible. Instead of the form factors $W_{1,2}(Q^2,\nu)$, the dimensionless form factors $F_{1,2}$ are commonly used which depend on the Bjorken-x and can be expressed as follows:

$$F_1(x,Q^2) = Mc^2 W_1(Q^2,\nu) \tag{1.50}$$

$$F_2(x,Q^2) = \nu W_2(Q^2,\nu). \tag{1.51}$$

Measurements of the electric form factor F_2 of the proton have shown that it does not depend on Q^2 which suggest scattering at a point-like charge [4]. This was a strong hint for the proton and the neutron being composite of point-like objects. These constituents have been identified with the predicted quarks.

In addition, the Callan-Gross-Relation which connects the electric and magnetic form factor for

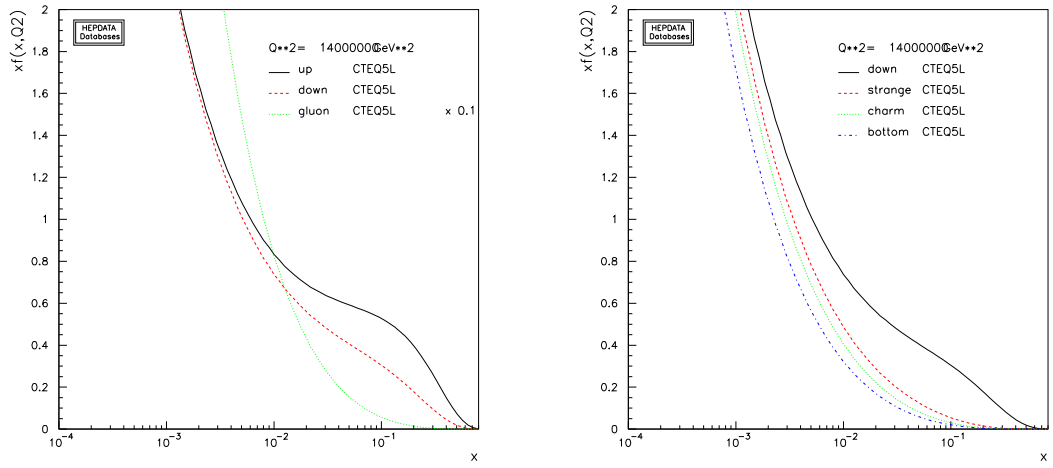

Figure 1.1: *CTEQ5L parton density functions for different partons. The plot on the left hand side shows the functions for up- and down-quarks as well as for gluons (scaled by a factor of 0.1). The plot on the right hand side contains the functions for down-, strange-, charm- and bottom-quarks [7].*

spin-1/2 particles to

$$2xF_1(x) = F_2(x) \tag{1.52}$$

has been proven experimentally. The combination of the results implies that quarks are charged, point-like particles with spin 1/2.

This finally led to the QCD parton model and the structure functions $F_{1,2}$ can be expressed in terms of the Parton Distribution Functions (PDFs). Here, the quantity $f_i(x, Q^2)dx$ represents the probability to find a parton of type i carrying a fraction between x and $x + dx$ of the total nucleon momentum. Here, i stands for the quark and anti-quark flavour as well as for the gluon.

Therefore, these parton density functions are important ingredients for any kind of cross-section prediction for proton interactions and they are determined in many experiments. The combination of the different experimental results is done by independent working groups like CTEQ [5] or MRST [6]. As an example, the CTEQ5L parton distribution functions are shown in Figure 1.1.

Assuming a transverse momentum of the hard parton interaction larger than 10 GeV, the total

proton-proton cross-section for a certain subprocess can be expressed as

$$\sigma_{hadronic} = \sum_{ij} \int \int f_i(x_1, Q^2) f_j(x_2, Q^2) \hat{\sigma}_{\text{partonic}} dx_1 dx_2 \; . \tag{1.53}$$

where x_1 and x_2 are the momentum fractions of the interacting partons. $\hat{\sigma}_{\text{partonic}}$ describes the cross-section of the respective hard subprocess.

1.7 The Strong Interaction

In quantum field theories, the strength of an interaction is described by the corresponding coupling like the well-known fine-structure constant α for quantum electrodynamics. However, as an effect of vacuum polarization, the charge of the electron is partially screened. These virtual processes are called self-energy terms and contribute to the experimentally measured charge and mass of the electron. In addition, the screening and thus the measured charge depend on the distance between the interacting particles, which is correlated with the transferred momentum Q^2.

To be able to compare calculation with measurements, these virtual processes have to be taken into account. However, there is no restriction on the momentum k of the particles in the loop diagrams, leading to logarithmically divergent terms. To derive finite results, renormalisation methods have been developed. Applying these methods, the renormalised coupling becomes dependent on the chosen renormalisation scale μ_R.

For the electromagnetic interaction, the effective coupling increases only very slowly with the energy-momentum transfer. For small Q^2, the coupling can be approximated by $\alpha(0) = \frac{1}{137}$ and increases only slightly to about $\alpha(M_Z) = \frac{1}{128}$ for energies in the order of the mass of the Z boson.

This situation changes for the strong interaction. Quark-antiquark clouds lead to a similar screening of the quark colour-charge. However, in contrast to the photon which is electrically neutral, the gluons themselves are carrying colour, which has to be considered for this process. It turns out that their contribution is larger than the screening and works in the opposite direction leading to the effect of antiscreeing for the colour-charge.

After renormalisation, the dependence of the strong coupling can be expressed as

$$\alpha_S(Q) = \frac{\alpha_S(\mu_R)}{1 - \frac{\beta_0}{2\pi} \alpha_S(\mu_R) \ln(\frac{Q}{\mu_R})} \tag{1.54}$$

with $\beta_0 = \frac{1}{3}(2n_f - 33)$. Here, the number of contributing quark flavours is denoted with n_f. Equation 1.54 allows the extrapolation with respect to Q^2 of the coupling constant once it has been measured experimentally for a fixed $Q^2 = \mu_R^2$. In addition, the strong coupling vanishes for large

transferred momenta

$$\lim_{Q^2 \to \infty} \alpha_s(Q^2) = 0 \tag{1.55}$$

leading to the phenomenon referred to as asymptotic freedom. The interaction between partons which are close together becomes very small and the partons forming the nucleus behave similar to essentially free particles. However, the theory framework above is only valid for small values of α_s as otherwise perturbation theory will not work.

To describe the long-range behaviour of the strong interaction which corresponds to small values of Q^2 the string model can be used. In analogy with the electric field lines between two electric charges, quarks are held together by colour lines. However, due to the gluon self-coupling, these colour lines are pulled together by the gluon-gluon interaction in the form of a tube or string. Assuming a constant energy density of this string leads to what is called the quark confinement. As long as the distance between two quarks gets larger the energy of the connecting string is increasing. At some point, the energy becomes sufficiently large to create a quark-antiquark pair, which explains why only colour-neutral objects are observed experimentally.

1.8 Z Boson Production at Hadron Colliders

Since having been discovered at the Super Proton Antiproton Synchrotron at CERN in 1983, the properties of the Z boson have been studied in different experiments, especially at the four experiments of the Large Electron Positron Collider (LEP). Their results for the mass and the width of the Z boson including the combined best fit of all four experiments are summarised in Figure 1.2 and Figure 1.3. Both LEP values

$$\begin{aligned} M_Z &= 91.1876 \pm 0.0021\,\mathrm{GeV} \\ \Gamma_Z &= 2.4952 \pm 0.0023\,\mathrm{GeV}\ . \end{aligned}$$

represent the current world average [3].

As the cross-section for Z boson production increases with the centre-of-mass energy, a huge number of Z bosons will be produced under LHC conditions. According to next-to-leading order (NLO) calculations the cross-section is about 55.5 nb [9] which increases to 60.0±1.9 nb in NNLO calculations [10]. The branching ratio for the decay of the Z boson into muons is (3.366±0.007)·10^{-2} [3] which leads to an overall NNLO cross-section for the process $pp \to Z \to \mu\mu$ of 2.02 nb.

A detailed overview of the expected cross-sections of different processes at the LHC and the Tevatron is presented in Figure 1.4. Additionally, the flavour decomposition of the total cross-section for Z boson production at both colliders is presented in Figure 1.5.

The coupling of the Z boson to fermions is not purely of the V-A structure as already discussed

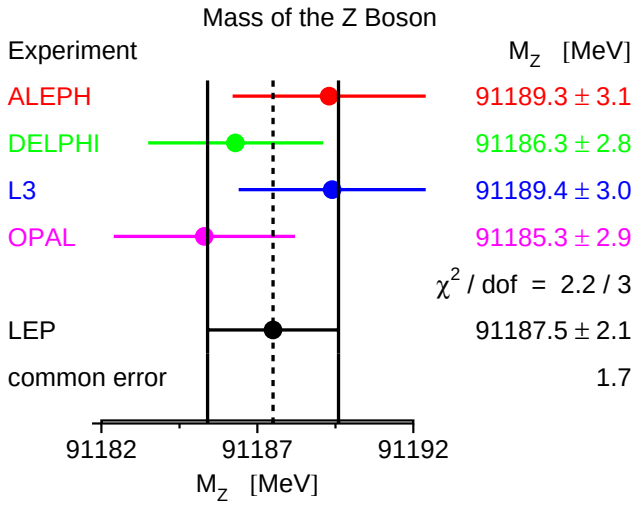

Figure 1.2: *Results of the measurement of the Z boson mass from all four LEP experiments including the result of the combined best fit [8].*

and for the corresponding Drell-Yan process, interference terms of the Z boson and the photon have to be taken into account. To lowest order, the amplitude for the quark anti-quark fusion can be described by the exchange of a photon and Z boson:

$$q\bar{q} \rightarrow \gamma^*/Z^0 \rightarrow l^+l^- \,. \tag{1.56}$$

Because of its clear signature, only the decay of the Z boson into two leptons is investigated in this study. The corresponding cross-section for this process can be derived in the effective Born approximation to be

$$\frac{d\hat{\sigma}}{d\Omega} = \frac{1}{64\pi^2\hat{s}}\left[\mathcal{A}_0\left(1+\cos^2\theta\right)+\mathcal{A}_1\cos\theta\right] \tag{1.57}$$

where θ is the scattering angle in the parton centre-of-mass system. The two functions $\mathcal{A}_0$ and $\mathcal{A}_1$ contain already several aspects of higher order calculations as for example the effective electromagnetic charge, which takes the photon vacuum polarisation into account [12]. The Z propagator term together with the overall normalisation factor of the neutral-current as well as the vector and axial-vector coupling constants are included. The latter is sensitive to the effective mixing angle. In addition, a forward-backward asymmetry of the cross-section in Equation 1.57 can be derived and is

$$\hat{A}_{FB} = \frac{\hat{\sigma}_F - \hat{\sigma}_B}{\hat{\sigma}_F + \hat{\sigma}_B} = \frac{3}{8}\frac{\mathcal{A}_1}{\mathcal{A}_0} \,. \tag{1.58}$$

at the parton level.

Having in mind the high accuracy measurements of its parameters like mass or width of the LEP experiments, the large number of events containing Z bosons at the LHC can be used for many purposes. High precision measurements of the parameters of the Standard Model allow to decrease

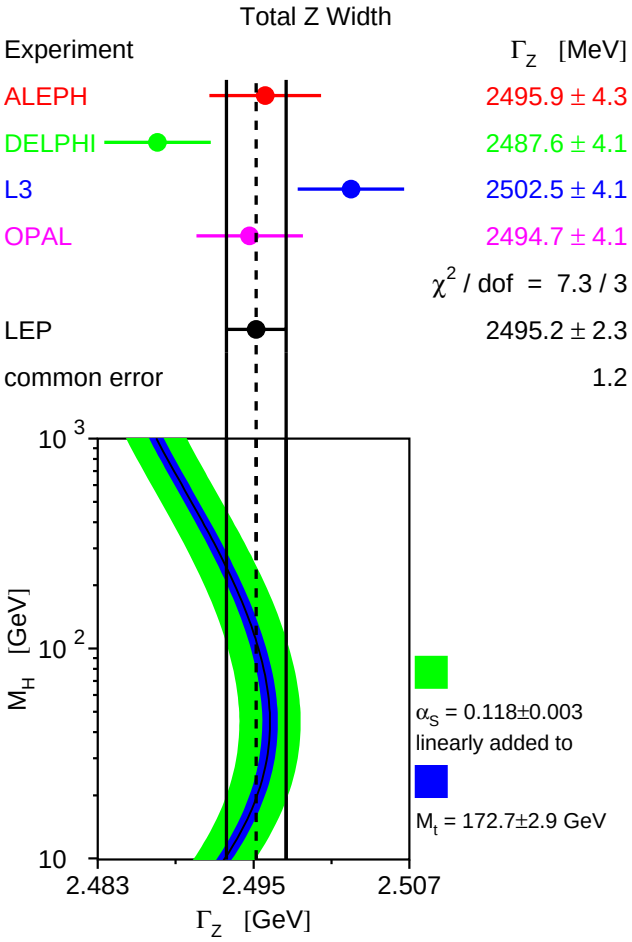

Figure 1.3: *Results of the measurement of the width of the Z boson from all four LEP experiments including the result of the combined best fit. The lower part of this figure shows the influence of the width of the Z boson on the Standard Model prediction for the mass of the Higgs boson [8].*

the uncertainties of the prediction and observed deviations can be interpreted as a hint to physics beyond the Standard Model. Also additional heavy neutral gauge bosons might be discovered at the LHC.

In addition, it becomes possible to monitor the luminosity of the collider or to calibrate different detector components like the electromagnetic calorimeter using the decay of the Z boson into two electrons.

Events which contain a Z boson and one additional jet open a window to a variety of interesting measurements. As the subprocesses

$$qg \rightarrow Z + q \tag{1.59}$$

$$\bar{q}g \rightarrow Z + \bar{q} \tag{1.60}$$

$$q\bar{q} \rightarrow Z + g \tag{1.61}$$

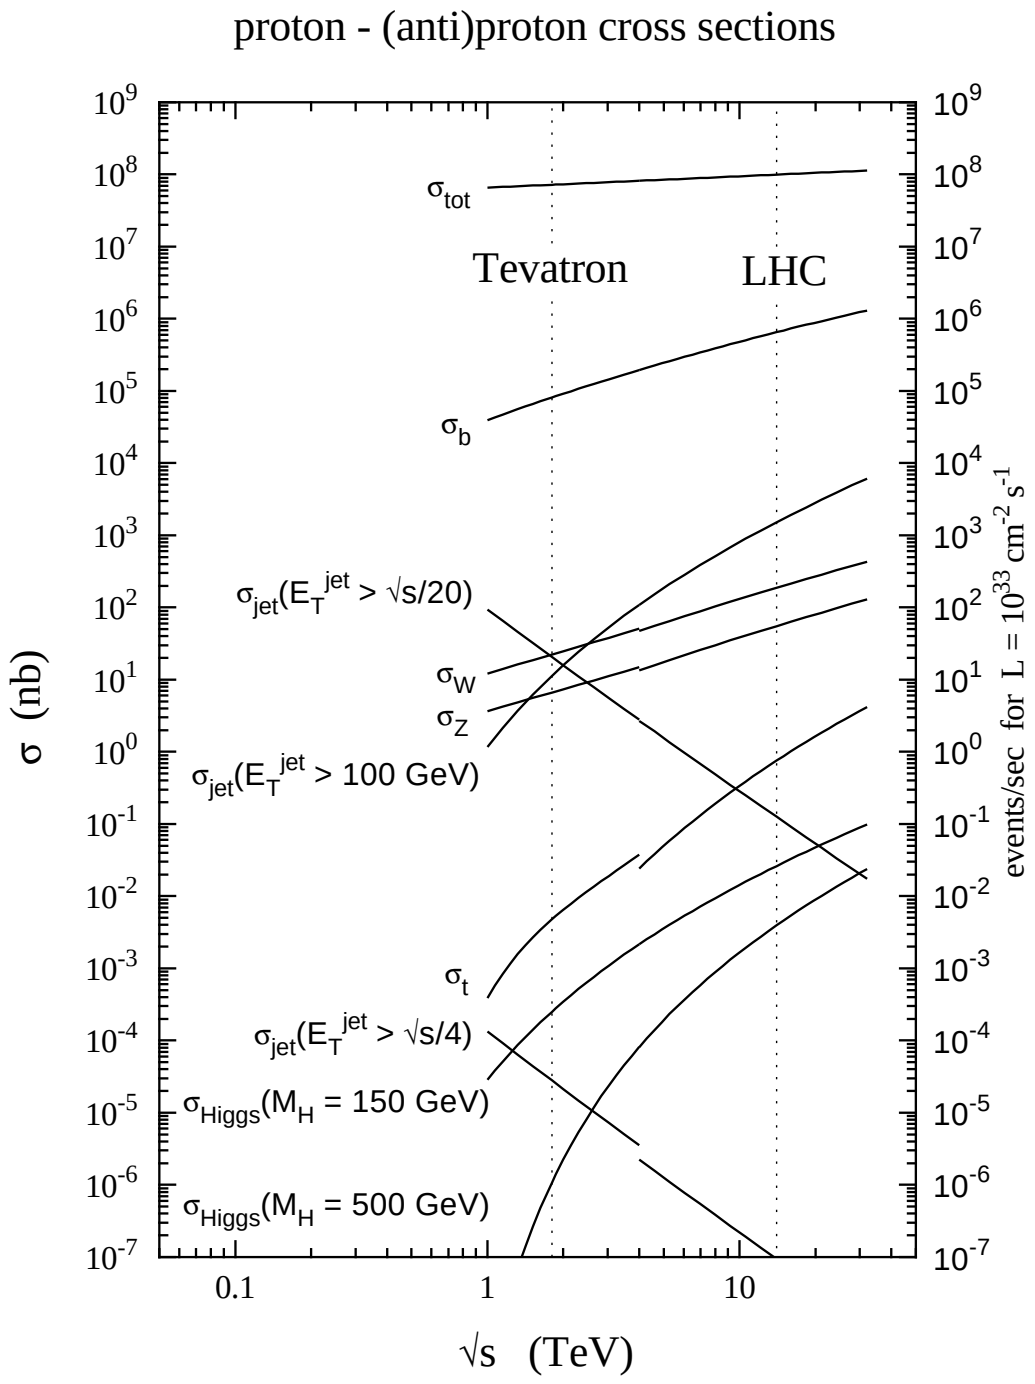

Figure 1.4: *NLO cross-section for hard scattering versus the centre-of-mass energy for proton-(anti)proton collisions. The vertical dashed lines indicate the design centre-of-mass energy of the Tevatron and the LHC. The cross-section for the Z boson production is σ_Z =55.5 nb [9].*

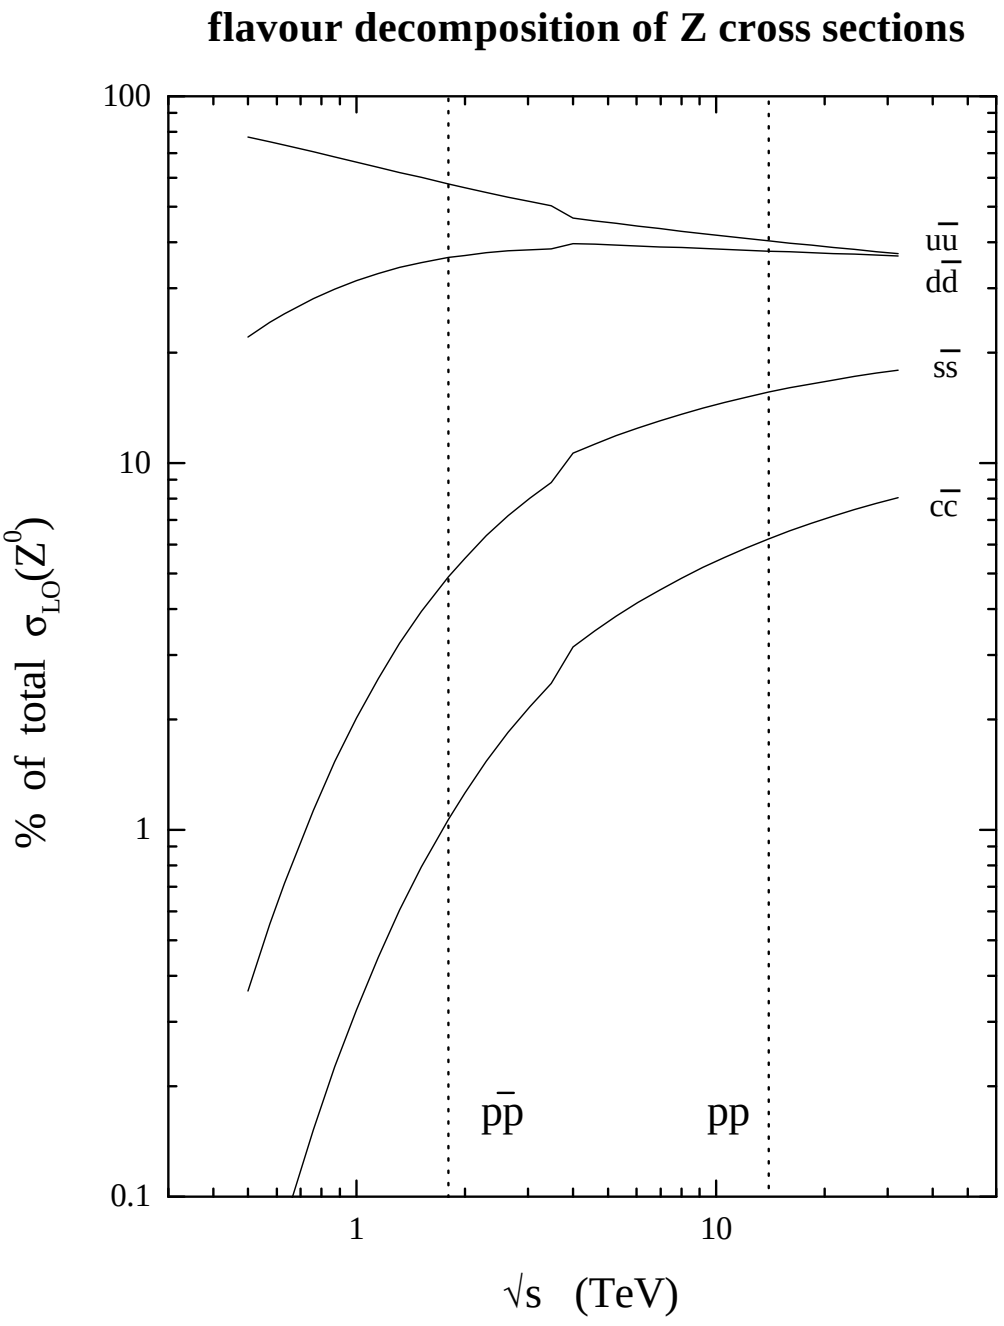

Figure 1.5: *Overview of the parton decomposition of the total cross-section for Z boson production at $p\bar{p}$ and pp colliders. The individual contributions are displayed as a percentage of the total leading-order cross-section [11].*

are sensitive to the gluon density the precise measurement of these processes can help to constrain the gluon PDF in the proton. The corresponding Feynman graphs are shown in Figure 1.6. In addition, these events are an excellent candidate for jet calibration. As the momentum of the involved parton is balanced by the Z boson, its kinematic is linked to the parton. The momentum of the Z boson decaying into two muons can be measured with high accuracy and so be used to calibrate the jet originating from the balanced parton.

$qg \rightarrow Z + q$

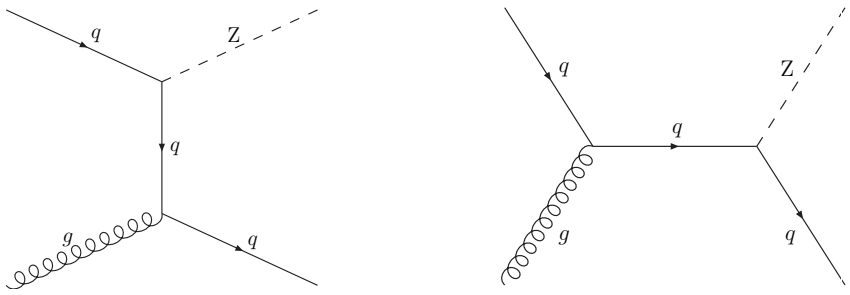

$\bar{q}g \rightarrow Z + \bar{q}$

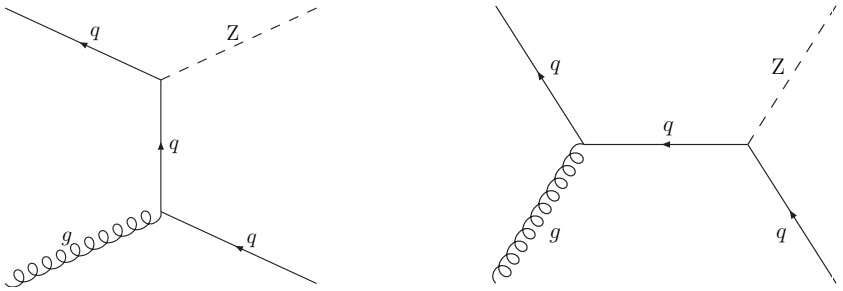

$q\bar{q} \rightarrow Z + g$

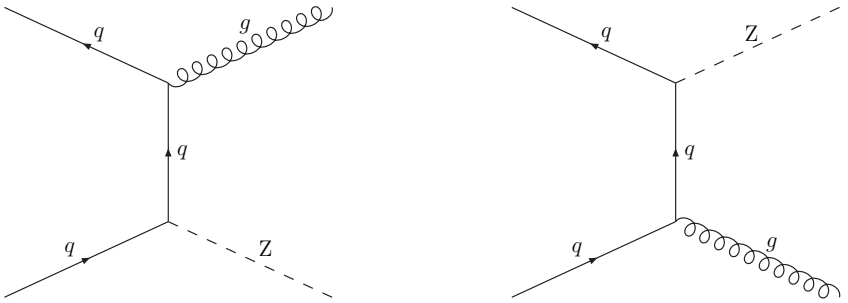

Figure 1.6: *Feynman diagrams for the Z boson production with one additional jet.*

The CMS Detector at the Large Hadron Collider

Over the last decades, the Standard Model of particle physics has been tested in numerous experiments and several of its free parameters, which have to be determined experimentally, have been measured with high accuracy. Predicted particles have been discovered and many questions are solved. Nevertheless, the experiments up to now have left some questions unanswered like the origin of mass, a unified description of all forces as well as the search for particles or phenomena responsible for dark matter and dark energy.

The Large Hadron Collider (LHC) will start operation this year at CERN near Geneva and is expected to answer several of the questions raised above. In this chapter, the basic properties of the collider are discussed and the associated experiments are introduced focussing on CMS.

2.1 The Large Hadron Collider

The LHC is a hadron-hadron collider which has been constructed in the 27 km long tunnel of the former Large Electron Positron Collider, being operational at CERN from 1989 until 2000 and provided electron-positron collisions with a centre-of-mass energy of up to 209 GeV. The LEP experiments precisely measured the properties of the W and Z boson and confirmed the predictions of the Standard Model with high accuracy. However, the design value of the centre-of-mass energy at LEP was mainly constraint by the loss of energy due to synchrotron radiation, which increases significantly with the energy of the accelerated particle. The energy loss per turn follows

$$-\Delta E = \frac{4\pi\alpha}{3R}\beta^3\gamma^4 \quad \text{with} \quad \beta = \frac{v}{c} \approx 1 \quad \text{and} \quad \gamma = \frac{E}{mc^2}. \tag{2.1}$$

According to this equation, an electron with an energy of 100 GeV at LEP radiates about 2.3 GeV of its energy per turn, which has to be compensated for by the accelerator facility.

To decrease this loss of energy due to synchrotron radiation and to reach higher centre-of-mass energies, two approaches are possible. Either the radius of the facility has to be increased - the optimal way would be a linear collider - or particles with a higher mass than electrons have to be accelerated. The first concept is followed by the International Linear Collider (ILC) project [13] which is currently in its planning phase.

The second approach can be realised by the acceleration of protons. As these particle are about 2000 times heavier than electrons and the energy-mass fraction enters Equation 2.1 at a power of four, synchrotron radiation becomes negligible. But in contrast to the electron, the proton is a composite object of quarks and gluons, each carrying only a fraction of the proton's momentum. Therefore, the beam energy at hadron colliders has to be foreseen to be well above the energy scale of the desired interactions. The Tevatron collider at FermiLab [14] is based on this concept. A proton and an anti-proton beam are accelerated in the 6.3 km long ring to an energy of 980 GeV per beam. With its centre-of-mass energy of 1.96 TeV it is today's most powerful proton collider hosting the two experiments CDF [15] and DØ [16].

In 2008, the Large Hadron Collider will start operation and provide two proton beams, circulating in opposite directions with a projected energy of 7 GeV each. This time, the beam energy is limited by the maximal strength of the magnetic dipole field compensating the centrifugal force of the charged particles due to the curvature of the ring. For the LHC, superconducting dipole magnets have been developed yielding a nominal magnetic field of 8.33 T, which corresponds to a 7 TeV beam that can be hold on track in the tunnel. The magnet coils are made of niobium-titanium (NbTi) cables which become superconducting below a temperature of 10 K. Therefore, they are cooled down with superfluid liquid helium of a temperature of about 1.9 K for operations. With this set up, a centre-of-mass energy of the proton-proton collisions of 14 TeV can be reached at the LHC.

As the cross-section of most of the physics processes to be investigated at the LHC energy scale is very small, only a tiny fraction of the collisions will yield such events of interest.

In general, the event rate $\dot{N}$ of a certain physics process is proportional to its cross-section σ and the luminosity $\mathcal{L}$ of the collider:

$$\dot{N} = \mathcal{L} \cdot \sigma. \tag{2.2}$$

At colliders, the beams consist of several equidistant bunches of particles. The luminosity for head-on collisions of two such bunches consisting out of n particles with a collision frequency f is given by

$$\mathcal{L} = f \, \frac{n^2}{4\pi\sigma_x\sigma_y}, \tag{2.3}$$

where σ_x and σ_y characterize the approximately Gaussian beam profiles in horizontal and vertical

directions [17]. Taking more realistic assumptions into account, the luminosity for a proton-proton collider can be expressed to be

$$\mathcal{L} = f \frac{\gamma N_B N_P^2}{4\pi\epsilon_n \beta^*} F, \tag{2.4}$$

where ϵ_n stands for the normalized transverse emittance, a measure of the phase space area associated with either of the two transverse degrees of freedom of the beams. At the LHC, the design value for this parameter is 3.75 μm. β^* is the amplitude function at the interaction point with a design value of 0.55 m and F is the reduction function due to the crossing angle. γ is the Lorentz factor, f the revolution frequency and the number of bunches per beam consisting of N_P protons is indicated as N_B.

At the LHC, the nominal bunch intensity may reach up to $1.15 \cdot 10^{11}$ protons and a beam will consist of 2808 bunches with a spacing of 25 ns which corresponds to a collision frequency of 40 MHz [18]. Per bunch crossing, between 10 and 20 inelastic proton-proton collisions will occur as a result of the large number of protons per bunch. This leads to the problem that a physics process of interest occurring during a bunch crossing has to be extracted out of 20 accompanying collisions. These additional collisions are called pile-up events.

After a ramp up phase, the LHC will be operated at a luminosity of $2 \cdot 10^{33}$ cm^{-2}s^{-1} during the first three years. This time is commonly referred to as the low luminosity running. Later, the design value of luminosity of 10^{34} cm^{-2}s^{-1} will be reached and accordingly this period is called high luminosity running. A detailed overview of the LHC machine parameters is presented in Table 2.1.

The CERN accelerator complex, which is schematically presented in Figure 2.1, consists of several facilities for the subsequent acceleration of the protons. The LINAC 2 (50 MeV), the Booster (1.4 GeV) and the Proton Synchrotron (PS) deliver protons with an energy of 25 GeV to the Super Proton Synchrotron (SPS). Here, the protons are accelerated in 24 cycles to an energy of 450 GeV before being injected into the Large Hadron Collider. The filling of each of the two LHC rings takes about four and a half minutes. Once the filling procedure is finished, the protons are accelerated for 20 minutes to their nominal energy of 7 TeV [19]. After this procedure, the experiments can take data at the interaction points for about 15 hours before the decrease in luminosity gets too high and the ring has to be refilled. Beside the proton runs, the LHC is designed to collide heavy ions like lead (Pb) with a centre-of-mass energy of 1148 TeV.

At four interaction points, particle detectors are installed to observe the collisions. Two of them are designed for special purposes: the ALICE[1] experiment for heavy ion physics and the LHCb[2] experiment for b-physics. In contrast, the ATLAS[3] and CMS detectors are two complementary general purpose detectors designed to explore new physics and to probe the electroweak symme-

[1] A Large Ion Collider Experiment [20]
[2] Large Hadron Collider beauty experiment [21]
[3] A Toroidal LHC ApparatuS [22]

Quantity	Number
Circumference	26 659 m
Design luminosity	10^{34} cm^{-2}s^{-1}
Nominal energy of protons	7 TeV
Injection energy of protons	450 GeV
Number of bunches per proton beam	2808
Number of protons per bunch	$1.15 \cdot 10^{11}$
Nominal bunch separation	25 ns
Dipole operating temperature	1.9 K
Number of magnets	9593
Number of main dipoles	1232
Peak magnetic dipole field	8.33 T
Number of main quadrupoles	392
Number of RF capacities	8 per beam
Frequency	11246 Hz
Crossing angle at interaction point	300 μrad
RMS beam radius at interaction point	16.7 μm
Stored energy per beam	362 MJ
Average number of collision per bunch crossing	≈ 20

Table 2.1: *Listing of the most important design LHC machine parameters for the phase of high luminosity running.*

try breaking. LHCf is a small experiment which will measure particles produced very close to the direction of the beams. Their detectors are located 140 m from the ATLAS interaction point. TOTEM is designed to measure the total proton-proton cross-section using extreme forward detectors which are installed at distances between 150 m and 400 m from the CMS interaction point. In the following, the CMS experiment is introduced in detail.

2.2 The CMS Experiment

CMS is one of the largest scientific collaborations. More than 2000 people from 181 institutes in 38 countries[4] are involved in the construction of the detector, the development of the software framework and the preparation of physics analyses. The name Compact Muon Solenoid already reflects the major design goals of the detector:

- **Compact:** With its overall length of 21.5 m and a diameter of 15 m, the spatial dimensions of CMS are considerably smaller than those of ATLAS, the second general purpose detector at the LHC. For comparison, ATLAS is 46 m long, has a diameter of 25 m and a total weight

[4]The numbers are from May 2007.

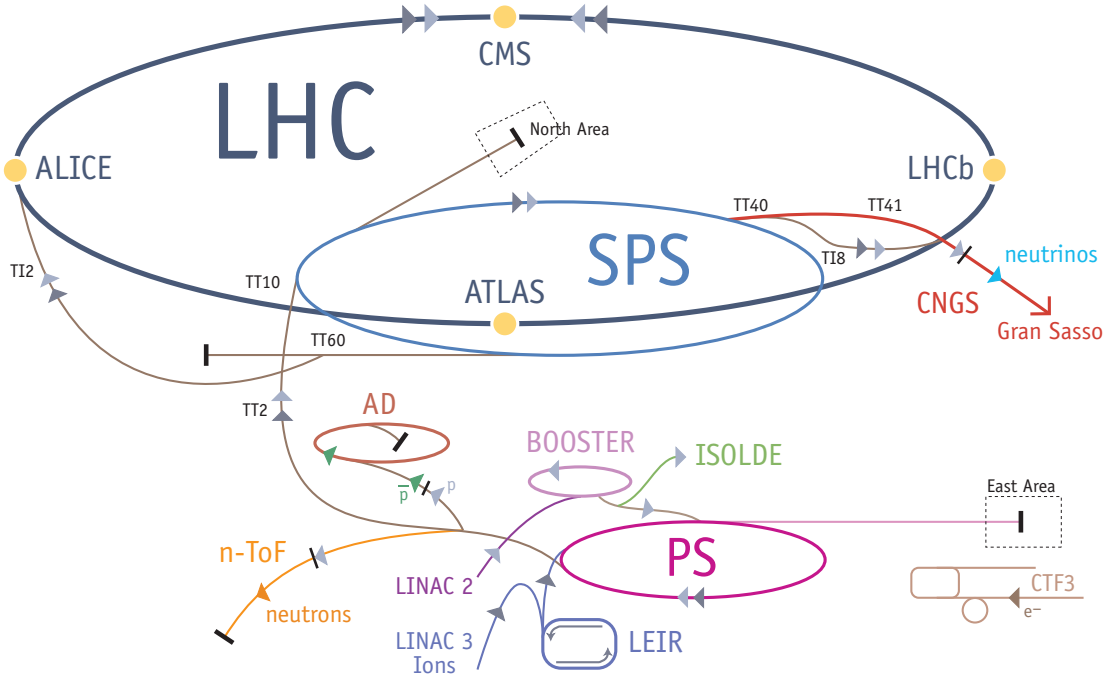

Figure 2.1: *Overview of the CERN accelerator complex. The protons are accelerated by several successive facilities to the LHC injection energy of 450 GeV before being accelerated to the nominal LHC energy of 7 TeV [19].*

of 7000 t, which is lightweight compared to CMS with its 12500 t. The latter is the result of the huge return joke which is made out of steel and guides the magnetic field outside of the magnetic coil.

- **Muon:** One of the key features of the detector is the identification, triggering and reconstruction of muons with high precision at high luminosities. It is designed to provide an efficient identification and good reconstruction of the momentum for muons up to a transverse momentum of 1 TeV over a wide geometric coverage.

- **Solenoid:** The design of the CMS detector follows a traditional approach with one single superconducting magnet coil providing the magnetic field for both, the inner detector as well as the muon system.

The detector is constructed in a modular way which has allowed to assemble the different components of the detector aboveground at the LHC access point 5 located at Cessy (France), near Geneva. Once a component has been assembled, it could be lowered down into the cavern, which is in a depth of about 100 m. A detailed overview of the detector and the collaboration is presented in Figure 2.2 and a brief description of the different detector components is given in the following [24].

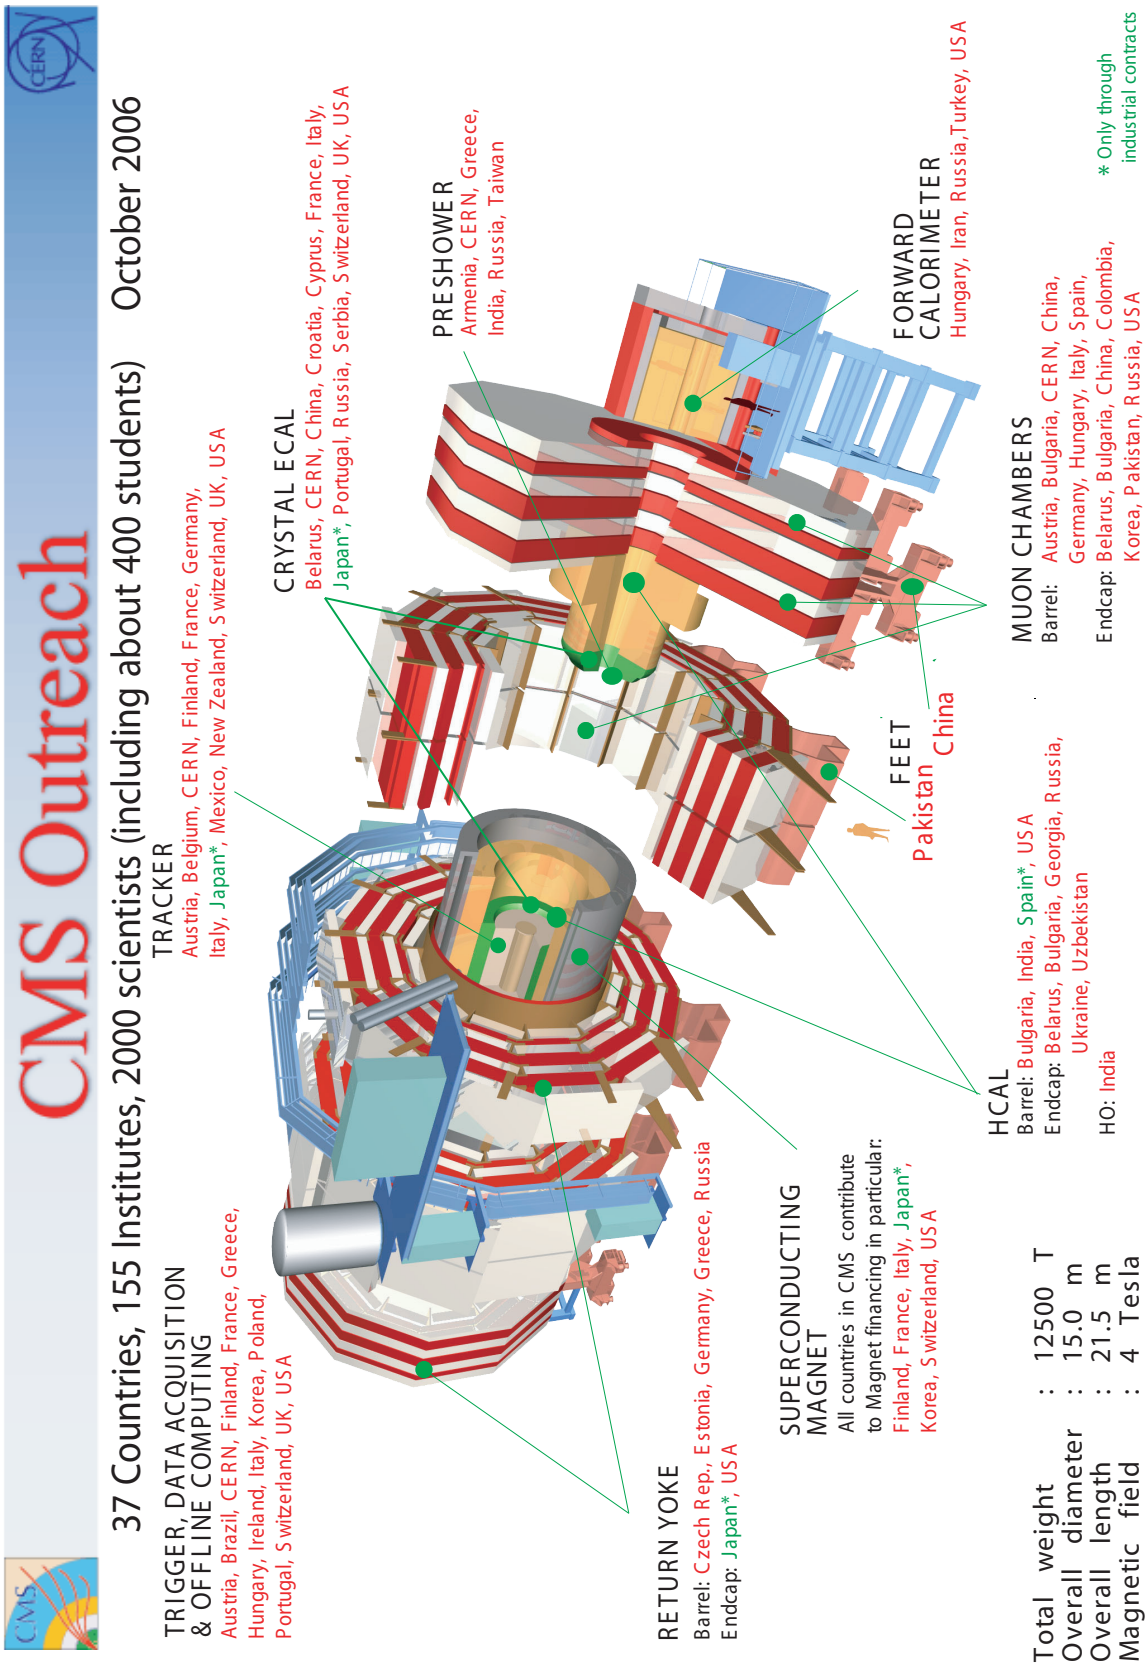

Figure 2.2: *Overview of the CMS detector and the collaboration [23].*

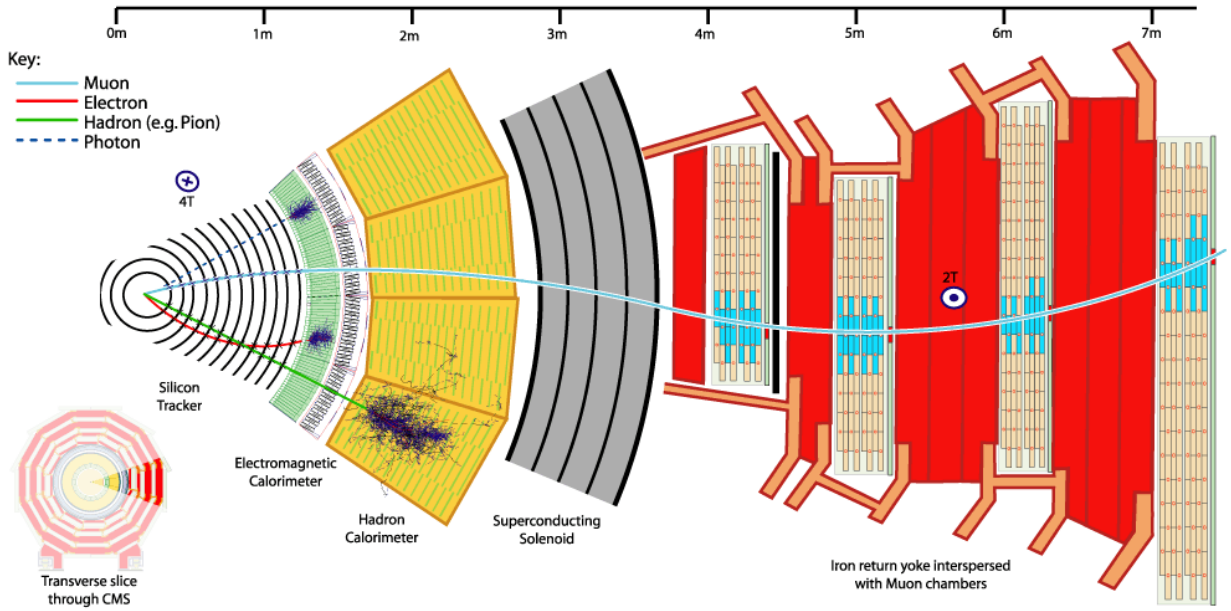

Figure 2.3: *Slice of the CMS detector with the tracks of an electron, a photon, a hadron (e. g. pion) and a muon. The electron and the photon deposit their whole energy in the electromagnetic calorimeter. The hadron passes this region of the detector without losing much of its kinetic energy and reaches the hadron calorimeter where it is stopped. Only muons are able to pass the detector material including the superconducting solenoid and can be detected in the muon system.*

2.2.1 Coordinate System of the Detector

The coordinate system adopted by the CMS collaboration has its origin at the nominal collision point inside the detector. Depending on the geometry of the component to be described, cartesian or polar coordinates are used as defined below:

Cartesian Coordinates

- x-axis: horizontal, pointing radially inwards towards the centre of the LHC
- y-axis: vertical, pointing upwards
- z-axis: horizontal, pointing tangentially to the beam pipe in westward direction in order to build a right-handed coordinate system

Polar Coordinates

- r: radial distance to the beam pipe

- Θ: polar angle with respect to the z-axis as defined above. In this definition, the +z-direction corresponds to $\Theta = 0$ and the -z-direction to $\Theta = \pi$

- ϕ: azimuthal angle measured in the x-y-plane with respect to the x-axis.

Another commonly used quantity related to the polar angle is the rapidity y. This dimensionless quantity, which is invariant under longitudinally boosts in the z-direction, is defined as

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right). \tag{2.5}$$

Instead of the rapidity, the purely geometrical defined pseudorapidity

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right] \tag{2.6}$$

is favoured which becomes identical to y for massless particles. For massive particles the difference between both quantities is negligible at high energies. The angular distance between two point-like objects can be expressed as

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} \tag{2.7}$$

2.2.2 Inner Tracking System

The inner tracking system is designed for an efficient measurement of trajectories of charged particles emerging from the interaction point of the proton-proton collision as well as the precise reconstruction of secondary vertices. At the LHC design luminosity of 10^{34} cm^{-2}s^{-1} there will be about 1000 particles from up to 20 collisions per bunch crossing. Therefore, besides the precise and fast reconstruction of these tracks, the intense particle flux is challenging concerning radiation damages of this innermost detector layer which has to operate in this harsh environment for an expected lifetime of 10 years. To master the challenges of a sufficient high granularity, fast response and radiation hardness, CMS employs an all-silicon approach for this detector component.

The tracking system is composed of a pixel and a strip detector, surrounding the interaction point with a total length of 5.8 m and a diameter of 2.5 m. The whole volume of the tracker is placed in the homogenous magnetic field of 4 T and provides an acceptance up to a pseudorapidity of $|\eta| < 2.5$. With about 200 m^2 of active silicon area, it is the largest silicon tracker ever built [25] [26]. A schematic overview of its different subsystems is drawn in Figure 2.4.

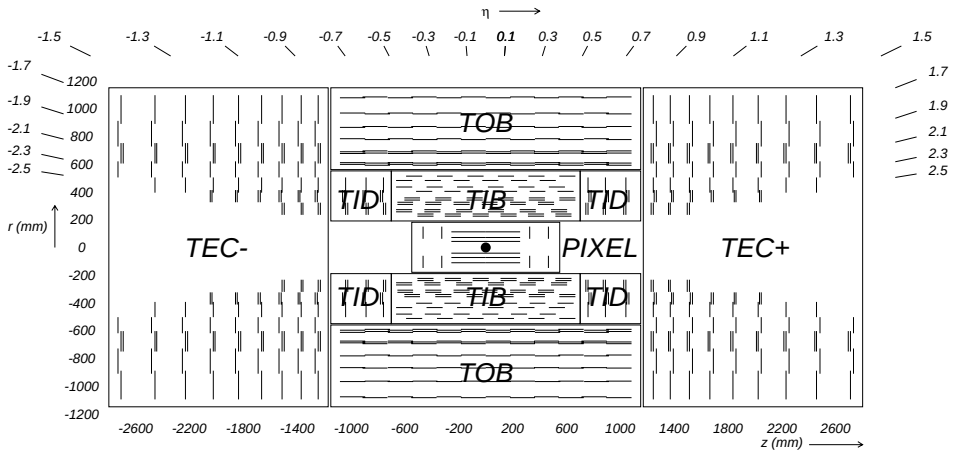

Figure 2.4: *Cross-section through the CMS tracking system [24].*

Silicon Pixel Detector

The silicon pixel detector consists of three barrel layers with a cell size of $100\times150\,\mu\text{m}^2$ measured in x and y coordinates. The layers are arranged in concentric rings around the beam pipe at radii of 4.4 cm, 7.3 cm and 10.2 cm and have a length of 53 cm. It is completed with two endcaps of a turbine-like geometry at each site of the barrel, covering radii between 6 and 15 cm and are located at z-positions of ±34.5 cm and ±46.5 cm.

The 1440 modules cover an area of 1 m^2 and provide with 66 million pixels precise tracking points in r-ϕ resulting in a excellent resolution of the impact parameter, which is required for a good secondary vertex reconstruction.

Silicon Strip Detector

At intermediate radii, the reduced particle flux allows the usage of cost-saving strip detectors. The silicon strip detector covers an area of about 200 m^2 and splits into four different subsystems. To obtain information about all three coordinates r, ϕ and z, some of the layers employ double-sided modules with a stereo angle of 5.7°.

- **Tracker Inner Barrel** (TIB): 4 layers at 20 cm $<$ r $<$ 55 cm, $|\text{z}| < 70$ cm
- **Tracker Outer Barrel** (TOB): 6 layers at 55 cm $<$ r $<$ 110 cm, $|\text{z}| < 120$ cm
- **Tracker Inner Disk** (TID): 2×3 disks at 20 cm $<$ r $<$ 55 cm, 70 cm $<|\text{z}|<$ 120 cm
- **Tracker EndCap** (TEC): 2×9 disks at 20 cm $<$ r $<$ 110 cm, 120 cm $<|\text{z}|<$ 180 cm

The first two layers of the TIB and TOB are made of the stereo-modules and the strips are parallel to the beam pipe. Two rings of the TID and three rings of the TEC are also stereo-modules with strips oriented perpendicularly to the beam pipe.

2.2.3 Electromagnetic Calorimeter

Electrons, positrons and photons create an electromagnetic shower, a cascade of photons and electron-positron pairs, in the calorimeter material through bremsstrahlung and pair production. Via Compton scattering and the photoelectric effect, these particles can interact with the electrons of the detector material and so deposit their energy in the calorimeter. The energy deposit is proportional to the kinetic energy of the incident particle and can be converted to an electrical signal by photodetectors.

In CMS, the Electromagnetic Calorimeter (ECAL) is a hermetic homogeneous calorimeter made of 61200 lead tungstate ($PbWO_4$) crystals in the central barrel part and closed by 7324 crystals in each of the two endcaps. The high density of 8.28 g/cm^3, the short radiation length X_0 of 0.89 cm and the small Molière radius of 2.2 cm make lead tungstate the appropriate material. In addition, the scintillation decay time is in the order of magnitude of the bunch crossing time. Within 25 ns about 80% of the light is emitted [27].

The barrel part of the ECAL (EB) is composed of crystals with a cross-section in η-ϕ of approximately 0.0174×0.0174 and a length of 23 cm which corresponds to 25.8 X_0 and avalanche photodiodes are used as photodetectors[5]. The EB is divided into 36 supermodules which are composed of four modules, each containing between 400 and 500 crystals. It covers pseudorapidities of $|\eta| < 1.479$.

The endcaps (EE) cover the rapidity range between $1.479 < |\eta| < 3.0$ and are composed of identical crystals grouped in mechanical units of 5×5 crystals, the supercrystals. Each crystal has a rear face cross-section of 30×30 mm^2, a front face cross-section of 28.62×28.62 mm^2 and a length of 220 mm (24,7 X_0). They are connected to vacuum phototriodes as photodetectors as they can be operated in regions with high radiation exposure. An overview of the ECAL is given in Figure 2.5.

One of the driving criteria of such a design have been the requirements concerning the spatial and energy resolution, which are particularly needed to discover the decay of the Higgs boson into two photons. Yet, a technical challenge for the operation of the ECAL is the calibration which can be separated in a global component related to the absolute energy scale and a channel-to-channel relative component. The latter is dominated by the scintillation light varying for the different crystals. A first inter-calibration is already performed by operating the supermodules exposed to cosmic rays and is expected to provide a precision of better than 2%. The final precision will be achieved in situ with physics events. Here, mass reconstruction or comparison of charged track

[5]This cross-section corresponds to of 22×22 mm^2 at the front face and 30×30 mm^2 at the rear face in the x-y-plane.

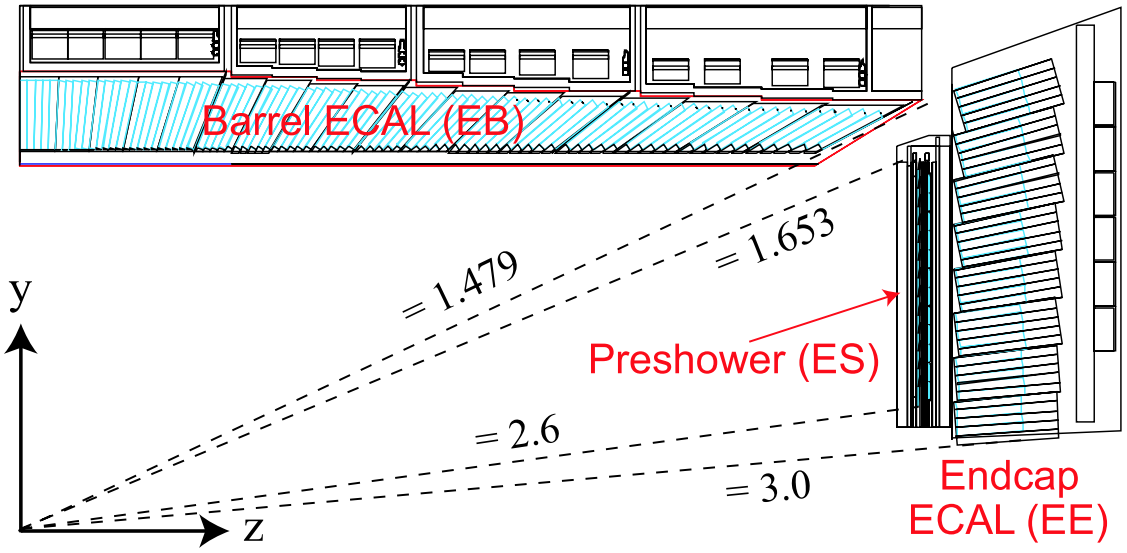

Figure 2.5: *Schematic overview of the electromagnetic calorimeter [18].*

measurements in the tracking system can be used for the calibration of the ECAL.

2.2.4 Hadron Calorimeter

Hadrons will pass the ECAL without losing much of their energy as a result of their larger radiation length. These particles are stopped in the Hadron Calorimeter (HCAL) where they create a hadronic shower. This cascade results from the interaction of the incident particle with the nuclei of the typically heavy material and consists mainly of pions, K mesons, protons, neutrons and fragments of nuclei. As a result of the high penetrating power of hadrons, the hadron calorimeter consists of steel and brass absorber plates interleaved with scintillator tiles and wavelength shifting fibres which bring out the light to hybrid photodiodes. As most of the energy is deposited in the absorber material and only a fraction is measured in the scintillators, the resolution is worse compared to the ECAL.

As the hadron calorimeter is placed within the magnetic field, the design of this detector component is mainly driven by magnetic properties of the material and the requirement to provide a good overall coverage for the measurement of the missing transverse energy. Due to momentum conservation, the vector sum of all transverse momenta must be balanced in an event. The non-vanishing part of this sum is called missing transverse energy and is an indirect measure for neutrinos and exotic particles which do not interact with any detector material. The HCAL consists of four subdetectors which are described below. In addition, a schematic overview of the subsystems is presented in Figure 2.6. More detailed information on this calorimeter can be found in [28].

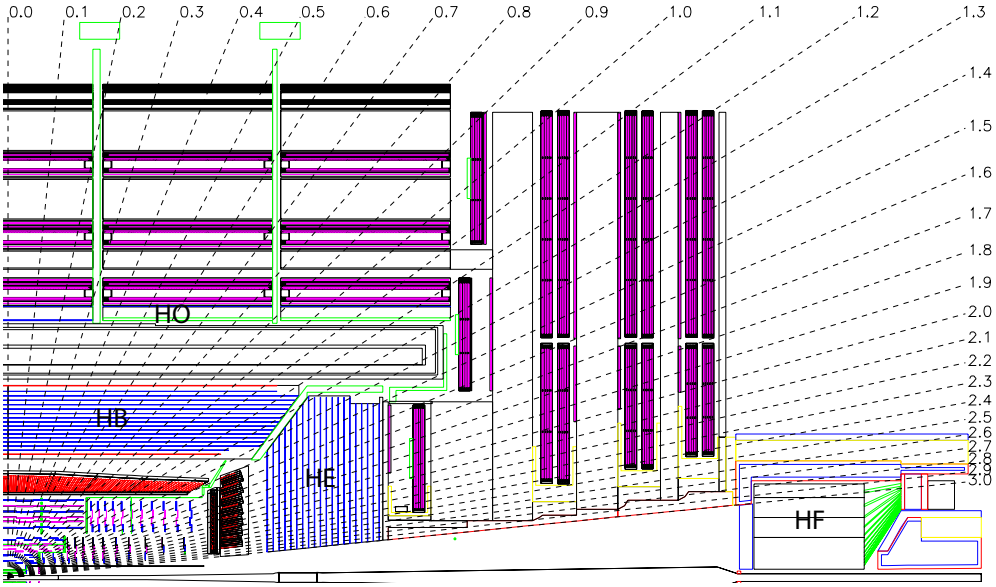

Figure 2.6: *Longitudinal view of the CMS detector with the location of the four subsystems of the hadron calorimeter [24].*

Hadron Barrel Calorimeter

The Hadron Barrel Calorimeter (HB) covers the pseudorapidity range $|\eta| < 1.3$ and consists of 36 identical wedges constructed out of flat brass absorber plates which are aligned parallel to the beam axis. For reasons of structural strength, the innermost and the outermost of the 16 plates of the absorber material are made of stainless steel whereas the remaining plates are made of brass. The latter have a thickness between 50.5 mm and 56.5 mm. As active medium, plastic scintillators with a thickness between 3.7 mm and 9.0 mm are used and divided into 16 η sectors to achieve a segmentation in the η-ϕ-plane of 0.087×0.087.

Hadron Outer Calorimeter

The Hadron Outer Calorimeter (HO) is the only part of the CMS calorimeter system which is located outside the solenoid. Placed in the central barrel between the solenoid and the muon system it covers a range in pseudorapidity of $|\eta| < 1.26$. Using the solenoid coil as additional absorber, this calorimeter part is designed to measure energy leaking out of the HB, which has only a thickness of 6.5 interaction lengths.

Hadron Endcap Calorimeter

The Hadron Endcap Calorimeter (HE) covers the endcap region with a pseudorapidity range of $1.3 < |\eta| < 3.0$. The granularity of this calorimeter is $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ for $|\eta| < 1.6$ and changes to $\Delta\eta \times \Delta\phi \approx 0.17 \times 0.17$ for increasing values of the pseudorapidity.

Hadron Forward Calorimeter

The Hadron Forward Calorimeter (HF) covers pseudorapidities between $3.0 < |\eta| < 5.0$. It is essentially a cylindrical structure with an outer radius of 130 cm consisting of 5 mm thick steel absorber plates. The front face of this cylinder is located at 11.2 m from the interaction point. As its innermost layer has a minimum distance to the beam pipe of only 12.5 cm, the rate of hadrons will be extremely high for this detector part leading to strong requirements concerning radiation hardness. Therefore, quartz fibres were chosen as active medium. As the signal in these fibres is generated by Cherenkov light, this subdetector is mainly sensitive to the electromagnetic component of the shower.

2.2.5 Superconducting Solenoid

The superconducting solenoid is about 12.5 m long and has an inner diameter of 5.9 m. It has been designed to provide a magnetic field of up to 4 T in the inner region of the detector. The flux generated by the coil is returned via a 1.5 m thick saturated iron yoke with a weight of 10000 t comprising 5 wheels and three disks on each site as endcaps. When the system is operated at 4.6 K and full current, a total energy of 2.6 GeV is stored.

2.2.6 Muon System

A precise and robust identification of muons is of central importance to CMS as the name of the detector already indicates. To provide a good momentum resolution, charge identification and trigger capabilities, the muon chambers are interleaved with the magnet return yoke providing a magnetic field of 1.8 T. In addition, the solenoid and additional absorber material ensure that only muons are able to enter the muon system, which consists of three different types of gaseous particle detectors for the identification: Drift Tubes (DT), Cathode Strip Chambers (CSC) and Resistive Parallel Plate Chambers (RPC). The latter are very fast and thus well designed to offer information for the first level trigger whereas the slower DT and CSC sensors yield a precise measurement of the position and thus the momentum of the muon. The layout of the muon system follows the shape of the solenoid and consists of about 25000 m^2 of detection planes.

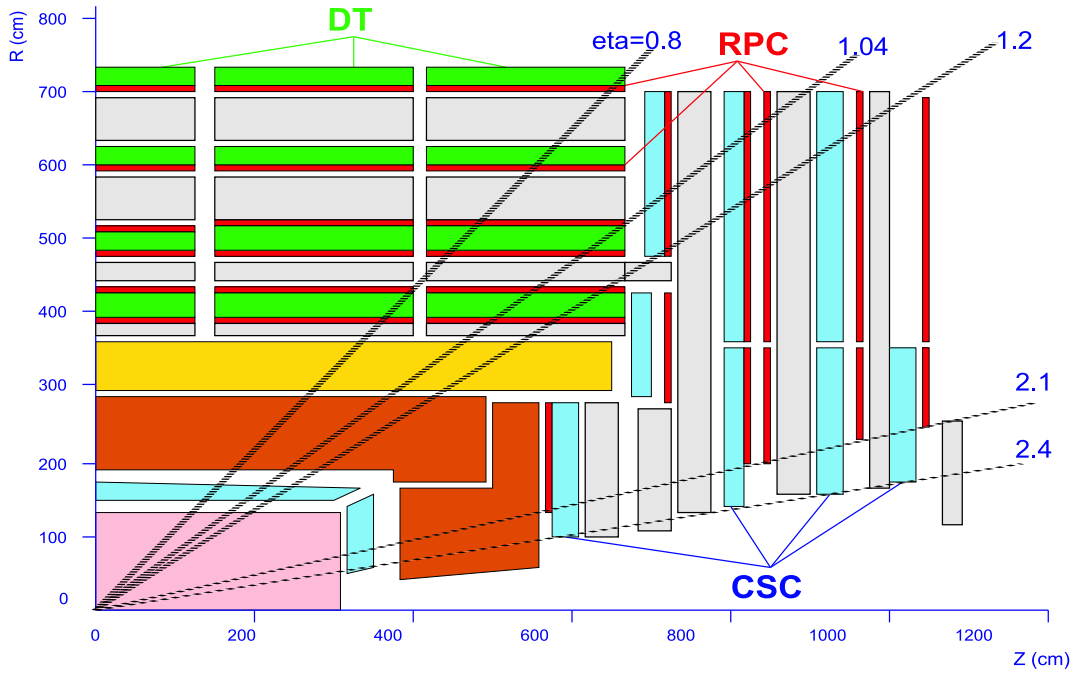

Figure 2.7: *Schematic overview of the muon system of the CMS detector [18].*

In the central barrel part of the detector, DT and RPC detectors are installed and provide two independent and complementary sources of muon information for the pseudorapidity region of $|\eta| < 1.2$. The muon endcap system consist of CSC detectors for a precise position and momentum measurement of muons with a pseudorapidity of $1.2 < |\eta| < 2.4$. The endcaps are completed by RPCs which cover the region up to $|\eta| < 1.6$ in the initial phase of the experiment and will be extended to $|\eta| < 2.1$ later. A scheme of the muon system is drawn in Figure 2.7.

With this layout, a reconstruction efficiency of 90% is achieved for muons with a transverse momentum larger than 100 GeV over the entire range. A detailed overview of the momentum resolution for the muon tracks, charge reconstruction efficiency and the performance of the different detector components is available in [29].

2.2.7 Data Acquisition

At the LHC interaction points, the bunches of the proton beams will cross each other every 25 ns corresponding to a rate of 40 MHz. During such a crossing, 20 interactions are expected at the LHC design luminosity of 10^{34} cm^{-2}s^{-1} resulting in a total number of about 10^9 events per second which have to be filtered and reduced to a storable and processable amount of data. To achieve such a drastic reduction a two level trigger system is used by CMS [30]. This is the start of the physics selection process and the system has to be designed carefully to avoid biases in the selection.

The Level-1 trigger (L1) is very fast and designed to reduce the rate to about 100 kHz. Special programmable hardware processors are searching for signs of an interesting event in each piece of the detector, for example a group of calorimeter cells with a high energy deposition or a dedicated number of muons in the event. While the Level-1 trigger is deciding to accept or reject an event, the high resolution data is kept in the memory pipelines of the front-end electronics.

If accepted, the event is passed to the High Level Trigger system (HLT) which has access to the complete read out data and can therefore combine information from different detector components. The software system of the HLT is implemented on a filter farm of about thousand commercial processors and the processing time per event may reach up to 1 s. The algorithms for this selection will evolve with time and experience. In this step, the rate is reduced to 150 Hz at high luminosity running. With an assumed event size of about 1.5 MB, 225 MB/s of data has to be stored for later offline processing [31]. A schematic overview of the CMS trigger system is presented in Figure 2.8.

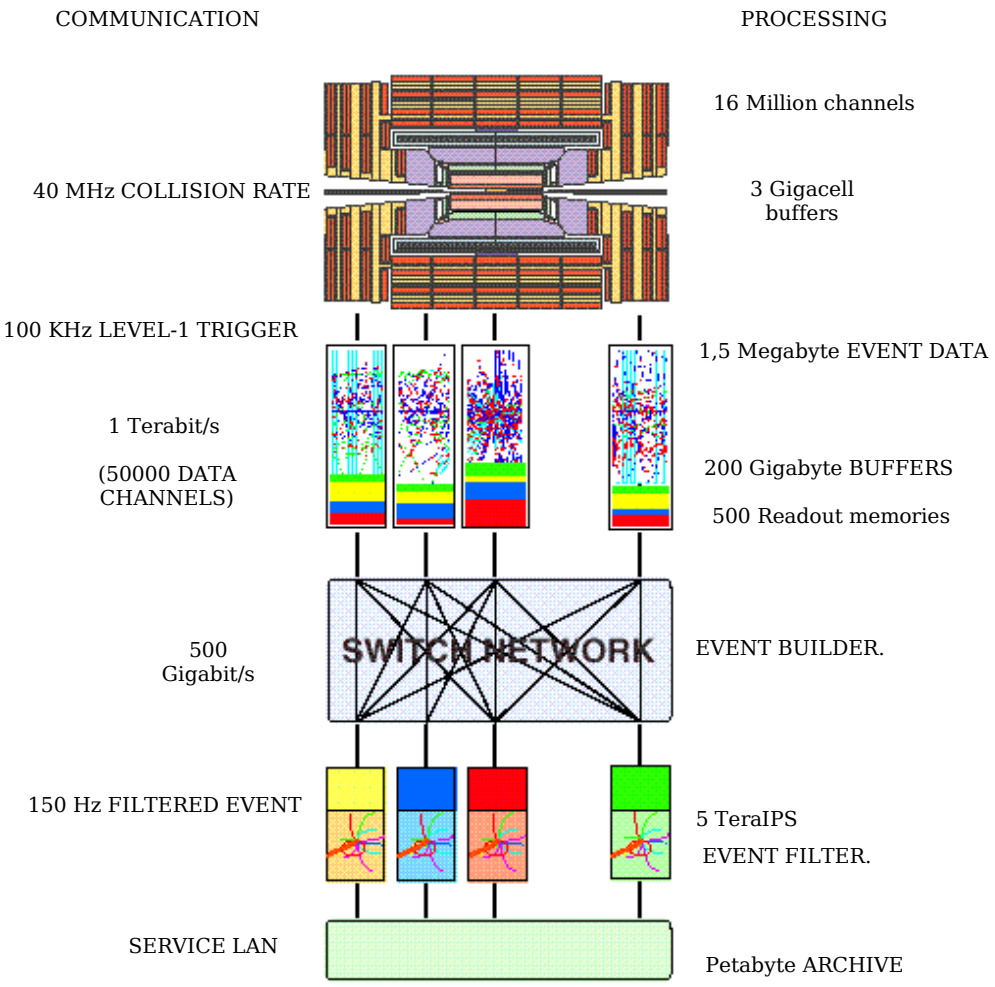

Figure 2.8: *Overview of the CMS trigger and data acquisition system. The raw data is filtered by the Level-1 trigger to a rate of 100 kHz. The High Level Trigger system then reduces the data to $\mathcal{O}(10^2)$ Hz.*

The Computing Infrastructure

Experiments at modern hadron colliders have to deal with data production rates and event sizes not known to date. As a consequence of the high centre-of-mass energy at the LHC, an unprecedented number of $\mathcal{O}(1000)$ final state particles will be produced per proton-proton collision. To achieve a precise reconstruction of these particles up to high momenta requires fine grained detector components, which necessitates about 10 million channels to read out several million times each second.

The four experiments at the Large Hadron Collider, which will start operation this year, will produce about 15 petabytes of data annually. Once the LHC is operating and the experiments have started to record data, a redundant and safe storage has to be guaranteed. In addition, the access to these data for several thousands of physicists working at institutes all over the world has to be provided for the different analyses. In the following, the LHC computing concept is presented, focussing on the requirements of the CMS experiment. This includes an introduction to the basic concept of grid computing including a discussion of the required services. At the end of this chapter, the particularities of the computing infrastructure of the IEKP with respect to the relation to the grid are discussed in detail.

3.1 The LHC Computing Concept

The traditional concept to cope with these challenges would be to centralise all needed computing resources near the experiments. However, due to funding constraints and the fact that the different analysis groups are distributed all over the world, the LHC experiments are using decentralised mass storage and computing resources. The LHC Computing Grid Project (LCG) was approved by the CERN council in 2001 to develop, build and maintain such an infrastructure [32].

Besides the advantages of distributed storage and computing power, there are a number of signifi-

Centre	ALICE	ATLAS	CMS	LHCb
TRIUMF (Canada)		x		
GridKA (Germany)	x	x	x	x
CC-IN2P3 (France)	x	x	x	x
CNAF (Italy)	x	x	x	x
NIKHEF / SARA (Netherlands)	x	x		x
Nordic Data Grid Facility (NDGF)	x	x	x	
ASGC (Taiwan)		x	x	
RAL (UK)	x	x	x	x
BNL (US)		x		
FNAL (US)			x	
PIC (Spain)		x	x	x

Table 3.1: *Listing of the LHC Tier 1 centres with the supported experiments [33].*

cant challenges:

- Network: Adequate levels of network bandwidth between the participating computing centres are required in order to ensure stable data transfers.
- Data storage: For the lifetime of the LHC, a redundant and safe storage must be guaranteed. This includes automated data transfers and checks between the different resources.
- Software: The needed software components have to be installed and maintained on a distributed and heterogeneous hardware infrastructure.
- Accounting: Advanced accounting mechanisms have to be implemented to allow different groups fair access to the resources, based on their quota.

The LHC Computing Grid Technical Design Report [32] summarises the requirements of the different experiments. All resources are connected and accessible via a computing grid, based on a distributed four-tiered hierarchical architecture. For all experiments, the Tier 0 centre at CERN accepts the raw data emerging from the data acquisition systems and a first reconstruction takes place. A copy of these data is distributed to the Tier 1 centres supporting the particular experiment as shown in Table 3.1. The defined tasks for the Tier 1 and Tier 2 centres differs for the four experiments and are therefore only discussed for CMS. A listing of the computing resources requested by the LHC experiments is listed in Table 3.2. The typical unit for the CPU specification is the Million SpecInt 2000 (MSI2k), based on the integer benchmark suite SPECint2000 which is provided by the Standard Performance Evaluation Corporation [34]. This benchmark is used as it has been found to scale well with typical applications of high energy physics. As an indication, a powerful Pentium 4 processor delivers about 1700 SPECint2000. The requested storage is listed in units of petabyte (PB).

	Year	CMS	Atlas	LHCB	Alice	Total
CPU	2007	21.9	9.1	7.8	14.0	52.8
	2008	43.8	50.6	13.0	35.0	142.4
	2009	67.2	82.9	14.5	45.5	210.1
	2010	116.6	138.5	17.9	59.2	332.1
Disk space	2007	4.1	4.8	2.0	5.7	16.6
	2008	13.8	25.4	3.3	14.1	56.6
	2009	23.2	41.2	4.0	18.4	86.8
	2010	34.7	70.5	4.7	23.9	133.9
Mass storage space	2007	5.4	2.6	2.1	4.2	14.3
	2008	23.4	15.2	3.4	10.6	52.6
	2009	41.5	29.2	7.1	13.8	91.6
	2010	59.5	49.1	11.6	17.9	138.1

Table 3.2: *Time profile of the computing resource estimates of all four LHC experiments for the years 2007 till 2010 according to [32]. The units are in MSI2k for CPU and PB for disk and mass storage.*

3.2 The CMS Computing Model

The CMS detector is designed to read out data 40 million times each second. Assuming an event size of about 1.5 megabyte (MB) would lead to 60 terabyte (TB) of data per second. This enormous flow is reduced by the CMS trigger system to a final recording rate of 150 Hz or 225 MB per second (see Figure 3.1).

Therefore, data transfer must happen in real-time with a rate of 225 MB/s from the HLT to the Tier 0 farm where a prompt first reconstruction is performed [35]. This reconstruction may happen with a short delay to allow the usage of updated information on the detector calibration but has to be finished within less than 24 hours. Both, reconstructed (RECO) and raw (RAW) data are archived on the Tier 0 mass storage system and a copy is transferred to a Tier 1. At these centres, a first version of the Analysis Object Data[1] (AOD) is produced, derived from the RECO events as well as additional processing like running dedicated filters (skimming) according to the requests of the physics groups as well as bulk re-processing of the RAW data. The latter is assumed to happen up to three times a year during the initial phase of the LHC and will decrease with time as a result of a better understanding of the detector and the data.

The skimmed datasets are then transferred to the Tier 2 centres which are smaller than the Tier 1 centres but offer in total substantial computing resources. A detailed listing of all Tier 2 centres can be found in [36]. These centres have to offer capacity for analysis, calibration activities and Monte Carlo simulation. Their access to larger data sets and secure storage of new data they produce is provided by the Tier 1 centres. The resource contributions and service levels of the Tier 1 and 2

[1] The Analysis Object Data is a data format within the CMS software framework which is designed to be sufficient for a large set of CMS analyses. It is a proper subset of the RECO data format.

Figure 3.1: *Overview of the CMS online trigger system with the corresponding trigger rate and data flow.*

centres are declared in a Memorandum of Understanding [37] between these centres and CERN. Both categories of grid centres are subject to the CMS accounting. An overview of Tier 0, 1 and 2 computing resources is presented in Table 3.3 and Table 3.4.

Finally, the Tier 3 centres are designed for interactive analysis of local groups but are also connected to the grid. On these and the Tier 2 resources most of the user analyses should be performed. This hierarchical tiered structure is schematically drawn in Figure 3.2.

In addition, CMS requires additional computing services directly at CERN, reflecting a combination of services similar to those provided by the Tier 1 and Tier 2 centres. The CMS-CERN Analysis Facility (CAF) which offers capacity for such services is hosted by the CERN IT Devision. For example, the CAF will provide functionality for short turn-around and latency-critical data processing, including alignment and calibration.

3.3 The Worldwide LHC Computing Grid

In the last years, several computing grids have been developed for the usage in science, e. g. the Scandinavian *NorduGrid* or the American *Grid3* and *Open Science Grid*. Also other fields of science have noticed the advantages of grid technologies such that most of the infrastructures have to deal with different needs of multiple fields like particle physics, biology or medicine.

	2008			2009			2010		
	Req	06 Req	Pledge	Req	06 Req	Pledge	Req	06 Req	Pledge
T0 CPU	5.3	3.9	3.9	9.8	6.1	5.4	19.1	10.6	7.6
CAF CPU	2.1	3.8	3.8	3.9	5.8	5.0	9.7	11.5	7.4
CERN total	**7.3**	7.7	7.7	**13.7**	11.9	10.4	**28.8**	22.1	15.0
T0 disk	0.4	0.3	0.3	0.2	0.3	0.4	0.4	0.5	5.5
CAF disk	1.8	1.3	1.3	2.3	2.0	2.3	3.4	3.3	3.8
CERN total	**2.2**	1.6	1.6	**2.5**	2.3	2.7	**3.8**	3.8	9.3
T0 tape	4.4	3.6	3.6	7.3	6.7	7.7	11.1	10.9	13.8
CAF tape	0.9	1.5	1.5	2.0	2.7	3.4	3.2	4.1	5.1
CERN total	**5.3**	5.1	5.1	**9.3**	9.4	11.1	**14.3**	15.0	18.9
T1 CPU	**9.6**	12.4	11.3	**16.3**	16.9	15.2	**34.8**	36.9	30.0
T1 disk	**7.2**	5.6	5.3	**9.7**	8.5	7.7	**15.4**	13.7	11.3
T1 tape	**9.8**	13.1	9.2	**15.0**	23.5	16.7	**23.2**	36.6	25.5
T2 CPU	**13.4**	15.2	16.8	**28.1**	25.6	25.3	**76.6**	45.2	43.0
T2 disk	**5.1**	4.2	4.2	**5.7**	8.4	7.1	**7.6**	13.3	10.5

Table 3.3: *Overview of the CMS computing resources for the years 2008 - 2010. The three columns per year show the request from 2007 (Req), the request from 2006 (06 Req) as well as the pledge (state 2007) [31]. The units are in MSI2k for CPU and PB for disk and mass storage.*

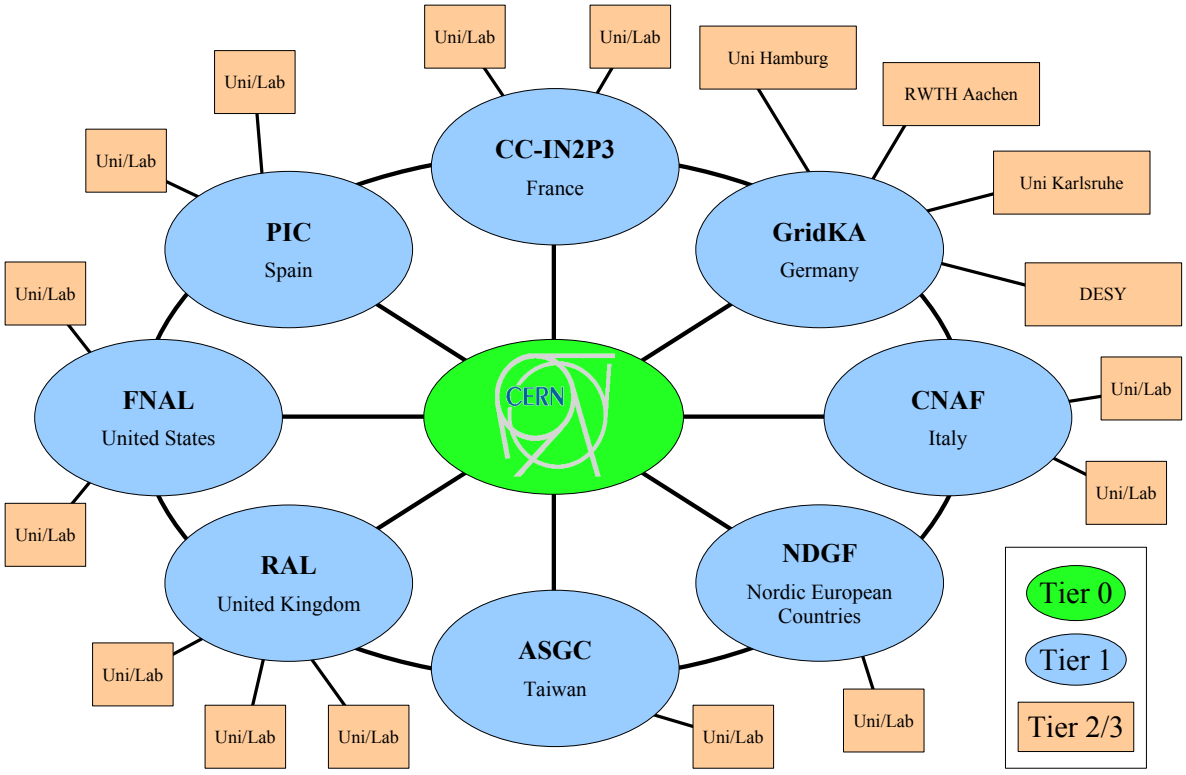

Figure 3.2: *Schematic overview of the multi-tier hierarchical structure of the LHC Computing Grid.*

	2011		2012	
	Request	Pledge	Request	Pledge
T0 CPU	19.1	7.6	19.1	
CAF CPU	13.4	10.6	17.0	
CERN total	32.5	18.2	36.1	
T0 disk	0.4	5.5	0.4	
CAF disk	4.6	5.3	5.7	
CERN total	5.0	10.8	6.1	
T0 tape	14.9	21.4	18.7	
CAF tape	4.8	7.1	7.1	
CERN total	19.7	28.5	25.8	
T1 CPU	42.2		50.3	
T1 disk	19.4		23.5	
T1 tape	32.3		41.2	
T2 CPU	105.7		134.7	
T2 disk	11.0		14.4	

Table 3.4: *Overview of the CMS computing resources for the years 2011 and 2012. The two columns per year show the request and pledge for the different Tier centres (state 2007) [31]. The units are in MSI2k for CPU and PB for disk and mass storage.*

The grid infrastructure used by the four LHC experiments is based on the grid middleware developed by the Enabling Grids for E-sciencE (EGEE) project [38]. Having started with a middleware from its predecessor, the European DataGrid, the LCG middleware stack has been developed and used in the early phase of the EGEE project. While gaining experience in running a grid in production quality, most of the middleware components have been re-engineered to build the new middleware framework gLite, which is now deployed and used in the production environment in its release 3.1.

In the following, the basic concepts of this middleware are discussed, including the mechanisms of authentication and authorisation as well as the interaction of the different services involved.

3.3.1 Concept of Identity Management on the Grid

The users of a grid are members of Virtual Organisations (VO). These abstract entities group users and resources in a common administrative domain. In the Worldwide LHC Computing Grid (WLCG), the VOs are associated with the different experiments.

Authentication and Authorisation

The authentication on the grid is based on the Public Key Infrastructure (PKI) realised by certificates. To obtain such a certificate, the user generates a pair of public and private keys according to the X.509 standard [39]. To ensure, that the information like name and organisation of the requester are authentic, the public key must be signed by an independent institution, the Certification Authority (CA), according to predefined policies [40]. In Germany, Certification Authorities are hosted at the DFN [41] and GridKA [42]. Once the identity of the user and the corresponding information of the certificate are checked and approved, the signed public key is sent back to the user, now including information concerning the CA as well as its digital signature. The certificate is now the equivalent to a passport for the grid.

The certificate is typically valid for one year and should be protected by a passphrase. For the work on the grid, the user creates a proxy certificate with a dedicated lifetime, which is based on the certificate and is no longer protected by a password.

Now being able to authenticate, the user can request the membership of a Virtual Organisation. Once the request is accepted, the user is authorised to use the grid resources of the VO. The access privileges to the resources depend on the task the user is performing on the grid and differ between personal physics analysis, official central software installation or other duties for the Virtual Organisation. Therefore, a fine-grained structure with different roles a user can play is needed.

VO Server

With the first versions of the middleware, the authorisation of the different users at the grid sites was realised using a VO Server, acting as registry office of a Virtual Organisation. On this LDAP[2] server a list of all users belonging to a VO and their role and permissions are stored. The VO server distinguishes only two kinds of members: normal grid users and the software managers. Originally, the role of a user was a static entry and it was therefore not possible to switch between different roles. To enable such functionality, a new concept of user authorisation is now used and discussed below.

VOMS server

The Virtual Organisation Management Service (VOMS) server offers a fine grained definition of user roles [43]. To enable such functionality, the role of the user is no longer stored statically on a central server but added to the corresponding proxy certificate. To get such an extended certificate, the user specifies the desired role while creating the proxy certificate. During this process, the

[2]LDAP is an acronym for Lightweight Directory Access Protocol. It is a lightweight client-server protocol for accessing directory services over the TCP/IP protocol or other connection oriented transfer services.

VOMS server is contacted automatically and signs the proxy certificate, if the user is authorised for the requested role. This allows the determination of the user's role according to the information in the proxy while working on the grid.

Some of the roles, available for members of the CMS VO are for example:

- Basic membership (default for all members of the VO)
- Software installation and administration
- Dataset and Tier 1 and 2 administration
- Monte Carlo production
- Analysis subgroups, e. g. Higgs or Standard Model physics
- Regional subgroups, e. g. Germany, Italy or the USA.

3.3.2 Grid-wide Services

One can classify the different services required to operate a grid in two groups: the grid-wide and site-wide services. The latter are services each site has to provide in order to offer computing power and disk space to a Virtual Organisation. In contrast the grid-wide services relate to the infrastructure of a VO and are therefore provided centrally by the Virtual Organisation itself. Both, the VO and the VOMS server are infrastructural services and belong to the second group of services. An overview of the different services and their interactions as described in the following section, is displayed in Figure 3.3.

The Information System

The Grid Information System provides detailed information about the different services and resources available on the grid. It is organised in a three-level hierarchical structure based on the Berkeley Database Information Index (BDII) as a fundamental building block [44]. The BDII is constructed of two or more standard LDAP databases which are populated by an update process. This concept allows to use one database to serve data while the other is refreshing [45].

In this architecture, the resource level BDIIs represent the lowest layer and provide information concerning the different resources like number of CPUs available, number of running and waiting jobs or the free and occupied storage space. At each grid site, one site level BDII is running to aggregate all information from the connected resources of the specific site. Finally, the top level BDII is collecting the information from all site level BDIIs to gather information on all grid services and resources available. In order to provide a fault tolerant system and to enable load

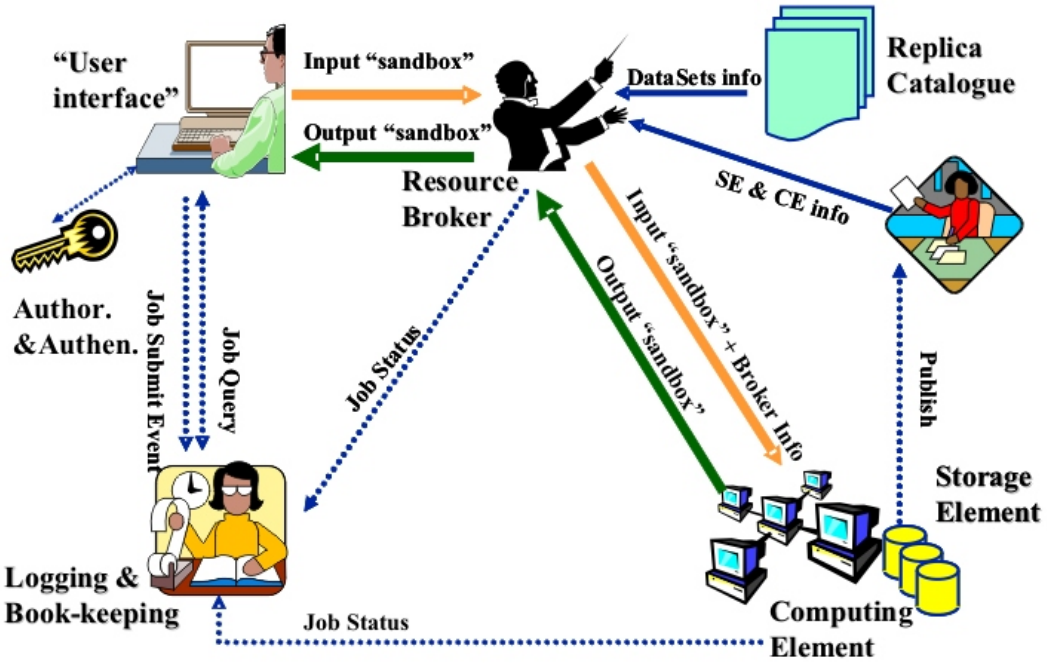

Figure 3.3: *Schematic overview of the interaction between the different grid services which are necessary for the execution of a job on the grid: The user, authenticated by the proxy, submits the request for a job with all requirements to the Resource Broker. Having compared these specification with the resources available at the different sites (SE & CE info, DataSets info), the Resource Broker submits the job to a matching Computing Element. From now on, the Logging & Book-keeping service provides information on the job status. When the job has finished, its output will be transferred to the Resource Broker and will be stored there until the user retrieves it.*

balancing, multiple instances of the top level BDII are offered. An overview of the information system is schematically shown in Figure 3.4.

The information is provided following the Grid Laboratory Uniform Environment (GLUE) schema that was defined by a collaboration between European and US grid projects [46]. This scheme offers a standardised description of a grid system allowing the exchange of information between resources and services on the one hand and user and external services on the other hand.

LCG File Catalogue

The LCG File Catalogue (LFC) is responsible for the mapping between logical and physical file names. Once a file is copied to the grid and registered to this catalogue it gets a unique identifier,

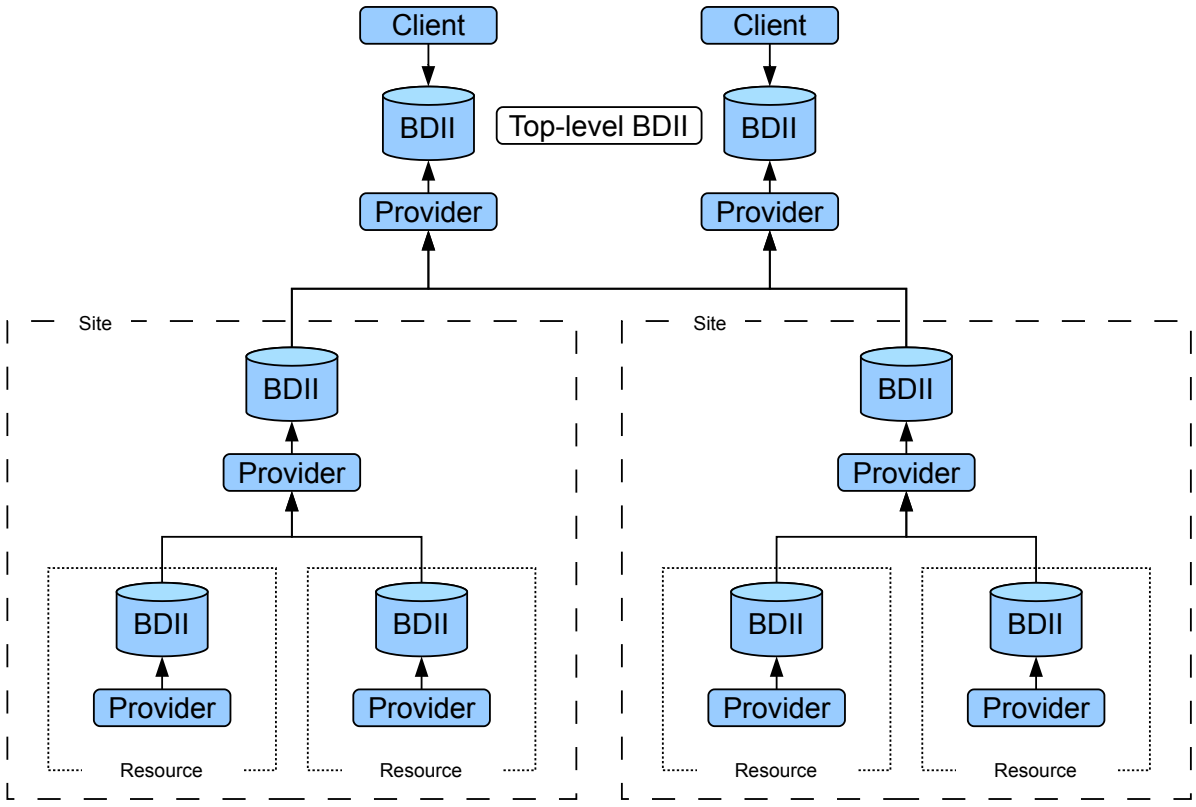

Figure 3.4: *Schematic overview of the information system. It is organised in a three-level hierarchical structure based on the Berkeley Database Information Index (BDII) as fundamental building block. The site level BDII aggregates information from the connected resources. The top level BDII collects information from all sites and holds information related to all grid services and resources.*

the Globally Unique Identifiers (GUID), which is mapped to the physical location and file name. As this identifier consists of arbitrary numbers and characters not easy to remember, a Logical File Name (LFN) can be assigned. In its database, the LFC stores the mapping between LFN, GUID and the physical file name on the dedicated grid storage resource.

It depends on the deployment models of the different Virtual Organisations whether the LFC is installed centrally or locally. Authentication and authorisation are based on the concepts of grid certificates and VOMS. The catalogue provides a hierarchical namespace and namespace operations, Access Control Lists (Unix Permissions and POSIX Access Control Lists) and a user exposed transaction Application Programming Interface (API) [47].

System attributes of the files like the time of creation or last access, file size and checksum are stored as attributes to the LFN. User-defined metadata are restricted to one field as such metadata should be stored in a separate metadata catalogue. Additional LFNs per GUID are realised as symbolic links to the primary LFN. A schematic overview of the LFC architecture is drawn in Figure 3.5.

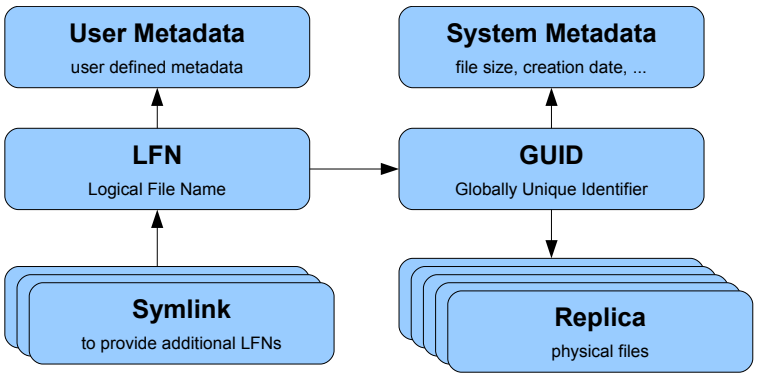

Figure 3.5: *Schematic overview of the LCG file catalogue.*

Resource Broker

The Resource Broker (RB) is the service which distributes incoming user job requests to the most appropriate resource according to the information obtained from the Information System. In gLite, this functionality is provided by the Workload Management System (WMS).

In order to specify job requirements, the Job Description Language (JDL) has been developed. Such specifications are for example the name of the executable to be executed including needed parameters, the files which should be moved to and from the site and any requirements concerning the hardware or datasets available at the site.

The Network Server (NS) is a generic network daemon which provides support for the job control. It is responsible for accepting incoming requests from the user like job submission or removal and passes the request, if valid, to the Workload Management System.

Once the job is submitted, the match-making process on the WMS queries the Information System and selects only resources fully matching the requirements of the job. The job is finally send to the resources with the highest rank, a quantity derived from the Information System expressing the "goodness" of a resource. In the current implementation, the rank is typically a function of the numbers of running and queued jobs [48].

The Logging and Book-keeping service (LB) tracks the jobs managed by the WMS. It gathers events from various WMS components in a reliable way and processes them to a single quantity, the status of the job. The WMS records the job's status and retrieves its output once finished successfully. Below, the different possible job states including their meaning is summarised:

- **Submitted:** The job is submitted by the user but not yet transferred for processing to the Network Server.
- **Waiting:** The job has been accepted by NS and is waiting for the resource allocation by the WMS. In cases where no matching resources are identified, the job remains in this state.

- **Ready:** The job has been processed by the WMS but has not been transferred to the resource.
- **Scheduled:** The job has been successfully transferred to the site and is now waiting in the local batch queue.
- **Running:** The job is running on the local resource.
- **Done:** The job has finished or is considered to be in a terminal state.
- **Aborted:** The job processing was aborted by WMS. Reasons for this might be that the job has been waiting in the Workload Manager queue or on the site for too long, the number of possible resubmissions of the job has been exceeded, over-use of quotas or the user credentials are expired.
- **Cancelled:** The job has been successfully cancelled on user request.
- **Cleared:** The output of the job was transferred to the user or removed due to a timeout.

The current implementation of the WMS supports the submission of 30,000 jobs per day. Nevertheless, the different VOs have multiple Workload Management Systems for load balancing. It is also possible to support different VOs on one WMS.

3.3.3 Site-wide Services

User Interface

The User Inteface (UI) is the access point for the user to the grid. This middleware component provides several programs, e. g. for submitting a job to the grid, getting information on its status or retrieving the output of the job. The User Interface is also the portal for file access. The user can copy and register files as well as replicate existing files to other destinations. To be able to work in the grid, a valid proxy certificate is needed.

Computing Element and Worker Nodes

The Computing Element (CE) is the interface to the local batch system of a site. It is not - as the name suggests - the place where jobs are computed; these are the Worker Nodes (WN) of the local batch system. The creation of this layer of abstraction between the computing resources and the user is necessary due to the fact that different sites may use different batch queuing systems. The CE offers standardised access to the local resources and provides the required information by running a resource level BDII.

A new gLite-flavoured CE is available since gLite 3.0 supporting the VOMS extension of the proxy certificate. However, the current release is not recommended to be used in a production environment.

Therefore, the usage of the LCG-flavoured Computing Element is recommended but this is not able to interpret the VOMS extension of the proxy certificate and authentication and authorisation of a user are still based on the VO Server concept. In the current implementation, it is only distinguished between few dedicated roles and the mapping between them and the different user certificates is static. Once a job is submitted to the CE, the middleware software looks up in the grid mapfile[3] to find out whether the user is authorised by a VO to run a job on this site. Furthermore, this file contains the information to which local account the user has to be mapped. Having successfully checked the authorisation of the job, the CE submits it via the local batch system to the Worker Node where the job is executed. The Computing Element copies the input data of the job (Input Sandbox) to a dedicated directory on the Worker Node. In addition, the Worker Node has access to files stored on the grid and the web. To access these files, several protocols and programs are available on the Worker Nodes like gridftp, http and wget. After the job has finished, the CE sends the job output (Output Sandbox) back to the Workload Management System.

Storage Element

Like the Computing Element, the Storage Element (SE) provides uniform access to a sites services, in this case data storage resources like disk servers, large disk arrays or tape-based Mass Storage Systems. To provide access to the local resources, different data access protocols and interfaces are provided. It is important to emphasize that data storage on the grid must be considered as read-only. To change a file, it must be deleted and has to be replaced with the new file.

For the file transfer, the GSIFTP protocol is used, offering the functionalities of the File Transfer Protocol (FTP), but with additional support for Grid Security Infrastructure (GSI) [49]. For direct remote access to files via the SE, the Remote File Input/Output protocol (RFIO) and the GSI dCache Access Protocol (gsidcap) are available in gLite 3.1 [50].

The Storage Resource Manager (SRM) provides capabilities like transparent file migration from disk to tape, file pinning, space reservation and many more. However, in the current implementation of the middleware, the supported SEs may require different SRM versions or do not support SRM at all.

Currently, the following SE flavours are supported by the gLite middleware:

- The **Classic SE** consists of a GridFTP server and an insecure RFIO daemon enabling access to a physical single disk or disk array. The support for this SE flavour will end in near future.

[3] /etc/grid-security/grid-mapfile

Type	Resources	File transfer	File I/O	SRM
Classic SE	Disk server	GSIFTP	insecure RFIO	No
CASTOR	MSS	GSIFTP	insecure RFIO	Yes
dCache	Disk pool/MSS	GSIFTP	gsidcap	Yes
DPM	Disk pool	GSIFTP	secure RFIO	Yes

Table 3.5: *Summary of the Storage Elements supported in gLite 3.1.*

- **CASTOR** provides a disk buffer frontend to a tape Mass Storage System. A virtual file system (namespace) shields the user from the complexities of the underlying setup and file migration between disk and tape is managed automatically.
- **dCache** consists of a server, providing a single point of access to the SE, and multiple pool nodes. The latter allow to store and access huge amounts of data on a large number of heterogeneous nodes under a single virtual file system tree [51].
- **LCG Disk Pool Manager (DPM)** is a lightweight disk pool manager, suitable to offer up to 10 TB of total disk space and therefore only usable for smaller sites. It also provides a virtual file system like dCache and CASTOR in order to hide the underlying complexity.

A summary of the supported SEs is presented in Table 3.5.

Monitoring System Collector Server

The Monitoring System Collector Server (MonBox) publishes information about the jobs, which have been executed at a site. This information is provided to the VO and can be used for accounting. The publication follows the R-GMA standard which is an implementation of the Grid Monitoring Architecture [52] proposed by the Global Grid Forum [53].

3.3.4 CMS specific Grid Services

Beside the basic grid functionality provided by the gLite middleware, additional services specific for the different Virtual Organisations are needed. In the following a brief overview of the additional services developed and used by the CMS collaboration is given.

Dataset Book-keeping System

The Dataset Book-keeping System (DBS) provides a standardised way of describing the event data available [54] by relating datasets and event collections. In the CMS computing model, a dataset

consists of one or more file blocks, each holding a large number of files. The DBS API enables the user to query the catalogue for the following information:

- Name and physics content of the simulated or recorded dataset.
- Software version used for the processing of the data.
- History of which event collections are derived from which other event collections.
- Rough estimate of integrated luminosity (the exact luminosity is stored outside the DBS).
- Information on the conditions / calibrations used for the processing.
- Data quality flags.

For example, direct access to the information available in the DBS is provided by a web page[4].

Data Location Service

The DBS described above only holds information on the registered datasets but not on the location where replicas of these datasets are located on the grid. To find the locations where the dataset of choice is available, the Data Location Service (DLS) has been developed [55]. This system keeps the mapping of the file blocks to the different Storage Elements where they are located.

Local File Catalogue

However, the DLS does not contain any information on the physical location of a file at the dedicated site as it only deals with logical file names. The Local File Catalogue (LFC) is a service running at each site which holds the mapping between logical and physical file name [35]. As all CMS applications are able to deal with logical file names, the LFC is queried once the application needs access to a physical file.

Physics Experiment Data Export

The software of the Physics Experiment Data Export (PhEDEx) project arranges the data placement and the file transfers for CMS [56]. Once the replication or transfer of a file block is requested in the Transfer Management Database, agents running at the involved sites are managing the file transfer. This includes the book-keeping of already transferred files and checksumming in order to ensure data set consistency and integrity.

[4] http://cmsdbs.cern.ch/DBS2_discovery

Submission and Monitoring Tools for Grid Jobs

To enable an easy and efficient way of creating, submitting and monitoring CMS jobs on the grid, several tools are available. The CMS Remote Analysis Builder (CRAB) is the tool for grid-based analyses, recommended by the collaboration [57]. Specifying the necessary parameters, CRAB collects all related files from the local working environment of the CMS software, creates the JDL and submits the job to the grid. It allows for automated job monitoring and re-submission if a job failed. Unfortunately, CRAB is designed to handle only jobs based on the CMS software framework and does not support other grid jobs.

Therefore, an additional submission and monitoring program has been developed at the IEKP. Grid-control is designed in a modular way for automated submission and monitoring of any grid job. In addition, similar functionality like CRAB is offered to the user including the interface to DBS for data access. The development and maintenance of grid-control is still ongoing.

A more detailed description of grid-control is given in Appendix C.3 including an example configuration as used for the analysis presented in this thesis.

3.4 The Karlsruhe Analysis Infrastructure

The computing infrastructure at the IEKP is shared between the different working groups of the two high energy physics experiments CDF [15] and CMS. A third group is involved in the AMS [58] experiment to be installed at the International Space Station. Up to now, these three groups are using the local computing cluster of the institute for offline analysis. This cluster provides the following resources:

- About 30 computing nodes, based on x86-32 or x86-64 architectures
- 5 file servers offering a total capacity of about 33 TB
- 12 portal machines for software development and local job submission
- Batch system and scheduling: Maui and Torque [59]

Beside the local usage, this cluster has been fully integrated into the WLCG as a Tier 3 centre. Therefore, also a Computing Element, a Classic Storage Element as well as a Monitoring System Collector Server have been deployed. This integration allows the sharing of idle resources. In addition, the SE provides direct access from the grid to the local storage resources for users of the IEKP. This is needed as the size of the OutputSandbox for a grid job is limited and larger files have to be transferred directly to a SE at the end of a grid job. For reasons of stability and maintenance, it is recommended to install the different services on isolated machines. A detailed

description of this local computing infrastructure and its integration into the WLCG is available in [60]. Nevertheless, the administration of the grid infrastructure including the deployment of new releases has shown that the production environment should be extended by a second installation to test release changes. As the different services do not fully load the most recent server machines, old hardware is used for the deployment of the services, which has dramatically increased the administrational effort. Therefore, the consolidation of this installation was investigated during the scope of this thesis. The basic concepts of consolidation as well as the final realisation are presented in detail in Chapter 4.

Since beginning of 2008, the local resources have been extended by an additional computing cluster. It is designed and maintained centrally by the computing centre of the University of Karlsruhe and the funding is shared between eight different departments of the university. Therefore, this cluster is called Institute Cluster (IC) from here on. It provides a substantial increase in computing power and storage capacity for offline analysis of the LHC data which are expected to be available end of this year. In the following, the key attributes of the IC cluster are presented:

Hardware

- Number of worker nodes: 200
- Peak Performance: 17.5 Tera Flops
- CPUs per node: 2×Intel Quadcore 2.66 GHz Xeons [61]
- Memory: 2 GB RAM per core, in total 16 GB per node
- Disk space per node: 750 GB
- Storage: 350 TB

Software

- Operating system on the worker nodes: SuSE Linux Enterprise Server 5 [62] with Xen enabled kernel
- Cluster file system: Lustre [63]
- Batch system: JMS (Job Management System) based on SLURM [5]

In addition to the computing nodes and the file servers, six portal machines are deployed for the user groups. Their hardware configuration is similar to the worker nodes except for the memory

[5] Simple Linux Utility for Resource Management [64]

which is doubled. As the IEKP contribution to the funding of this infrastructure is about one third of the total costs, a similar amount of computing and storage resources can be used by the IEKP members. This also includes the exclusive usage of two of the portal machines, offering capacity for code development and job submission to the user.

In consideration of the notably larger amount of resources available at the IC and the administrational effort of grid installations, it has been decided to reorganise the WLCG resources in Karlsruhe. The IC will be included as the Karlsruhe Tier 3 centre in the WLCG to share idle computing resources with the collaboration. The local batch system of the institute will no longer be connected to the grid. Nevertheless, the access to local storage resources of the IEKP besides the IC is realised by a Classic SE.

Integrating the Institute Cluster into the grid is quite challenging. In the current implementation, the gLite middleware requires a special operating system, ScientificLinux [65], a Linux distribution mainly maintained by the FermiLab [14] and CERN computing division and largely used in science. It is based on a recompiled RedHat Enterprise Server [66]. However, it has been decided to use the SuSE Linux Enterprise Server as operating system on the Worker Nodes as the computing centre requires an OS with professional support. In addition, encapsulation of grid jobs has been requested due to security issues. As users from the authorised Virtual Organisations are able to access this cluster from all over the world, the meaning of the Linux permission "world readable" really means that such files or directories are accessible from all over the world. Beside this aspect of privacy, the risk of misbehaving users increases. For both areas of concern, the concept of operating system virtualisation offers solutions which have been investigated and developed in the scope of this thesis and are presented in detail in Chapter 4.

3.5 Accounting and Billing

In the current implementation of the gLite middleware, jobs of a dedicated Virtual Organisation are distributed to all participating sites. If all available resources are fully loaded, the jobs are scheduled at the site it has been sent to until a free slot is available.

Beside the Tier 0, Tier 1 and Tier 2 centres, there is a substantial number of smaller computing centres participating in the WLCG without being bound to contracts with the corresponding Virtual Organisations. Currently these Tier 3 centres offer temporarily free resources to other grid members, hoping that they will be able to use a comparable amount of grid resources later in time in case of need. But up to now, this cannot be guaranteed.

The implementation of a billing method could help to ensure sites that they will get back a comparable amount of computing resources to what they have offered once they need it. Having such a mechanism in hands could motivate more institutes of different disciplines to cooperate using grid technologies and really profit from sharing their resources in order to optimise the load of the

different installations including a better handling of peak loads.

Billing the Grid, a project involving three institutes of different departments of the University of Karlsruhe, namely institutes of economics, informatics and particle physics, have started to discuss, design and implement a prototype of a billing mechanism which addresses this issue [67]. The goal of the project is to exemplarily identify the requirements of billing procedures in the WLCG grid context and to design and implement a prototype installation based on the gLite middleware framework.

To establish such an accounting and billing mechanism, the demands of particle physics usage with respect to billing scenarios have been investigated in detail. The key points are that it should be transparent to the user and that there is only a minimal overhead due to the new concept compared to the current situation. It is mandatory that the grid money has no exchange rate to existing currencies and that international privacy policies are taken into account. The model should ensure that computing power which has been offered to an external grid user enables the members of the institute to get a comparable amount of computing time at another site. In addition, the billing concept should be used to organise the usage of the WLCG tiered architecture within a collaboration.

All this led to the proposal that each institute has a certain amount of tokens, reflecting the grid currency. The account for the tokens of an institute is managed by a local site manager. Before a user sends a job to a site, the Workload Management System is queried for sites matching the job requirements. These sites are then contacted by an agent running on the users User Interface for the price of a CPU interval. Besides the price, information of a reputation database can be taken into account to weight the different offers. This database provides a history of the stability and reliability of the different sites. Based on this information, the user can design criteria for the final job submission.

Having submitted the job to the remote site of choice, the equivalent of tokens to the maximum time, the job may stay in the queue is blocked on the account of the users institute. When the job has finished, the accurate amount of tokens corresponding to the time the job has been executed on the Worker Node is debited from the account.

Two different scenarios are possible:

- The user sends a job to the local site:
  This is the simplest case. The tokens are transferred to the same account they have been debited from. One may regard this as a neutral cash flow. There is no difference to the current concept.
- The user sends a job to a remote site:
  Now, the tokens for the job executed at the remote site are transferred to the account of the institute providing the resource.

The big advantage of such a concept is that it intrinsically realises the concept of load balancing

over many sites and periods in time. Therefore, it represents an added value for the participating sites and may attract more institutes to share their resources using grid based environments. In the current phase, the basic concept is under development and a prototype implementation is prepared. It is planned that the billing mechanism is provided as external package to gLite and that no adaptations of the middleware have to be performed. Therefore, it should be possible to extend the concept to other grid middleware components on the market.

Virtualisation

It the 1960s, IBM offered a mainframe being able to run multiple operating systems in parallel and so to cope with the challenge of full CPU load [68]. While one instance is waiting for data, another can process data which is already available. This can be seen as the first implementation of operating system virtualisation.

In the year 1978, Intel presented the first CPU with the x86 instruction set and the success story of this architecture started. The trend changed from large mainframes shared by many users to one PC per user and operating system virtualisation was no longer needed on systems based on this architecture, which is commonly used in high energy physics.

However, with recent developments in the field of x86 based computing, the problem of not fully loaded resources re-appeared and therewith the need for operating system virtualisation. Unfortunately, the original x86 instruction set does not support multiple operating systems on the same hardware and several techniques based on additional software layers and modifications of the x86 processors have been developed to overcome this limitation.

After a short introduction to the x86 architecture with respect to its limitations concerning virtualisation, the different techniques to enable such functionality are presented. This section is followed by a detailed discussion of use cases of virtualisation techniques in high energy physics experiments, which have been investigated and developed in the scope of this thesis.

4.1 The x86 Architecture

In contrast to the mainframe architecture, the x86 based operating systems are designed to run directly on the hardware and are therefore the only instance needing special privileges. The x86 architecture distinguishes between four different levels of privileges for the access to the hardware

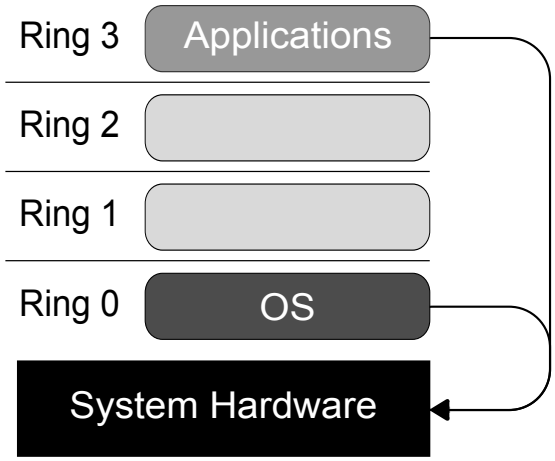

Figure 4.1: *Schematic overview of the x86 privilege level architecture. The operating system is executed in the privileged Ring 0 whereas the user applications run in Ring 3.*

resources of the computer, named Ring 0, 1, 2 and 3 [69].

The Ring 0 allows the execution of privileged instructions for direct access to the hardware layer. Thus, the operating system executes in this level. The user level applications are usually restricted to the unprivileged Ring 3 as schematically drawn in Figure 4.1. It is obvious that a parallel execution of several operating systems is not easily possible on such architectures. To enable this, either the required functionality is provided by an additional software layer or the x86 architecture itself has to be extended. Both approaches are realised and described in detail below.

4.2 Virtualisation of Operating Systems

Virtualisation of operating systems concerns the execution of multiple, isolated and independent operating systems on one single host computer at the same time. The different concepts of such virtualisation are discussed in the following based on [70] and [71].

Hardware Emulation

The concept of hardware emulation realises virtualisation by emulation the required hardware components for the guest operating system. This includes the processor, memory, disks, display, network devices, BIOS and common hardware peripherals. For example, Bochs [72] is able to emulate 386, 486, Intel Pentium, PentiumII, PentiumIII and Pentium4 or x86-64 CPUs including optional MMX, SSEx and 3DNow! instructions. The list of different operating systems, Bochs

is able to execute contains Linux, DOS and the whole family of Microsoft Windows. Another product based on this approach is QEMU [73].

The big advantage of the hardware emulation is the large variety of possible system configurations for the guest operating system. One of the key features is the complete decoupling of the guest operating system and its configuration from the hardware set up on the host system. The guest system can therefore be investigated from the host operating system all the time, even in the case of a system crash. Due to both features, hardware emulation is predestined for the development or testing of operating systems. However, as all hardware components including the processor, the memory and the hard disk are emulated, these products do not reach high emulation speed, which decreases the performance of the guest operating system significantly.

Operating System Container

The parallel execution of several different instances of one OS can as well be achieved using the concept of containers. Only one operating system kernel is running and the different containers are a subgroup of the host OS. To achieve isolation, the internal data structure of the kernel is replicated and a context is defined for each container to offer a full runtime environment based on the host OS including the functionality required by the dedicated applications.

On the one hand, this tight integration into the host operating system eases administration and bears only a small performance overhead. On the other hand, it is limited to one base kernel for all instances. This implies that it is not possible to load different device drivers in the containers. For instance, the commercial product Virtuozzo [74] and the related open source project OpenVZ [75] are following such an approach.

Hardware Virtualisation

In contrast to the emulation, the concept of hardware virtualisation does not emulate any arbitrary configuration for the guest operating system. Instead of holding several different components for each class of hardware for the guest operating system, only a small number of virtual components are offered, independent of the actual configuration of the host system. This allows, like for the hardware emulation, to execute one virtual machine on different host systems with different configurations but the performance is increased compared to full emulation. Products which offer such a kind of virtualisation are the Virtual PC [76] from Microsoft and the products of VMware [77] for which the concept of hardware virtualisation is exemplarily discussed.

The VMware products allow to prepare a special container including virtual CPUs, memory, hard disk, network interfaces, USB ports and other common hardware components. The operating system of the guest is executed in the unprivileged program space and a software layer, called the hypervisor or Virtual Machine Monitor VMM, translates calls of the guest OS requiring privileges.

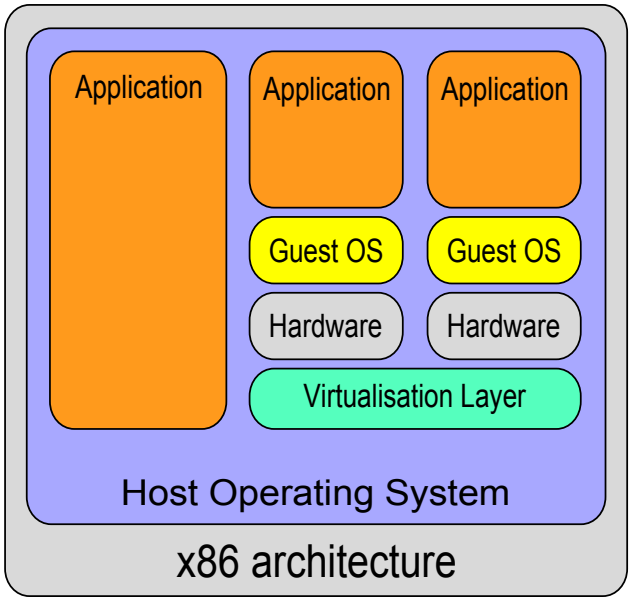

Figure 4.2: *Schematic overview of the hardware virtualisation by the VMware Workstation. The hypervisor (virtualisation layer) is executed as a program of the host operating system, allowing the execution of additional unmodified guest operating systems.*

This process is called binary translation. In addition, the hypervisor organises the usage of memory and CPU cycles enabling the definition of service levels for each virtual machine. The access to other components of the system like USB, sound or graphic card is provided by pass-through drivers to the real device of the host system.

VMware offers two different product lines for hardware virtualisation. The VMware Workstation is a program to be installed and executed on a Microsoft Windows or a Linux host enabling to start additional virtual machines in parallel to the running host operating system.

As the hypervisor is executed as an application, the performance of the virtual machines is limited as substantial resources are always required by the host operating system. Nevertheless, this product is well suited to execute additional operating systems on a single PC without the need for rebooting. A commodity version of this product is available as VMware Player which is designed to operate any virtual machine created by VMware products but does not support the creation of new virtual machines. A schematic overview of the VMware Workstation is presented in Figure 4.2.

The VMware ESX server consists of a hypervisor kernel which directly runs on the hardware. By requiring supported hardware components, it is ensured that special pass-through drivers, optimised for the dedicated hardware components are available for the best possible performance. In addition, the host operating system of the ESX server including the hypervisor does not require large resources which results in a minimal loss of performance for the virtual machines. Beside the near-native performance, advanced management tools are available. In the following, some

key features of the VMware ESX server are presented:

- Near-native performance: performance of the virtual machine approximately comparable with the native execution on the system hardware.
- Memory over-commitment and ballooning: It is possible to assign more virtual memory to the virtual machines as the physical host offers. In addition, the memory of the virtual machines can be adjusted during runtime.
- Suspend-to-disk and cloning: A virtual machine can be suspended to disk during runtime. This includes the duplication of the virtual machine by a simple copy process.
- Live-migration on the fly: A virtual machine can be migrated to another host system without any interruption. For this, the memory and the file system are synchronised between the two host machines. Only for the final synchronisation, the machine is not reachable for some microseconds. In addition, migration of the storage of virtual machines between different storage arrays is supported.
- Execution of a large number of different operating systems in parallel: Up to 128 unmodified OS are supported.

A schematic overview of the VMware ESX server is presented in Figure 4.3. In addition, several products are offered to extend the functionality of the VMware ESX server and to optimise administration. For example, the VMware Virtual Infrastructure adds features like high-availability and automated dynamic load balancing to the server's basic functionality. The Virtual Center package enables the central administration of multiple VMware servers.

Paravirtualisation

Using paravirtualisation, a hypervisor, running in the most privileged and lowest ring, organises the direct access to the native hardware components for the virtual machines. In contrast to the previous presented full virtualisation, no hardware is virtualised or emulated. In the following, this concept is discussed in detail with respect to Xen [78], a free hypervisor.

The hypervisor is booted together with a privileged instance of an operating system, in the Xen terminology called the *dom0* out of which the unprivileged instances of additional operating systems, the *domUs*, are started. Also administrational tasks like the adaptation and administration of the virtual machines are organised via the *dom0*. As the access to the system hardware is managed by the hypervisor, the host and guest operating systems have to be modified for such a usage. For different Linux flavours and other Open Source products, this can easily be done by patching the kernels. Though, this is not the case for commercial products like Windows where the sources of the kernel are not available and it is therefore not possible to provide the necessary patches [79].

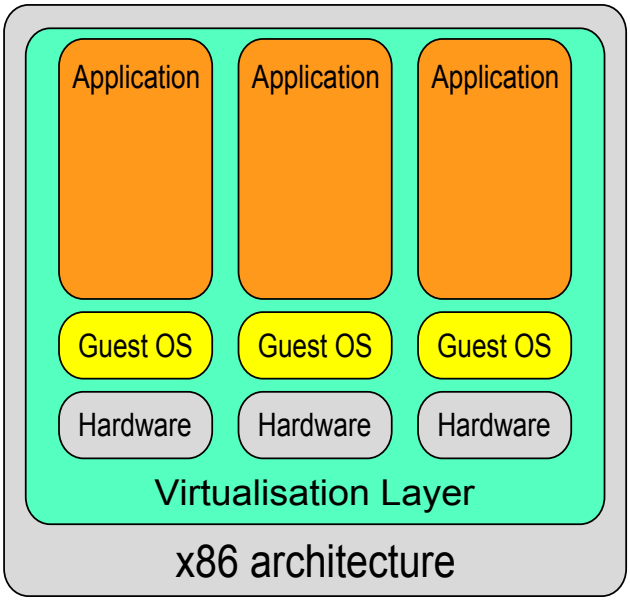

Figure 4.3: *Schematic overview of the hardware virtualisation by the VMware ESX server. The hypervisor runs directly on the bare-metal allowing to start several unmodified operating systems in parallel.*

To execute arbitrary operating systems without modifications under Xen, hardware virtualisation support is required. With the new processor generations of Intel and AMD that are discussed in Section 4.2, it becomes possible to run unmodified OS with Xen, but for the price of decreasing performance.

The big advantage of paravirtualisation is the near-native performance of the VMs. Additionally, all key features like adaptation of the memory of virtual machines, suspend to disk and migration between different host systems without interruption are provided. However, a dependence on the hardware architecture of the host remains as the components are not virtualised or emulated. In addition, the over-commitment of memory for the virtual machines is not possible by design [80]. A schematic overview of a Xen based system is presented in Figure 4.4.

Enterprise versions of the Xen hypervisor including tools for administrations are provided by XenSource, Inc., which has been acquired by Citrix [81] end of 2007.

Also User Mode Linux (UML) supports paravirtualisation by enabling the execution of several Linux systems in parallel. Isolation and hardware access are organised through execution of the guest Linux kernels in the user space of the host system for which the guest kernel has to be modified [82]. Since release 2.6 of the Linux kernel, the UML functionality is included in the main kernel tree.

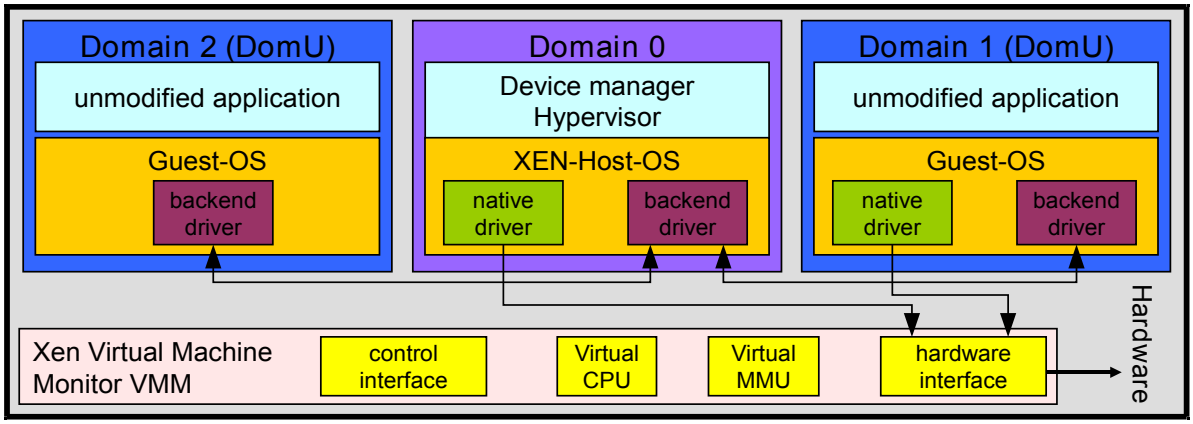

Figure 4.4: *Schematic overview of the paravirtualisation by Xen.*

x86 Hardware Extension for Virtualisation Support

With the evolving market and needs for virtualisation, AMD and Intel have developed x86 processor generations directly supporting virtualisation on the hardware layer. These new processor generations of both vendors, Intel Virtualisation Technology (VT-x, also known as Intel Vanderpool [83]) and AMD-V (also known as AMD Pacifica [84]), offer a new CPU execution mode for a virtual machine monitor in a more privileged level than the operating system which remains in Ring 0, see Figure 4.5. In this new architecture, the privileged calls of the operating system are automatically directed to the VMM.

Therefore, modifications of the kernel as required for the paravirtualisation are not needed anymore and the execution of unmodified guest operating systems becomes feasible. Also fully virtualised systems can use the extensions of the new processor generations to avoid the step of binary translation. Nevertheless, the VMware products only use this extension to support 64-bit guest systems as for the other guest OS, the performance of binary translation is still better compared to this first generation of hardware virtualisation [85] [69].

Besides this extension of the processor architecture, vendors have developed motherboards with an embedded flash or USB devices. On these devices, a hypervisor can be installed to provide virtualisation techniques directly integrated within server systems.

4.3 Hardware Consolidation using Virtualisation Techniques

Computing centres have to provide a multitude of services for the different user groups. As a dedicated service should not be affected by the failure of another, they are often executed on

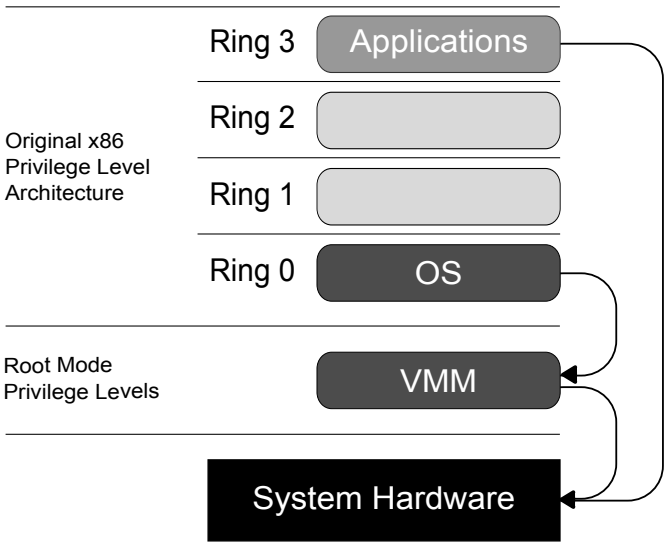

Figure 4.5: *Schematic overview of the x86 processor architecture including the VMM extension. To support virtualisation on the hardware level, a new level for the VMM is introduced. The operating system remains in Ring 0 whereas the VMM is executed in the more privileged Root Mode Level.*

different machines with additional advantages concerning security and maintenance.

The increase in computing power over the last years has created an environment in which the different machines are often not fully loaded by one single service. In addition, the server has to be designed to cover peak performance of the hosted service, which further decreases the load of the machine by definition. Once additional services are needed, new server machines are purchased or older machines are recycled leading to a heterogeneous hardware infrastructure, which complicates the administration of such an environment. Using virtualisation techniques, these servers can be consolidated on few machines. As an example, the computing situation of the IEKP as a typical university's institute is discussed in detail in the next section.

4.3.1 Typical Computing Situation at University Institutes

The computing infrastructure at the IEKP is shared between the three working groups of CDF, CMS and AMS. Besides the local computing cluster which consists of about 30 compute nodes (based on x86-32 or x86-64 architectures), five file servers with a total capacity of about 33 TB and 12 portal machines for software development and local job submission, the following additional services are required:

- LDAP server: The Lightweight Directory Access Protocol (LDAP) is used at the institute for the central administration of user accounts.

- Samba server: Samba, a free software re-implementation of the SMB networking protocol, offers file and print services for the Microsoft Windows clients at the institute.
- CUPS server: The Common Unix Printing System (CUPS) acts as the institute's print server to which the central printers are connected.
- Log server: The log server collects and archives log files from different machines.
- Trac server: Trac [86] is an enhanced wiki and issue tracking system for software development projects.
- Squid server: Squid [87] is a caching proxy for the Web supporting HTTP, HTTPS, FTP, and more. At the institute, it is used to get data from the CMS Conditions Database at CERN.

In this shared computing infrastructure, the three groups have different requirements concerning the operating system. For the desktop machines, an OS supporting newest hardware, multi-media applications as well as a prompt support with security updates is essential. Therefore, Debian [88] has been identified as the OS of choice for this class of machines.

The situation changes for the development machines and the compute nodes. Here, an operating system which supports the dedicated experiment specific software is required. Both, the CMS and CDF software are developed and tested under ScientificLinux.

Besides serving local users, the IEKP cluster is also connected to two grid infrastructures, the already discussed WorldWide LHC Computing Grid and the SAMGrid [89] used by CDF. Users from all over the world can, provided they have the permissions, submit their jobs to the IEKP computing cluster. To enable such grid functionality, the site-wide services for the WLCG middleware, a Computing Element, a Storage Element and a Monitoring System Collector Server have to be deployed by the grid centre to offer local resources to the collaboration. Also additional machines for the SAMGrid are necessary.

The requirements concerning the operating system of the different groups are summarised below:

- The CDF software is developed and tested with ScientificLinux FermiLab Edition version 3 and needs a 32 bit operating system. The SamGRID machines require also SLF 3.
- The CMS group favours the ScientificLinux CERN Edition version 4. The transition from ScientificLinux CERN Edition version 3 has been completed for the experiment specific software in 2007. Nevertheless, some grid services still require the older SLC version.
- The AMS group does not have special operating system requirements. They would, however, benefit from a 64 bit operating system.

As a result of these different requirements, the computing structure at the IEKP is inhomogeneous. In contrast to the cluster services like LDAP and CUPS which have no strong constraints concerning the operating system, the gLite middleware requires in the current implementation Scientific Linux 3 CERN Edition as operating system. The SAMGrid services run under ScientificLinux FermiLab Edition version 3.

To provide a reliable and stable production environment it is recommended to run each of these services with its own operating system, realised in the classical way by using separate machines. Nevertheless, experience has shown that the typical usage of these services at smaller clusters does not fully load the machines. Only sometimes peak performance is needed for short periods in time. As a consequence, older unused machines are recycled to host the different services. The resulting hardware became heterogeneous which complicates the installation procedure of the OS and the maintenance. In addition, there is a significant risk of failure and the peak performance is limited. Using virtualisation techniques can solve both, the different requirement concerning the operating system and the deployment of the different services in an efficient way. The general idea of hardware consolidation using virtualisation is presented in Figure 4.6.

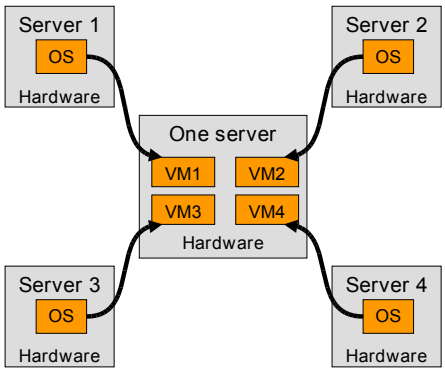

Figure 4.6: *Hardware consolidation using virtualisation techniques. In this example, four former physical machines now run as VMs on one host system.*

4.3.2 Benefits from Virtualisation

The appliance of virtualisation techniques offers plenty of advantages which will be discussed briefly below:

- Reduction of the hardware overhead. Many different operating systems may be executed on one single host machine. The whole infrastructure becomes easier to maintain and better loaded.
- Easy setup of additional virtual machines by simply copying a VM template image containing an OS which is fully adapted to the specific computing infrastructure.

- Easy migration of VMs to a new host system in case of hardware maintenance or load limits.
- Possibility to easily backup a whole OS and to restore a working state of a service.
- Cloning of virtual machines allows the replication of a service during runtime. This enables testing of installations without touching the production environment. Once the new installation is tested and confirmed to be stable, the production environment can simply be switched. This leads to fewer service interruptions and more effective administration.
- Migration of virtual machines to other host machines without a shutdown of the OS. If the file system of the virtual machine is located on shared storage, the migration can be performed without interruption of the machine. The maintenance of host machines becomes possible without affecting the services executed in virtual machines. Also load balancing and different service levels can be achieved using this feature.
- Possibility to set up test environments without the commissioning of new machines. Further resources are only necessary in the case of capacity limitations of the existing hardware, eventually hosted by a service provider.

Besides the advantages of virtualisation, one has to keep in mind that the host machine of the VMs becomes a single point of failure. Therefore, the system has to be designed carefully to protect or recover virtual machines in case of hardware failures of the host machines.

This can be realised through a connection of the server on which the virtual machines are actually executed to a highly available storage space for the file system of the VMs. In case of maintenance or failure of one of the host machines, its virtual machines can be migrated to the remaining hosts without interruption. The different virtual machines and with them the dedicated services are always available. Only the service level may be reduced due to the resource shortage. Assuming the worst case in which the failure of the host machine appears so quickly that a migration is not possible, the VMs can easily be restarted on the remaining server infrastructure on short time scales. One possibility to build up such a system is the combination of a server farm connected to a Storage Area Network (SAN) as shown in Figure 4.7.

For smaller sites it may not be possible to build up such a system due to financial constraints. An alternative is to connect two machines with identical processor architectures and to mirror the disks between these machines over Ethernet. Such functionality is for example supported by the Distributed Replicated Block Device (DRBD) under Linux [90]. This technique ensures that the file systems of the different machines are always available in an identical way on two different server machines. As the mirroring using DRBD is only possible between two machines, this approach does not scale and is therefore only suitable for sites with a limited number of virtual machines. A schematic overview of this concept is presented in Figure 4.8.

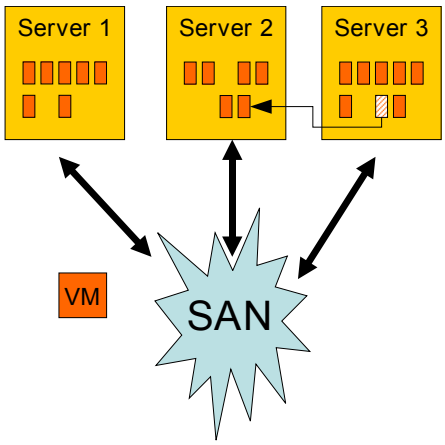

Figure 4.7: *Schematic overview of server consolidation in combination with a SAN. The different host machines are connected to a SAN, which holds the file system for the virtual machines. In case of maintenance or failure of one of the host machines, the virtual machines can be migrated to another host without interruption.*

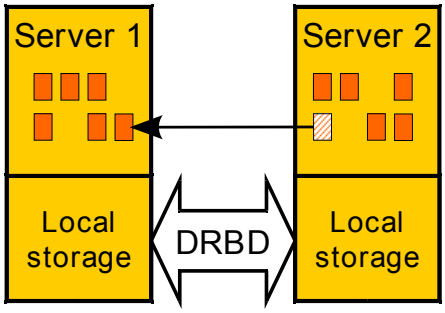

Figure 4.8: *Overview of server consolidation using mirrored disks. DRBD is used to mirror the disks of the two servers, holding the file system of the virtual machines. But such an approach is limited to a two server architecture.*

4.3.3 Choice of Virtualisation Technique

In order to identify the virtualisation technique of choice for the consolidation at the IEKP, the following products have been investigated:

- VMware Player, Workstation and ESX server
- Xen
- User Mode Linux

The VMware Player is freely available, but no administration tools are available and it is not possible to steer the virtual machines using scripts. In addition, the performance overhead of the VMware Workstation does not match the requirements. The ESX server includes high end administration tools and provides a near-native performance of the virtual machines. Due to financial

reasons, a small scale usage of this product at smaller institutes is improbable. For larger sites, the VMware ESX server is an excellent product for hardware consolidation.

The Open Source products Xen and UML support the steering of the VMs using scripts and are freely available. However, a performance comparison, presented in Figure 4.9, shows that Xen performs considerably better than UML. The circumstance that both products require kernel modifications of the virtualised guest system is negligible in the case of the IEKP as only Linux based systems are consolidated.

Therefore, the approaches presented in the following are based on Xen but can easily be transferred to other products like the VMware ESX server.

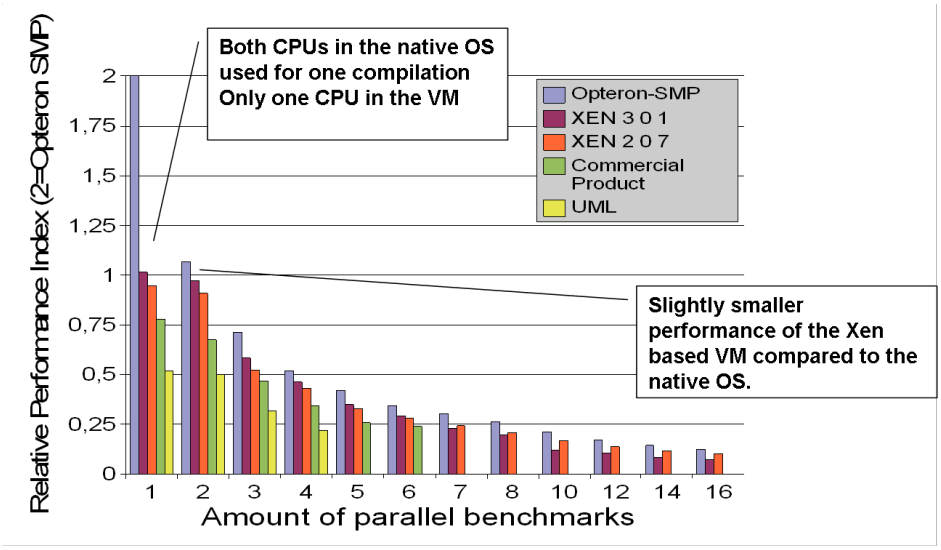

Figure 4.9: *Performance comparison of different products. As benchmark, a Linux kernel compilation allowing four jobs simultaneously (make -j 4) is used. Virtualisation using the Xen hypervisor brings a near-native performance. User Mode Linux, also following the concept of paravirtualisation, shows a significant loss in performance.*

4.3.4 Hardware Consolidation at a University's Institute

Due to financial constraints and the number of services to be consolidated, the IEKP has decided to use Xen in combination with DRBD for the consolidation of the different services. In the current implementation, the two identical host systems are AMD Athlon 64 X2 processor machines with 6 GB RAM, providing 400 GB of Raid 10 disk space each. This disk space is mirrored between them over Ethernet via DRBD. To ensure a stable connection for the mirroring, a dedicated network device is used. Both machines are installed with Debian as base system running a 2.6.19 kernel patched with Xen version 3.0.4.

At peak period, the following services have been migrated to virtual machines on this infras-

tructure. Before the consolidation, each of the services has been executed on a dedicated server machine.

- WLCG production environment:
  The three grid services CE, SE and MON are installed on the SLC 3 operating system. The version of the grid middleware is gLite 3.1 with a LCG-flavoured computing element.
- WLCG test environment:
  In order to test new releases of the middleware, a test system which is identical to the gLite production environment is deployed.
- SAMGrid production environment:
  The SAMGrid environment consists of two machines: One SAM station handling the data transfer to the local site, and a GlideCAF service which plays the role of the CE in the CDF middleware implementation. Both systems run ScientificLinux FermiLab Edition 3.0.5.
- Local cluster management services including LDAP, CUPS, Samba, Trac and the Squid server.

Using virtualisation techniques, the services, originally executed on more than ten different machines are now hosted on the two server architecture. Besides the economy of server machines, the service level for the different services with respect to peak loads is improved. In addition, the mirroring of the file system eases the maintainability of the system, including a better protection of the services against hardware failures. A summary of the results presented in this section is available in [91].

With the commissioning of the Institute Cluster at the computing centre beginning of this year, most of the grid services have been migrated to this cluster and are hosted again on a virtualised infrastructure.

4.3.5 Virtualisation of Training Environments

Server consolidation using virtual machines is an ideal candidate for the deployment of training infrastructures. Here, often a large number of identically configured operating systems have to be deployed in short time cycles, adapted to the purposes of the different courses. In addition, such school systems are often characterised by a small load. Here, virtualisation allows for easy deployment of many identical machines within few hours on a moderate number of server machines. One simply has to prepare a virtual machine containing the configured OS followed by a replication of this machine as often as needed to the different host servers. As an example, the computing infrastructure for the GridKA School 2006 [92] is discussed in detail below.

One of the main challenges during the school was the deployment of the infrastructure for the two gLite courses. For the gLite introduction course, several User Interfaces have been required to offer access to the grid infrastructure for the students. In addition, the existing resources of the used Virtual Organisation dgtest[1] have been increased by five additional sites hosted on the school's hardware, each consisting of a computing element, a storage element and some worker nodes.

Having finished the course, the infrastructure had to be changed for the gLite installation course, taking place the next morning. During this course, 15 groups should practise the deployment of a whole gLite site. Therefore, isolated operating systems for a CE, SE, WN, MonBox and a UI had to be deployed for each group, 75 machines in total.

To offer the computing environment for both courses, the infrastructure has been deployed in virtual machines on 20 Xen servers. The virtual machines have been distributed by simply copying the image of a template SLC machine to the hosts. The machine specific network configuration was realised using a DHCP server, assigning the IP and hostname according to the MAC address of the virtual machines. The latter was generated taking the information of the service and the number of the student group into account. Using virtualisation techniques, it became possible to offer isolated operating systems to each group including the deployment of the whole infrastructure within two hours.

At the GridKA School 2007 [94], the infrastructure for all courses was based on virtual machines, either using Xen or VMware ESX server as hypervisor. In addition, an introduction course to virtualisation has been offered for the first time.

4.4 Virtualisation of Batch Systems

Up to now, the concept of virtualisation has been discussed and used for server consolidation and the deployment of a large number of identical operating systems. Beside these usage scenarios, virtualisation can be extended to the field of batch systems to overcome some intrinsic limitations.

Traditionally, the computing environment at research centres and universities is designed, deployed and maintained by the different institutes themselves or centralised at the computing centre. The main arguments for the first approach are, that the institutes are able to design the infrastructure exactly according to their requirements and that the control remains within their responsibility. The price, however, is that each institute has to provide manpower and the basic infrastructure such as network and cooling space.

Though, the expertise for designing and running such High-Performance Computing (HPC) clusters is available at central computing centres and only few experts are necessary for the mainte-

[1] The Virtual Organisation dgtest of the d-Grid project [93] is offered for test and training purposes.

nance of different installations. In addition, it becomes possible to bundle tenders leading to better prices. This scenario becomes even more interesting once the installation is shared between different user groups. As a cluster exclusively used by one group is designed to offer sufficient resources in times of peak loads, the cluster usage is not at its optimum on average by design. This changes for installations used by different groups as it is unlikely that all groups need the full resources at the same time. Taking this into account, the size of the cluster will be smaller compared to the sum of the resources required for isolated clusters. The average load of such a shared cluster will therefore be higher.

Besides these advantages of a shared infrastructure, drawbacks of such an installation have to be considered. To share a computing infrastructure between different groups and to profit from load balancing, the cluster has to be homogeneous. Therefore, the groups have to agree on a common infrastructure with respect to hardware, operating system, installed software components and security issues. The latter comprises the situation that users may not want to execute their jobs on a system which serves users from other groups. This topic gets even more critical once the cluster is accessible via the grid.

There are, however, cases, where a common operating system for all computing nodes has disadvantages or is even not possible:

- Different groups are bound to different operating systems.
- The hardware has a mixed 32/64 bit architecture. In this case, the uniformity requirement forces the OS to be 32 bit, which does not make use of benefits of the 64 bit extensions. Some users could use these benefits for their projects.
- The hardware consists of 64 bit processors. One group can only use a 32 bit operating system, whereas the other groups could benefit from a 64 bit operating system.
- The requirements concerning the security level are excluding.

4.4.1 Concept of Horizontal Cluster Partitioning

For the reasons presented above, it may be difficult for groups to agree on one common computing environment. Partitioning the cluster into different subsets of nodes is then the only feasible way of sharing at least some resources of the common infrastructure like storage, network or cooling. If the partitioning is static or varies only seldom, opportunistic use of the computing resources is impossible. Dynamically changing the installed operating system to always match the required computing resources to the installation is a difficult task and switching from one operating system to another is rather time consuming. A schematic overview of the different scenarios is presented in Figure 4.10.

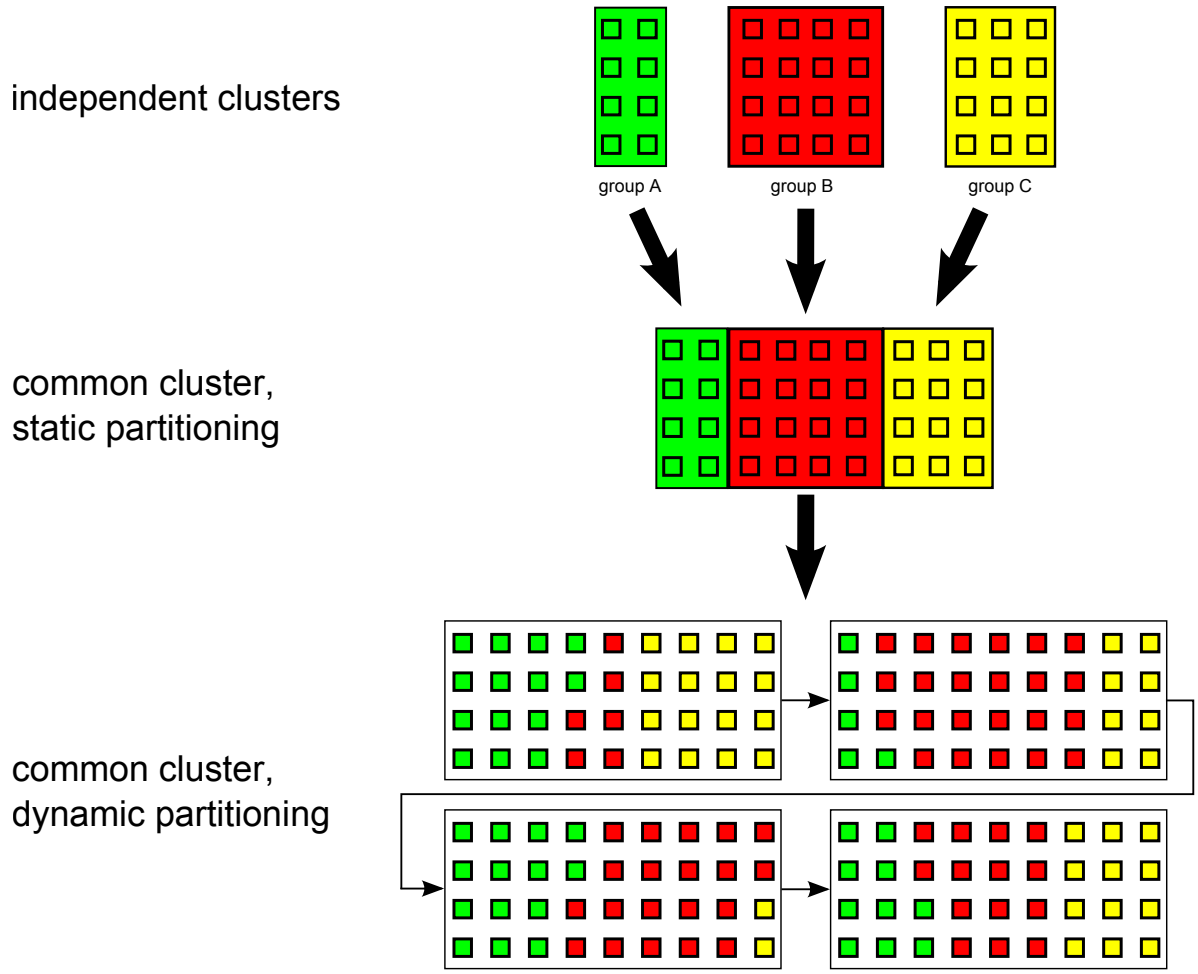

Figure 4.10: *Three typical possibilities to run a HPC infrastructure. The scenario on the top shows independent clusters maintained by the specific user groups. The second and third scenarios show shared, centralised clusters which run on the same hardware infrastructure. In these cases it can either be offered with statically or dynamically partitioned sub-clusters.*

Here, virtualisation comes into the play: using virtualisation techniques allows to install and run different operating systems in virtual worker nodes on the same physical host at the same time. In contrast to the "vertical partitioning" of a cluster, a "horizontal partitioning" using virtualisation techniques is much more flexible, can be better integrated into existing clusters and enables an optimal sharing of the computing resources. Figure 4.11 schematically shows the two different partitioning paradigms.

4.4.2 Performance Considerations

Such a cluster virtualisation only becomes feasible if the performance of the different applications does not suffer from virtualisation. For this, the performance of two benchmark scenarios reflecting

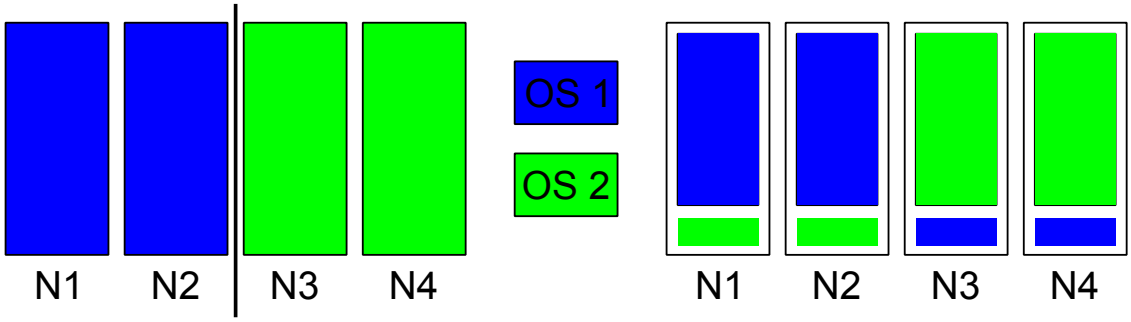

Figure 4.11: *A cluster of four physical hosts is partitioned vertically (left picture) and horizontally using virtualisation techniques (right picture). In the horizontal partitioning scheme, all VMs are running at the same time, the resource allocation can be adjusted dynamically.*

typical use cases of the IEKP groups have been investigated using the Xen hypervisor.

The first benchmark contains IO and CPU intensive standard applications of the CMS software framework like Monte Carlo simulations using Pythia [95] as well as simulations of the interaction of particles with matter with the Geant [96] program. The execution of this benchmark on a native Dual Opteron system with 4 GB of memory and in a virtual machine, hosted on the same system shows a decrease of performance below 4%.

A second benchmark measures the execution time of an application based on the GalProp program [97], which is used by the AMS group. This application can be compiled as a 32 bit and a 64 bit application. The benchmark was run on a Dual Opteron system with 2 GB of memory in a native 32 bit operating system. On an identical system a 64 bit version of the application was executed in a virtual machine. Both, the VM and the host system are 64 bit systems. In this scenario the job execution takes 22% less CPU time than in the native 32 bit scenario.

The results of both benchmarks show that the execution of standard applications of the IEKP does not suffer from virtualisation. The loss in performance is minimal for native 32 bit applications. Virtualisation may even bring a gain in performance as it allows to offer optimised environments for different applications.

4.4.3 General Concepts of Batch System Virtualisation

Ideally, virtualisation should not change the behaviour of the batch system, e.g. the number of jobs executed on a worker node or the prioritisation. Thus, three categories of batch queuing systems can be distinguished in terms of implementing a horizontal partitioning of a cluster:

- The batch queuing system support control mechanisms for the virtual machines and enables to relate several VMs to one physical host. In this case, the implementation a horizontally partitioned cluster reduces to the choice of the best fitting virtualisation technique and the

deployment of the different virtual machines.

- The batch queuing system does not offer native control mechanisms for the virtual machines. It has, however, the capability of grouping different VMs to one physical host. In addition to the choice of the virtualisation technique and the preparation of the operating systems, control mechanisms have to be implemented. Most batch queuing systems offer the possibility to run prologue scripts, which can be used to adequately prepare the virtual machine before the execution of the job.
- The batch queuing system does not provide native control mechanisms for virtual machines. Relating different virtual machines to one physical host is also not possible. In this case both features have to be implemented in around the batch queuing system.

In the following, two implementations of a batch system virtualisation are presented in detail.

4.4.4 Cluster Virtualisation without Modification of the Batch Queuing System

This first approach is mainly driven by the idea of a horizontal cluster partitioning without any modifications of the batch queuing system. This is motivated by the large number of commercial batch systems in use for which generally no source code exists.

The batch system of choice is the Maui/Torque framework as it is widely used in the high energy physics community. In addition, this framework is supported by the gLite middleware. The architecture of Maui/Torque does not hold the functionality of handling or grouping virtual machines. Therefore it is of the third category defined above and all functionality necessary for a cluster virtualisation has to be added. The virtualisation is realised using the Xen hypervisor as already motivated. In the next sections, the preparation of the worker nodes and the features, added to the batch system, are presented.

Preparation of the Worker Nodes

In a first step, the physical worker nodes are installed with a stripped-down Debian Linux including the Xen hypervisor. This privileged domain is only used for administrative purposes like maintenance and the management of the virtual machines. Allowing user jobs to be executed in the *dom0* would introduce a disparity between the *dom0* and the unprivileged domains, the *domUs*. In addition, restriction of user activity to the *domUs* grants the encapsulation of the different jobs in the virtual machines. The file system of the operating system for the virtual machines is stored on a central NFS server but could also be installed locally into an image or a partition. Each desired operating system is exists as virtual machine.

To offer the operating system of choice for the user job on the worker node, two different implementations are possible. For each job, the desired operating system is started in a virtual machine. Once the VM has booted, it has to be registered at the batch management server before sending the job to the virtual machine. Having finished and transferred back the output of the job, the VM is unregistered from the batch system and finally shut down. The worker node is now waiting for the next job to repeat this workflow. Unfortunately, the registration process of new machines is quite slow in the Maui/Torque framework.

In the second approach, all virtual worker nodes which could potentially be requested are running all the time on the physical worker node with a minimum of memory allocation. In the present case, a stable operation of the different virtual machines with only 32 MByte of allocated RAM was possible. Once a job is submitted to such a virtual machine, its memory gets a higher share of the total memory. In a scenario in which only one job per physical worker node is permitted, the VM executing the user job gets all memory available. Using Xen, the increase of memory for a *domU* is almost instantaneous. The CPU consumption of the memory minimized virtual worker nodes is negligible and the one executing the user job gets all CPU time. As all virtual machines are always running, there is not time lost due to the booting process and the registration with the batch server. Therefore this approach has been chosen.

Integration into a Batch Queuing System

After the virtualisation of the computing nodes, the different virtual nodes have to be registered with the batch queuing system. As Maui/Torque does not provide the grouping of virtual worker nodes to one physical host, this feature has been implemented around the batch queuing system. To avoid modifications of the batch system itself, this functionality is implemented in an external program.

The overall sketch of such a program being implemented as a daemon[2] and running continuously is drawn in Figure 4.12.

As a first step, the daemon ensures to have the full control over all nodes and all jobs. In the Maui/Torque wording, this is realised by setting all nodes offline and all jobs to hold. At this moment no job can be submitted to any virtual worker node.

The scheduler assigns a priority to each job waiting in the queue according to fair share, expected running time, waiting time, and other criteria. This information is queried by the daemon which then sorts the jobs according to their priority. It is important that the priority assignment policy is not touched by the daemon. This remains the task of the scheduler which is of utmost importance for computing centres with existing policies.

[2] A daemon is a computer program which is executed as background process, rather than under the direct control of a user.

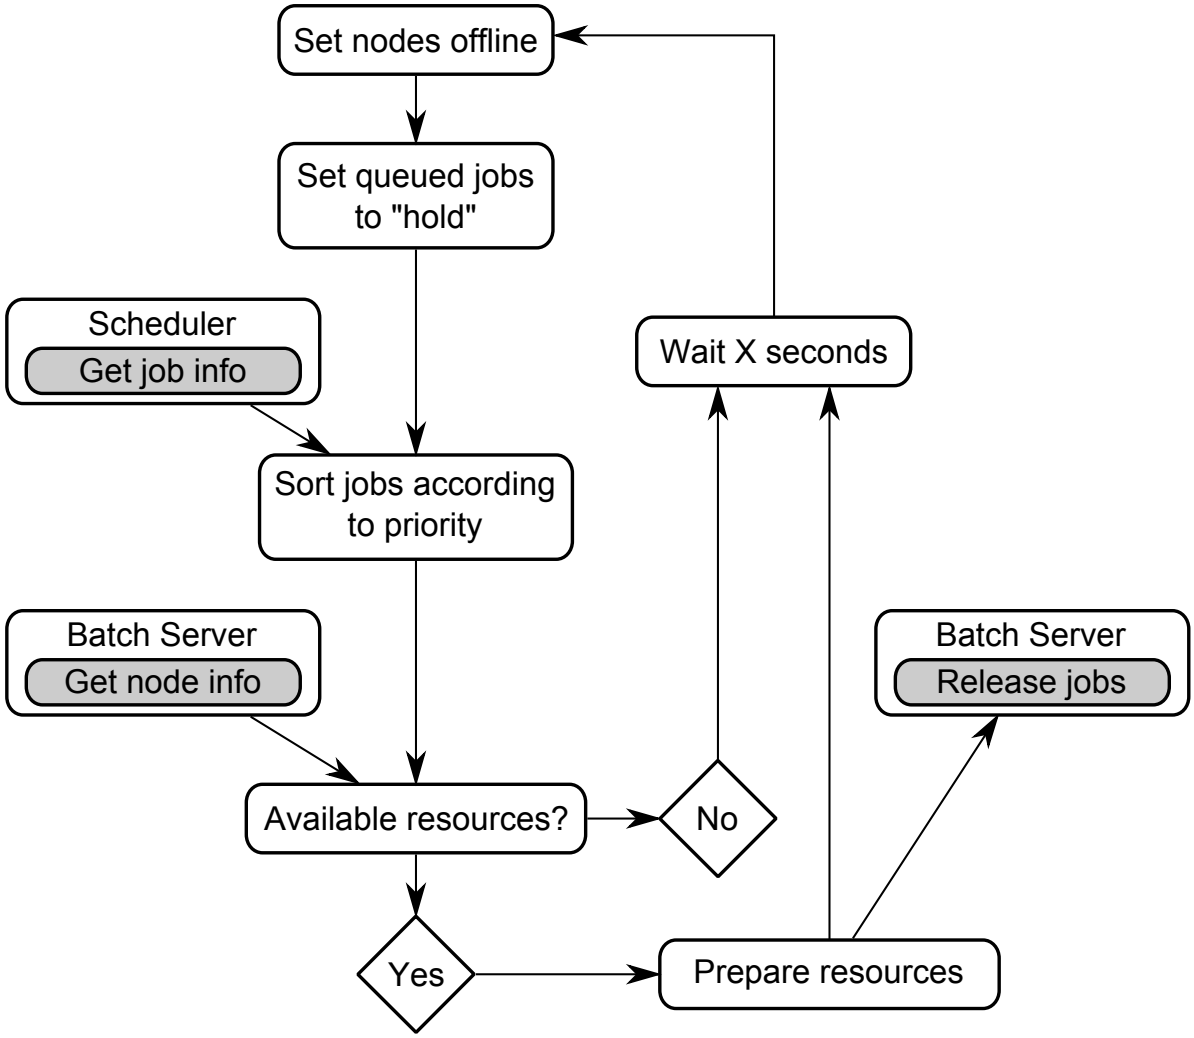

Figure 4.12: *General sketch of a daemon implementing the grouping of virtual machines to one physical host. The gray shaded areas represent standard interfaces to the external scheduler and batch system.*

In the next step, the daemon determines the free resources by querying the batch queuing server. In this context, a free resource is defined as a worker node that can accept another job. Again, the daemon supports the implementation of a site policy. One site could for example allow three jobs on one CPU, or one IO and one CPU intensive job per physical host and the daemon does not change this allocation policy. If there are free resources, the next jobs in the prioritised list are sent to them.

Before sending a job to a virtual worker node, the physical host and the virtual worker node needs to be prepared. When the decision is made that a job should run on one virtual worker node, its memory allocation is set to the maximum possible. In addition, the state of the virtual worker node is set to “free”, and the job is released. This way the daemon can control which job is submitted to which virtual worker node, ensuring that the job runs on the desired operating system, and that the physical host can offer resources to the virtual worker node.

After the job is submitted, the state of the node is set to “offline” again to ensure that no other job can be submitted to this virtual worker node. After a short waiting period, the whole procedure is repeated.

Such a daemon has been implemented in the Perl programming language [98]. It has been interfaced to the Maui/Torque batch queuing system and evaluated on the IEKP production cluster. A summary of the results presented in this section is available in [99].

Discussion

During the evaluation period of this daemon on the IEKP production cluster, the behaviour and particularities of this implementation have been investigated. The possibility to offer different operating systems adapted for the usage of different user groups is a desirable feature. Also maintenance and deployment of new operating systems is eased and becomes more efficient as shorter service interruptions are necessary. Unfortunately, the facts that an additional program bypasses some key features of the batch system and the scheduler causes problems, which cannot be avoided following such an approach.

The suggestion after that test phase is that the principle of horizontal partitioning should be implemented directly into existing batch systems. Providing the functionality to group virtual machines per physical host and to define rules on this group would be sufficient.

It is foreseen to implement the concept of batch system virtualisation as a key feature of the new Institute Cluster. Therefore, the future development of this concept is investigated in an additional diploma and PhD theses and the main enhancements of this project are briefly presented in the next section.

4.4.5 Cluster Virtualisation supported by the Batch Queuing System

One of the key features of the new Institute Cluster at the computing centre is the realisation of a batch system virtualisation. The approach presented in the following is therefore adapted to the boundary conditions of this installation including the chosen batch system.

To avoid bypassing the batch system, only functionalities which are already implemented are used. Such a common feature of batch or workload management systems are the prologue and epilogue scripts, executed before or after the job on a dedicated worker node. The standard usage of these scripts is to prepare and clean up the worker node. Therefore, they will be used for the preparation of the dedicated virtual infrastructure for the job followed by the shut down once the job has finished. Hence, the actual user job is encapsulated in an additional job executed as a prologue

script. From here on, this additional job is referred to as the wrap-job and its workflow is sketched in Figure 4.13.

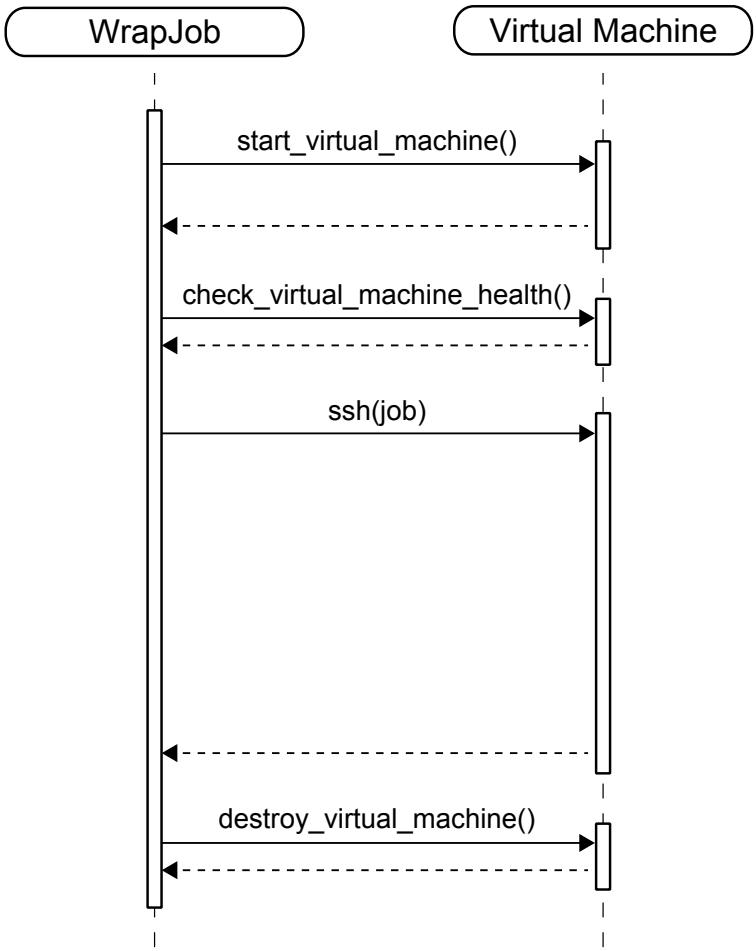

Figure 4.13: *Schematic overview of the workflow of the wrap-job. Executed as the prologue script of the batch system, it prepares the virtual machine for the job, passes the user job to it and organises the clean up once the job has finished.*

Once a user job is assigned to a free worker node, the wrap-job deploys the required virtual machine on the physical resource. In the current implementation, this is realised by cloning the virtual machine from a centrally stored template. As soon as the requested VM has booted successfully, the user job is passed to it using a pipe created by the Secure Shell protocoll (SSH). The network configuration of the virtual machine is again realised using DHCP. After the job has finished, the wrap-job takes care of the disposal of the virtual machine. In cases where the job is terminated by a system signal from the batch system, the termination of the VM is carried out by the epilogue script. As the virtual machine will not be used anymore in both cases, it is simply destroyed and its file system is deleted.

The advantage of this implementation is that each user gets a well defined, isolated operating system which has not executed other user jobs before and will not in future. Nevertheless, in order

to enable forensics for security issues and debugging, the log files of the virtual machines have to be stored for some period in time. Currently, the virtual machine templates are configured in a way that all required log files are stored on a central log server. The current implementation of this system is now ready to be deployed and tested at the Institute Cluster.

In the meantime, the concept of batch system virtualisation enjoys great popularity. Some of the features discussed above are included in commercial batch systems like Platform LSF [100] for a short time and should be evaluated for their suitability for the IC cluster. In addition, the Amazon Web Service EC2 (Elastic Computing Cloud) offers a web service interface to the user to request an operating system instance based on virtual Xen machines [101]. The dedicated VM can either be instantiated from an existing public image file or crafted by the user itself.

Simulation and Analysis Framework

Before being able to perform physics analyses, the output signals of the different detector components have to be processed in order to reconstruct the physics processes that have taken place inside the detector. This step is referred to as the reconstruction. In addition, a comparable amount of corresponding Monte Carlo data has to be produced according to the current knowledge of particle physics as well as of the design and expected response of the detector. To offer such functionality, the modular CMS software framework CMSSW has been developed.

In the following, the general architecture of this framework is presented including the discussion of the full simulation and reconstruction chain in CMS, which can be divided into three main steps. First, a Monte Carlo generator calculates all particles resulting from a proton-proton collision according to the theory of particle physics including the description of their decay until they reach the innermost layer of the detector. At this point, the detector simulation software starts simulating the interaction and trajectories of the incident particles with the detector material and the electronic output signal of the different detector components. In the last step, the trajectories and energy deposits of the particles will be reconstructed according to these signals. The latter is identical for Monte Carlo and collision data.

5.1 The CMS Software Framework

The CMS software framework CMSSW is a modular collection of software which is built around the Event Data Model (EDM) and includes the dedicated services which are needed by simulation, calibration and alignment. Modules to process the event data and for the reconstruction of the objects necessary for physics analyses are as well part of the framework. The architecture of CMSSW is based on a software bus model with a single executable called *cmsRun* and many plug-in modules which run the different algorithms. This allows to use the same executable to

process collision and Monte Carlo data [102]. The specification of required modules including the parameter settings is realised by configuration files which steer the executable at run-time. The requested modules are dynamically loaded at the beginning of the job. An example for such a configuration file is given in Appendix C.1.

The CMS event data model is based on the concept of an event as a C++ object container, the EventContainer, for all raw and reconstructed data of a physics event. During the processing, the data are passed from one module to the next via this container. These event data are completed by periodically updated information about the detector and accelerator status such as calibration, alignment, geometry descriptions, magnetic field and run conditions which are not only valid for a single event but for a certain period in time. They are therefore not part of the event container itself but stored and accessible via the EventSetup system.

In the current implementation, six basic types of dynamically loadable processing modules are available for the user:

- Source: Reads the EventContainer and the EventSetup from file and gives access to the information.
- Producer: Reads data from the EventContainer, processes them and adds this new information to the container.
- Filter: Reads data from the EventContainer and returns a boolean value which can be used to indicate whether the processing of that event should continue.
- Analyzer: Study the properties of an event. Typically, this type of module is used to produce histograms or other output.
- Looper: Control multi-pass looping over input data. This type of module is used for example in the track based alignment procedure.
- Output: This is the last module executed in a job. It writes out all processed EventContainer and EventSetup information to a file.

5.2 Monte Carlo Event Generation

The purpose of event generators is to produce hypothetical events of high energy particle collisions with kinematic and topological distributions as predicted by theory. For this, Monte Carlo techniques are used to determine the different variables according to the desired probability distributions and to generate events with all details that could be observed by a perfect detector. Over the recent decades, several packages have been developed by theory groups for a wide range of collider experiments including a tuning of the generators to provide the best possible description of

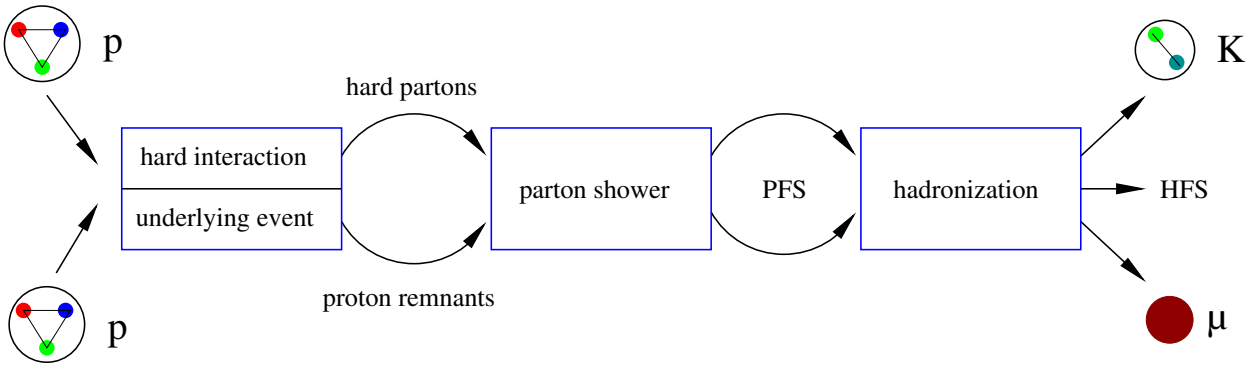

Figure 5.1: *Overview of the different steps from the hard process to the Hadronic Final State (HFS) for a multi-purpose Monte Carlo event generator. First, the hard interaction of two partons from the incoming protons is calculated according to the matrix element followed by the parton shower. The resulting Partonic Final State (PFS) is hadronised leading to the final stable particles [107].*

available collision data. Although the details of the implementation of the underlying physics processes are different for the available generators, the fundamental principles are the same. Interfaces to many generators like Pythia, Herwig [103], Alpgen [104] or MadGraph [105] are implemented in CMSSW. The study presented in this thesis is based solely on Pythia 6.4 [106] as event generator which is briefly presented in the next section.

Pythia

Pythia is the most frequently used multi-purpose generator in CMS for pp-, ep- and ee-collisions which is developed and maintained by the theory group of the Lund University since 1978. It is well designed to describe multi-particle production between elementary particles as detailed as experimentally observable. For this purpose the program contains models for a large number of physical aspects briefly presented in the following.

The initial interaction between two partons of the incoming protons is called the hard process. It can be calculated at leading order precision by evaluation of the matrix element of the investigated process. In this context, the component of transferred momentum of the hard process which is perpendicular to the beam axis is denoted as $\hat{p_{\mathrm{T}}}$. In the current version of Pythia, about 300 different hard processes are supported, including physics beyond the Standard Model. Radiative effects and higher order corrections are derived using the concept of parton showers starting from the particles of the hard interaction.

The resulting partonic final state is then left to hadronise according to the Lund string fragmentation model. This phenomenological model is based on the picture of linear confinement. Strings are modelled as quark-antiquark pairs moving apart and being bound together by the colour field. As a consequence of the strong interaction, the energy of this field increases with the distance be-

tween the two partons and at a certain point, the string breaks producing a new $q\bar{q}$-pair. Through the combination of a quark from one breaking with an antiquark from an adjacent one, colour-neutal hadrons are formed. This procedure is done until only hadrons with no net colour remain.

Furthermore, Pythia handles the behaviour of the beam remnants as well as the possibility of multiple interactions. Both processes as well as initial and final state radiation are referred to as the underlying event and cannot be derived from first principles. It is described by many free parameters, which have been determined based on model tunings to collision data. However, extrapolations of the different tunes to the LHC energy range show large discrepancies which indicate the need for a measurement of the underlying event at the LHC. More information on the different tunes used for this study are given in Appendix C.2.

5.3 Detector Simulation

In order to compare the events coming from the event generator, the interactions of the different particles with the detector material have to be taken into account as well as the response of the detector electronics. In the CMS software framework, this is realised in two subsequent steps. First, the physics processes that accompany the passage of all generated particles through the material of the detector and the response of the different subdetectors are simulated based on Geant4 [96]. Geant provides a rich set of physics processes to model the electromagnetic and hadronic interactions in the presence of a magnetic field based on the detailed geometry of the detector and the properties of the different materials. This includes energy loss through ionisation, multiple scattering, bremsstrahlung and the processes of electromagnetic and hadronic showering.

The result of each particle's interactions with matter like its energy loss in a dedicated volume of a subdetector is stored in the form of *simulated hits*. As final step, the response of the readout electronics of the different detector components is simulated which is called the *digitization*. To safe resources, additional collisions during a bunch crossing, the pile-up events, are simulated independently and are merged with signal events afterwards. This is done by a special module which superimposes the hit level detector information of the pile-up and signal events and merges the corresponding Monte Carlo truth information. An additional advantage of such a procedure is that the signal samples can be used for different run conditions like low and high luminosity running, which result in different pile-up scenarios.

5.4 Event Reconstruction

Reconstruction means the creation of physics quantities or objects based on the output of the detector. This step does not distinguish between data recorded by the data acquisition systems

Muon p_{T}	$\Delta p_{\mathrm{T}}/p_{\mathrm{T}}$ stand-alone muons	$\Delta p_{\mathrm{T}}/p_{\mathrm{T}}$ global muons
10 GeV	8-15%	0.8-1.5%
1000 GeV	16-53%	5-13%

Table 5.1: *Resolution of the transverse momentum of the muons using the stand-alone and global muon reconstruction [110].*

of the real detector or the simulated detector response described above. The algorithms for the different reconstruction steps are available as modules in the framework. In the following, the muon and jet reconstruction is presented in detail.

5.4.1 Muon Reconstruction

The muon reconstruction in CMS is divided in three steps. The reconstruction chain starts with the local muon reconstruction where only data from the three muon systems (DT, CSC and RPC) are considered [108]. Hits within the individual chambers are matched to form track segments. Starting from the innermost muon station, Kalman filter techniques are used to combine the different track segments and hits from all muon detectors to form the stand-alone muon tracks [109].

For the global muon reconstruction, the stand-alone muon tracks are combined with tracker information. Therefore the muon tracks are extrapolated from the innermost muon station to the outer tracker surface and to the nominal interaction point to search for compatible hits and tracks reconstructed in the silicon tracker. Finally, a global fit is performed which combines hits from the silicon tracker with those in the muon system originally forming the stand-alone tracks. A selection on the tracks of this final fit provides the global muons. As a result of the high resolution of the tracker combined with the strong magnetic field inside the solenoid, the transverse momentum resolution of the global muons is improved by a factor of ten compared to the stand-alone tracks.

The expected performance of the muon system has been investigated using a full detector simulation and the result for the resolution is summarised in Table 5.1. Whereas the resolution for muons with low transverse momentum is limited by multiple scattering in the return yoke, the limiting factor for high p_{T} muons is the resolution of the muon chambers [110].

5.4.2 Jet Reconstruction

As a consequence of the confinement in QCD, coloured partons of the hard process hadronise and only a collimated stream of colour-neutral objects can be observed with detectors. In order to create a link between these observables and the energetic partons created in the hard process, the concept of jets is introduced.

A jet is defined as the cluster of all particles which are supposed to originate from the same initiator.

This allows to link the properties like momentum and energy of the jet to the parton. For the clustering of the particles, different jet clustering algorithms have been developed. However, it is impossible to unambiguously map each hadron to a dedicated parton.

Whereas the basic idea of the clustering algorithms remains the same, the actual implementation differs. To be able to run an algorithm on objects coming from different levels of theory calculation, generation or reconstruction, the properties of the objects used as input to the algorithm should be comparable. A suitable choice is for example a four-vector like quantity, which provides a three dimensional direction and the energy as both information can be derived either from calculated particle properties or measured energy deposit in a calorimeter cell. For the clustering of the input objects, each jet algorithm has to provide the following basic functionality:

- a distance measure, which defines the separation of two objects,
- a procedure to decide whether objects are combined or not,
- and a recombination scheme, which defines how objects are combined.

It can be distinguished between two techniques used for the clustering: Cone type algorithms which cluster objects together being in a cone and clustering algorithms, which combine objects having the smallest distance of all possible pair wise combinations.

To compare jet results between different experiments and theory prediction, a precise non-ambiguous definition of the jet algorithms and the input variables is required. Such a possible definition as well as the actual implementation of the different algorithms, which have been used for this study including their basic properties are presented in the following.

Jet Nomenclature

A jet definition specifies a procedure by which an arbitrary set of (physical) four-vectors is mapped into a set of jets. To communicate experimental results and to compare them with theoretical prediction, the following elements have to be specified when dealing with jets according to the Standard Model Handles and Candles Working Group of the Les Houches Workshop 2007:

- **The definition of the jet clustering algorithm** specifies all details of the procedure which is used to map an arbitrary set of four-momenta from physical objects into a set of jets. This includes the name of the algorithm, all its parameters and the recombination scheme, which defines how the constituents of a jet are combined to calculate the energy and direction of the jet.
- **A specification of the final states truth level** is required for a consistent comparison between different experimental results or with Monte Carlo simulations. Therefore, the jet

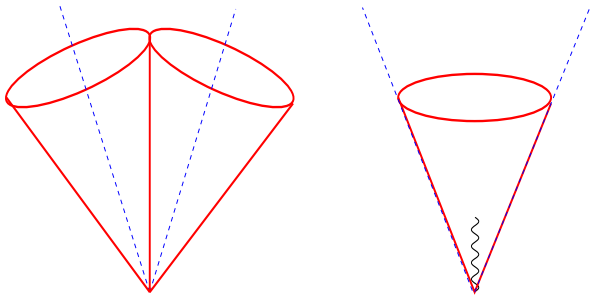

Figure 5.2: *Illustration of infrared sensitivity of a seed based jet clustering algorithm. The picture on the left hand side shows a configuration which is clustered into two jets. However, adding a soft parton to this configuration leads to the merging of the jets and only a single jet is clustered by the infrared unsafe algorithm [112].*

definition has to be supplemented with an exact definition of the set of physical objects used for the jet clustering. For example, it has to be specified whether measured input objects have been corrected for some aspect or if muons from resonance decays or neutrinos from Monte Carlo simulations are considered.

A more detailed discussion including recommendations for the jet definition and the specification of the final state truth level can be found in [111].

Infra-Red and Collinear Safety

One of the desired features of jet algorithms is collinear and infrared safe behaviour. Requiring that adding an infinitely soft parton does not affect the output of a jet clustering algorithm is called infrared safety. Such soft partons may be a result of soft gluon radiation during the parton shower or the hadronisation. Algorithms which start with the jet clustering around towers with high energy deposit, the seed tower or simply denoted as seed, can by sensitive to soft radiation as schematically illustrated in Figure 5.2. Adding a soft parton to a configuration with two separated jets could lead to a sufficient energy deposit that the two original jets are merged into one single jet. Such soft radiation may also come from pile-up and the underlying event.

Collinear safety denotes in this context that adding or replacing any massless parton by an exactly collinear pair of massless partons does not affect the output of the jet algorithm. This means that for example collinear gluon radiation emitted during the jet evolution does not lead to a different interpretation. An example for collinear unsafe behaviour are seed based algorithms as drawn in Figure 5.3. Starting with one jet produced in the hard process, two scenarios are possible. If the resulting energy concentration is sufficient to produce a seed for the jet clustering, the algorithm clusters it into one jet. The situation changes if the energy deposit is split into two equal parts,

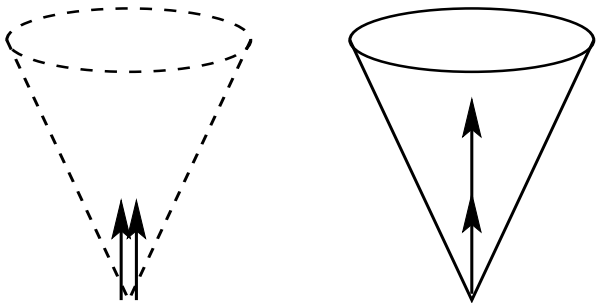

Figure 5.3: *An example for collinear-unsafe behaviour of a jet algorithm. For both scenarios, the production of one jet in the hard process is assumed. In the example on the left hand side, the energy concentration is sufficient to produce a seed for the jet clustering leading to one jet. However, if the energy deposit is split into two equal parts, the energy concentration is no longer sufficient to act as seed for the clustering and no jet is found [112].*

the energy concentration may no longer be sufficient to act as seed for the clustering and no jet is found. This may happen through collinear radiation effects or due to the granularity of the calorimeter. A cluster of particles which hits exactly the middle of two calorimeter towers would not be clustered into a jet whereas the same energy deposit in one tower results in a jet.

A more detailed discussion of the attributes of ideal jet algorithms from a theory and experimental point of view can be found in [113].

Iterative Cone Algorithm

The Iterative cone algorithm starts with a list of particles ordered by their transverse energy. The object with the largest transverse energy is taken as seed and all objects within a cone with the radius

$$\Delta R^2 = (\Delta\eta)^2 + (\Delta\phi)^2$$

in the η-ϕ-plane are merged into a jet called the proto-jet. The energy and direction of this proto-jet are thereby calculated according to

$$E_{\mathrm{T}} = \sum_i E_{\mathrm{T}}^i \tag{5.1}$$

$$\eta = \frac{1}{E_{\mathrm{T}}} \sum_i E_{\mathrm{T}}^i \, \eta_i \tag{5.2}$$

$$\phi = \frac{1}{E_{\mathrm{T}}} \sum_i E_{\mathrm{T}}^i \, \phi_i \tag{5.3}$$

which is called the E_{T} recombination scheme. The direction of this jet is taken as seed for the next proto-jet and this procedure is iterated until a defined ending criterion. This may be that a

change of the proto-jet axis or energy is below a certain threshold or that a maximum number of iterations has been performed. The result is a stable proto-jet which is added to the list of found jets. After having removed all constituents of this jet from the list of objects to cluster, the next proto-jet is build of the remaining objects. This is repeated until the input list does not contain any objects above a certain threshold [114]. However, as this cone based algorithm uses seeds, it is not collinear and infrared safe.

Midpoint Cone Algorithm

The Midpoint cone algorithm is designed to address some disadvantages of the Iterative cone algorithm. First, all objects above a certain threshold are taken as seeds and the objects inside this cone are clustered into proto-jets. In contrast to the previously discussed Iterative cone algorithm, the objects are not removed from the input list and can therefore belong to different proto-jets. For all proto-jets which are closer than the diameter of the cone, the midpoint is calculated as the direction of the combined momentum. Then, a second iteration is done based on the midpoints and the proto-jets as seeds leading to an increased number of proto-jets after this second step. As no objects are removed from the input list, objects may belong to several jets and the proto-jets overlap. This is resolved by splitting and merging routines.

This procedure starts with the proto-jet with the highest transverse energy. If the jet does not share objects with other proto-jets, it is defined as stable and removed from the list. If there is an overlap with other proto-jets, the further processing depends on the fraction of transverse energy shared between the two jets. If this is smaller than the splitting parameter E_{split}, the objects are assigned to the proto-jet which is nearest in the η-ϕ-plane. Otherwise, the two jets are merged into one according to the selected recombination scheme. The stable cones after these steps are the final jets.

Adding midpoints as seeds for the second iteration has been introduced to improve infrared and collinear safety compared to the Iterative cone algorithm. But this procedure cannot solve the problem completely.

SISCone Algorithm

The Seedless Infrared-Safe Cone (SISCone) jet algorithm is a cone clustering algorithm, which is designed to be infrared safe. Whereas infrared and collinear unsafety have been introduced into the Iterative cone and Midpoint cone algorithms by using a dedicated number of seeds for the jet clustering, the SISCone algorithm avoids this by searching for all stable cones. Using a brute force technique for this would be to test all subsets of N input objects for stability. However, such an approach is not suitable as the number of such distinct subsets grows with N^2, which has implications for the needed computing time.

Therefore, the SISCone algorithm follows a more advanced approach by exploiting the fact that a circle enclosing a set of two input objects can be moved around such that two of the remaining objects lie on its circumference. Reversing this allows the determination of all stable cones with a radius R by testing the circles defined by a pair of objects and radius R. With this procedure, all stable proto-jets can be found which may overlap. Splitting and merging is done by the same procedure as described above for the Midpoint cone algorithm except that the scalar sum of transverse momentum is used as ordering parameter. A detailed description of these procedures can be found in [115]. It has been demonstrated, that the hard jets are affected by the addition of soft objects in a fraction less than 10^{-9} of the investigated events. For comparison, for the Midpoint cone about about 15% of the events are affected.

$\mathrm{K_T}$ and Cambridge/Aachen Algorithms

Unlike the previously discussed cone-like algorithms, pairwise recombination algorithms are not based on a fixed geometrical shape of the jet. These algorithms use a resolution parameter which depends on the geometrical distance ΔR and additionally on the energy of the input object. Therefore, objects are clustered to a jet if they are close in four-vector space. For each object k of the input list, the distance d_{kl} to every other object l is calculated as well as its distance d_{kB} to the beam. To be able to control the size of the jets, the resolution parameter D is introduced.

- $\mathrm{k_T}$ algorithm: $d_{kl} = \min(p^2_{\mathrm{T}k}, p^2_{\mathrm{T}l})\Delta R^2_{kl}$ and $d_{kB} = p^2_{\mathrm{T}k} \cdot D^2$
- Cambridge/Aachen algorithm: $d_{kl} = \Delta R^2_{kl}$ and $d_{kB} = D^2$

The spatial distance of the two objects k and l is denoted with $R^2 = (\Delta y)^2 + (\Delta\phi)^2$. For all input objects, the distances d_{kl} and d_{kB} are calculated and the minimum of all these distances is searched and labelled $d_{\min}$.

If this $d_{\min}$ is a beam distance d_{kB}, the object is removed from the input list and added to the list of jets. However, if $d_{\min}$ is a distance between two objects, both are removed from the input list and merged according to the selected recombination scheme. This merged object is added to the input list and the procedure is repeated. If the input list still contains objects, the procedure begins anew. This procedure describes the inclusive mode to run the algorithm and is used for this study.

In the standard implementation KtJet [116], the computing time increases with N^3 with the number of input objects. However, a more performant implementation of the algorithm is available as the FastJet [117] package and is used within the CMSSW framework. An overview of the running times of different jet clustering algorithms is presented in Figure 5.4.

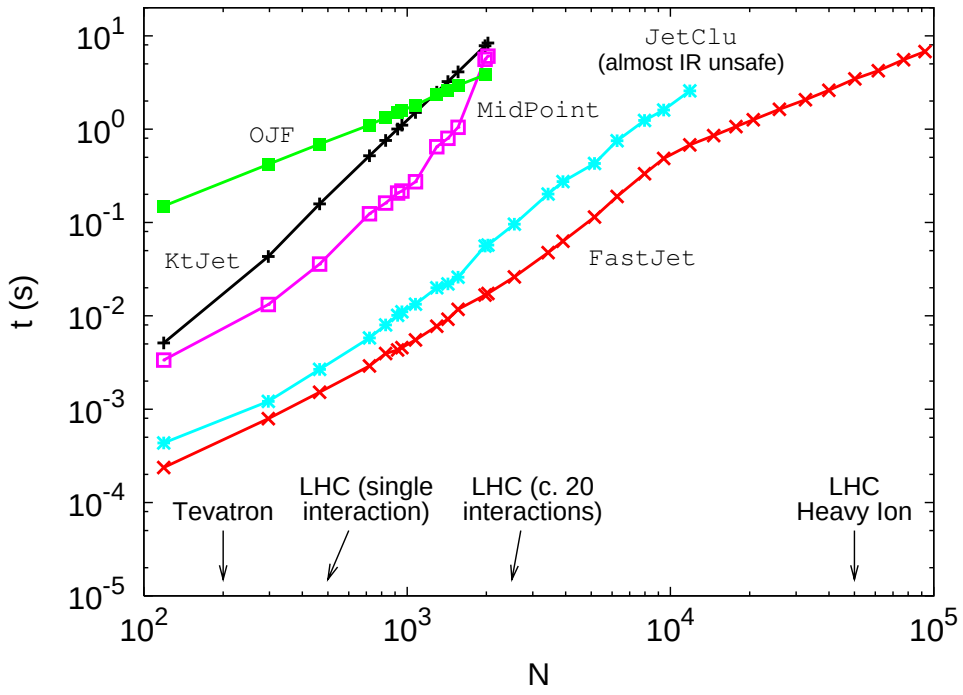

Figure 5.4: *Comparison of the running time of different jet clustering algorithms versus the number of input objects. Taken from [117].*

Actual Jet Definition for this study

In CMS, two different types of objects are used as input collection for the jet algorithms. The first type comprises all hadronic final state particles which are listed in the Monte Carlo truth information. This includes also neutrinos and muons from resonance decays. Clustering these objects leads to particle jets and are called GenJets in CMS.

The same algorithms can be applied on energy deposits in the CMS detector. These calorimeter jets are reconstructed based on energy deposits in calorimeter towers which are composed of one or more cells of the hadron calorimeter and the corresponding crystals of the electromagnetic calorimeter. For the barrel region of the CMS detector, a calorimeter tower is the unweighted sum of a single HCAL cell and the associated ECAL supercrystal. For the forward region of the detector, a more complex association between HCAL and ECAL is required [118]. The thresholds for noise suppression in the calorimeters are summarised in Table 5.2. Jets which have been clustered from this list of input objects are called calorimeter jets or CaloJets in CMS, see Figure 5.5.

The recombination scheme and the specific parameters for the different algorithms are listed in Appendix B. The cone size R or the resolution parameter D is specified in the name of the algorithm. For example the SISCone algorithm with a cone size of 0.5 is referred to as SISCone 0.5.

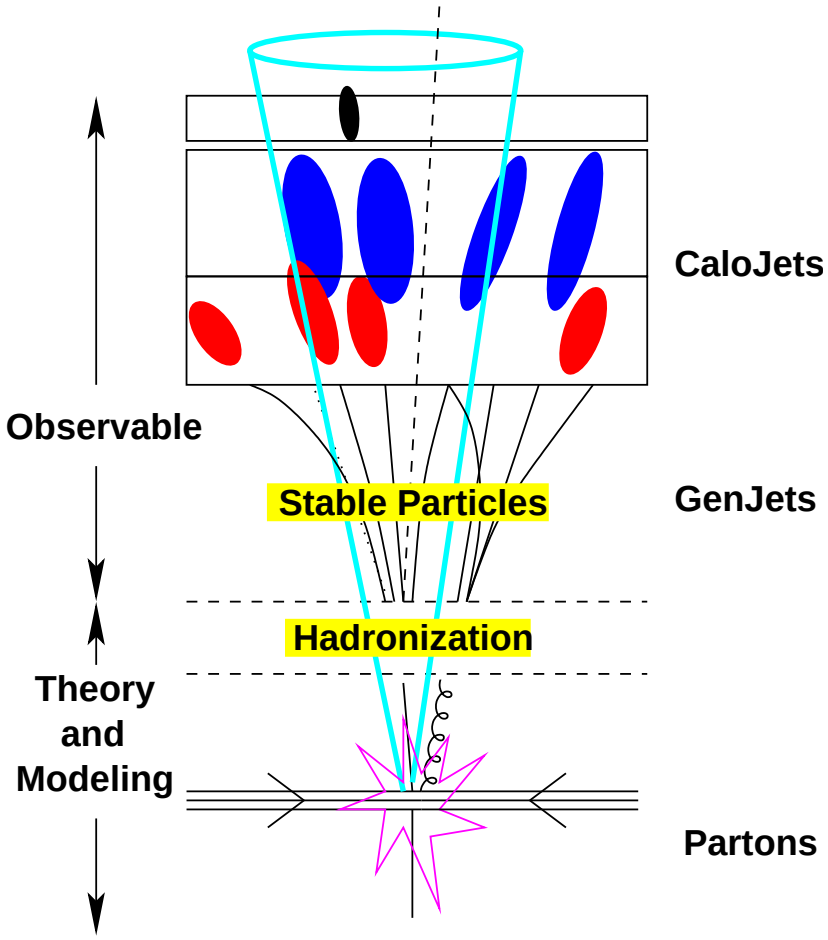

Figure 5.5: *Visualisation of the final state definition in CMS. Jets clustered from the hadronic final state based on Monte Carlo truth information are called particle jets or GenJets in CMS. Taking calorimeter information as input objects leads to calorimeter jets or CaloJets [112]*

Calorimeter	Threshold [GeV]
HB	0.90
HO	1.10
HE	1.40
$\sum$EB	0.20
$\sum$EE	0.45

Table 5.2: *Energy thresholds for noise suppression of the different calorimeters. $\sum$EB and $\sum$EE denote the sum of the energy deposit of all ECAL cells which are associated to a HCAL cell. These values are referred to as Scheme B in CMS [118].*

5.5 Physics Analysis Tools

ROOT

ROOT [119] is an object-oriented framework which is widely used in high energy physics. The project was started in 1995 and is officially supported by CERN since 2003. Written in C++ it offers a large collection of classes useful for different kinds of analyses. This includes four-vectors, geometry packages and a high-performance input/output system as well as functionality for histogramming, fitting routines and parallel data processing.

Jet Calibration using Z Balancing

Jet clustering algorithms are designed to group all particles originating from the same parton into a single object, the jet. However, the reconstructed energy of the jet does not correspond to the energy of the initial parton due to various effects. On the one hand, the performance of the energy measurement of the detector has to be taken into account, such as electronic noise of the calorimeter cells, threshold effects, absorber material and the non-linear response of the calorimeter. On the other hand, aspects related to the jet reconstruction itself have to be considered. This includes the fragmentation model, initial- and final-state radiation, the multiple parton interactions and particles coming from pile-up.

The difference between the transverse energy of a calorimeter jet compared to the corresponding particle jet is presented in Figure 6.1. In general, the energy measured in the calorimeter is underestimated compared to the particle jet and depends on its transverse momentum and pseudorapidity.

However, many physics analyses are interested in the properties of the originating parton or in comparisons of the measured events with Monte Carlo simulations of the hadronic final state. Therefore, to compare calorimeter jets with the original parton or particle jets, corrections of the measured energy have to be applied. The CMS collaboration envisages a factorised multi-level jet calibration procedure which is briefly described in the following.

6.1 Factorised Multi-Level Jet Correction

The overall jet calibration is subdivided into several steps, which are referred to from here on as calibration levels. A measure for the deviation of the transverse momentum of a calorimeter jet

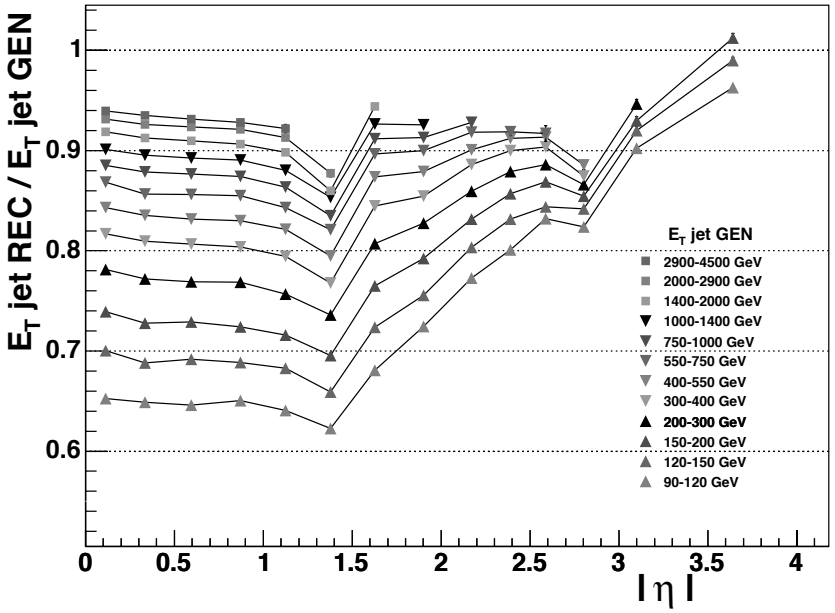

Figure 6.1: *Ratio of the transverse energy of uncalibrated calorimeter (REC) and particle (GEN) jets as a function of the generated absolute pseudorapidity for different jet energies clustered by the Iterative cone 0.5 algorithm [18].*

from the original parton or the particle jet is the response R which is defined as

$$R = \frac{p_{\mathrm{T}}^{\mathrm{calo}}}{p_{\mathrm{T}}^{\mathrm{gen}}}. \tag{6.1}$$

Here, $p_{\mathrm{T}}^{\mathrm{calo}}$ denotes the transverse momentum of the calorimeter jet and $p_{\mathrm{T}}^{\mathrm{gen}}$ the transverse momentum of either the corresponding particle jet or the original parton.

At present, seven calibration levels are foreseen, starting with uncalibrated calorimeter jets and ending with calibrated jets according to the requirements of a physics analysis [120]. The different levels are schematically drawn in Figure 6.2 and listed below:

1. **Offset:** corrections for pile-up, electronic noise and the jet energy loss due to the calorimeter thresholds.

2. **Relative:** corrections for variations in the jet response with pseudorapidity relative to a control region.

3. **Absolute:** corrections of the transverse energy of a calorimeter jet to particle jets in the control region.

4. **EMF:** corrections of the jet response taking into account the electromagnetic energy fraction.

Figure 6.2: *Schematic overview of the factorised multi-level jet correction in CMS. Starting with an uncalibrated calorimeter jet, several subsequent corrections can be applied. In this scheme, required correction levels are indicated with solid boxes [120].*

5. **Flavour:** correction to particle jets in dependence of the flavour of the original parton (light quarks, c, b or gluon).

6. **Underlying Event:** removes the energy of the underlying event contributing to a jet.

7. **Parton:** corrects the jet back to the energy of the parton, which includes hadronisation corrections and out-of-cone effects.

The advantage of such a factorised approach, which is also pursued by the Tevatron experiments, is that it allows to determine, refine and understand the corrections of the different levels almost independently. Additionally, systematic uncertainties can be investigated for each level, which yields a better overall understanding of the origins of systematic uncertainties on the jet energy scale.

In this context, the first three levels are required whereas the remaining are optional corrections, offering calibrated jets according to the requirements of the final physics analyses. In addition, such a modular approach enables the jet energy calibration to evolve with time and the available amount of data. With the beginning of data taking, mainly Monte Carlo driven methods will be applied, but with increasing statistics and an improved understanding of the detector, data-driven calibrations will become available.

In the next section, the first three calibration levels are discussed in detail focussing on both, Monte Carlo based and data-driven techniques. In addition, the basic aims of the optional correction levels are briefly presented.

6.1.1 Offset Corrections

The first level of jet calibration is intended to correct for the effects of pile-up, electronic noise and the threshold of the calorimeter cells. It is planned to measure these corrections directly from the early collision data using zero-bias and minimum-bias events.

The default for the data acquisition is to record the information of an event according to defined trigger criteria. Although this restriction is necessary to select only the most interesting events, it introduces a certain bias. Therefore, dedicated runs with a reduced number of bunches per beam

are planned and the read-out of the detector is triggered by the bunch crossing. Therefore, all events are recorded and called zero-bias.

However, the zero-bias data will contain events with no activity in the detector. To suppress such empty events, minimal trigger constraints are defined which demand a certain activity to record an event. The criteria are designed to keep the bias on the selection as small as possible. Therefore, these data samples are referred to as minimum-bias events.

The zero-bias data will be used to measure the offset related to the noise of the calorimeter and pile-up events from the early collision data. In these collisions, the energy deposit in a jet area is measured as a function of the pseudorapidity by placing cones at fixed values of the pseudorapidity and randomly in the azimuthal angle. The $\mathrm{k_T}$ and SISCone algorithms additionally provide techniques for pile-up estimation and can therefore be used to check or improve this offset correction. Also the measurement of the energy density in the calorimeter as a function of η, which is assumed to be equal to the energy offset of the jet divided by the jet area, is a proper quantity. Having estimated the energy gain through pile-up, the jet energy can be corrected for this effect.

To estimate the energy loss due to the threshold of the cells and the towers, these thresholds are turned off for some runs to store the whole calorimeter information of an event. The combination of both subcorrections finally yields the Level 1 correction in this scheme.

Up to now, both subcorrections have been investigated using simulations and will be derived by the procedures discussed above with the first collision data. The corrections for this level were not available at the timescale of this study and are therefore not included. However, the investigated data samples do not contain pile-up events which indicates, that the Level 1 correction factors would be small.

6.1.2 Relative Corrections

Due to the detector geometry, the reconstructed energy of calorimeter jets and therewith the correction depends on the pseudorapidity. The goal of the Level 2 correction (L2) is to arrive at a uniform response in η. This is done by correcting jets at arbitrary values of the pseudorapidity relative to a control region, which is defined to be the central barrel region of the CMS detector with $|\eta| < 1.3$.

The central barrel part of the detector is chosen as preferred control region as

- it is the best understood part of the calorimeter to calibrate in absolute terms,
- it provides the final states with the highest p_T reach and
- the response varies only little and smoothly with the pseudorapidity.

Assuming a flat response with respect to η and ϕ is the basis for correction factors of the absolute jet energy scale which do not depend on any angular variable.

The currently available Level 2 correction is based on Monte Carlo truth information and is used for this study. With the first LHC collisions, this calibration will be supplemented by data-driven methods.

Monte Carlo based Correction

In the current implementation, the Monte Carlo truth based L2 correction is derived by comparing the transverse momentum of particle and corresponding calorimeter jets. The matching is performed by requiring a minimum distance in the η-ϕ-plane for the particle and calorimeter jet. This is done by minimising the geometrical quantity ΔR. Those pairs of jets which do not satisfy $\Delta R < 0.25$ are not considered for the determination of the response

$$R(\eta, < p_{\mathrm{T}}^{\mathrm{gen}} >) = 1 + \frac{< \Delta p_{\mathrm{T}} >}{< p_{\mathrm{T}}^{\mathrm{gen}} >} \quad \text{with} \quad \Delta p_{\mathrm{T}} = p_{\mathrm{T}}^{\mathrm{calo}} - p_{\mathrm{T}}^{\mathrm{gen}} \tag{6.2}$$

which depends on the mean of the transverse momentum of the particle jet, $< p_{\mathrm{T}}^{\mathrm{gen}} >$, and its pseudorapidity. The response is determined in 82 bins of η covering a region up to $|\eta| < 5.191$ and 25 bins of the transverse momentum of the particle jets up to 6 TeV.

The relative correction factor is defined as

$$C(\eta, p_{\mathrm{T}}^{\mathrm{calo}}) = \frac{p_{\mathrm{T}}^{\mathrm{control}}}{p_{\mathrm{T}}^{\mathrm{calo}}}. \tag{6.3}$$

Here, $p_{\mathrm{T}}^{\mathrm{control}}$ and $p_{\mathrm{T}}^{\mathrm{calo}}$ denote the measured transverse momentum of a particle jet with the same $p_{\mathrm{T}}^{\mathrm{gen}}$ in the control region or an arbitrary region of η, respectively. The correction factor is deduced by the procedure described below:

1. Selection of a p_{T} and η.
2. Numerical inversion of the equation $p_{\mathrm{T}}^{\mathrm{calo}} = R(\eta, p_{\mathrm{T}}^{\mathrm{gen}}) \cdot p_{\mathrm{T}}^{\mathrm{gen}}$ to get $p_{\mathrm{T}}^{\mathrm{gen}}(p_{\mathrm{T}})$.
3. Calculation of $p_{\mathrm{T}}^{\mathrm{control}}(p_{\mathrm{T}}^{\mathrm{gen}})$ from the equation $p_{\mathrm{T}}^{\mathrm{control}} = R(\eta^{\mathrm{control}}, p_{\mathrm{T}}^{\mathrm{gen}}) \cdot p_{\mathrm{T}}^{\mathrm{gen}}$.
4. Determination of the relative correction factor according to Equation 6.3.

Data-driven Correction

Following the experiences of the Tevatron experiments, the Level 2 correction factors can be derived from dijet events once collision data are available. Assuming a 2→2 hard process with two

outgoing partons, the transverse momentum of the jets has to be balanced in the absence of hard QCD radiation. The relative jet energy correction factor $C(\eta, p_{\mathrm{T}}^{\mathrm{calo}})$ can be determined exploiting this balance.

In an event, the two calorimeter jets with the highest transverse momentum are selected. From here on, the transverse momentum of a calorimeter jet with a pseudorapidity inside or outside the control region is denoted with $p_{\mathrm{T}}^{\mathrm{control}}$ and $p_{\mathrm{T}}^{\mathrm{probe}}$, respectively. If both jets satisfy $|\eta| < 1.3$, the probe jet is chosen randomly. In this notation, the measure B of the dijet balance is calculated as

$$B = \frac{p_{\mathrm{T}}^{\mathrm{probe}} - p_{\mathrm{T}}^{\mathrm{control}}}{p_{\mathrm{T}}^{\mathrm{dijet}}} \tag{6.4}$$

with

$$p_{\mathrm{T}}^{\mathrm{dijet}} = \frac{p_{\mathrm{T}}^{\mathrm{probe}} + p_{\mathrm{T}}^{\mathrm{control}}}{2} \tag{6.5}$$

where $p_{\mathrm{T}}^{\mathrm{probe}}$ is recorded in bins of $p_{\mathrm{T}}^{\mathrm{dijet}}$ and η^{probe}.

The relative jet energy response from this balancing is expressed by

$$r(\eta^{\mathrm{probe}}, p_{\mathrm{T}}^{\mathrm{dijet}}) = \frac{2+ < B >}{2- < B >} \tag{6.6}$$

and is conceptually equivalent to $r =< p_{\mathrm{T}}^{\mathrm{probe}}/p_{\mathrm{T}}^{\mathrm{control}} >$. The final correction factor $c = 1/r$ is expressed in terms of $< p_{\mathrm{T}}^{\mathrm{probe}} >$ for the same $p_{\mathrm{T}}^{\mathrm{dijet}}$ and η^{probe} bin.

A comparison of the derived relative correction factors shows good consistency with those of the Monte Carlo truth matching, discussed above. Both methods lead to a flat response of the calorimeter jets with respect to the control region in which the Level 3 correction will be derived. Besides the dijet balancing, also momentum conservation of the whole event can be exploited for a L2 correction and is described in detail in [120].

6.1.3 Absolute Correction of the Transverse Momentum

The goal of the Level 3 correction (L3) is to remove the variation of the calorimeter jet response of the detector as a function of the measured transverse momentum of the jet. These variations mainly result from the non-linear response of the calorimeter. Beside Monte Carlo truth based correction factors, which are derived from a QCD dijet sample, several physics processes permit to extract data-driven calibration factors. Both procedures are briefly discussed in the following. To reduce systematical and statistical uncertainties it is planned to combine the different Level 3 calibrations.

The advantage of this factorised approach is that the different calibration techniques can evolve with the available number of events as well as the understanding of the detector and can be easily

replaced. With the beginning of data taking, only L3 corrections based on Monte Carlo truth information will be available to calibrate the jets of the first data. With increasing statistics, the data-driven techniques using the balancing of a jet by a photon or a Z boson will arise and finally replace the Monte Carlo truth corrections. Once the statistical precision is sufficient, also QCD three jet events may be used to extend the correction to jets with higher transverse momenta.

However, it is important to consider the different levels of correction. Whereas the photon and Z boson balancing correct back to the transverse momentum of the original parton, the Monte Carlo truth based calibration corrects to the particle jet level. This difference has to be taken into account for the combination of the calibrations.

Monte Carlo based Correction

In the current implementation, the Monte Carlo truth based L3 correction is derived by comparing the transverse momentum of particle and corresponding calorimeter jets in the control region $|\eta| < 1.3$. The matching of the jets is again based on a purely geometrical criterion. If the two particle jets with the highest transverse momentum in an event fall in the control region, they are matched to the calorimeter jets by requiring a minimum distance in the η-ϕ-plane by minimising the geometrical quantity ΔR. Those pairs of jets which do not satisfy $\Delta R < 0.25$ are not considered for the calibration.

The difference of the transverse momentum of the particle and calorimeter jet

$$\Delta p_{\mathrm{T}} = p_{\mathrm{T}}^{\mathrm{calo}} - p_{\mathrm{T}}^{\mathrm{gen}} \tag{6.7}$$

is recorded in 25 bins of the transverse momentum of the particle jet. Due to the non-Gaussian tails of this distribution, its peak value is considered and determined using a Gaussian fit in the range of $\pm 1.5\sigma$ around the maximum and are plotted versus the mean of the corresponding $p_{\mathrm{T}}^{\mathrm{gen}}$ bin.

However, the response defined this way is a function of the transverse momentum of the particle jet, but the calibration factors have to depend on transverse momentum of the calorimeter jets. Therefore, the response curve is inverted numerically to get $p_{\mathrm{T}}^{\mathrm{gen}}$ as a function of $p_{\mathrm{T}}^{\mathrm{calo}}$. A fit of this distribution yields the corresponding calibration.

To offer calibration factors for jets up to the highest transverse momenta, a large number of QCD dijet events have been simulated in bins of the hard transverse momentum $\hat{p_{\mathrm{T}}}$. For the determination of the calibration factors, these events have been used without applying weights according to their cross-section. From here on, this calibration is referred to as Level 3-dijet calibration or $\mathrm{L3}^{\mathrm{dijet}}$. A more detailed description of this method is presented in [121].

Data-driven Correction

With the increasing number of recorded events and the understanding of the collision data, events containing a photon or a Z boson can be used to determine calibration factors for the absolute jet energy scale. Both methods are based on the concept that the momentum of a parton of the hard subprocess is balanced by a photon or a Z boson.

Whereas the measurement of the jet transverse momentum is linked to both, the electromagnetic and the hadron calorimeter, the determination of the energy of the photon only requires information from the ECAL. This detector component can be calibrated using electron test beams supplemented by in situ calibration using $\pi^0 \rightarrow \gamma\gamma$ or $Z \rightarrow e^+e^-$ decays.

Therefore, a precise measurement of the transverse momentum of the photon is possible which can be related to the momentum of the balanced parton. The latter can be compared with the measured energy of the balanced calorimeter jet. However, this process is affected by a large background from QCD dijet events in which one jet is misidentified as a photon. In order to enhance the signal to background ratio, strong isolation criteria on the photon are applied, reducing the achievable statistical precision.

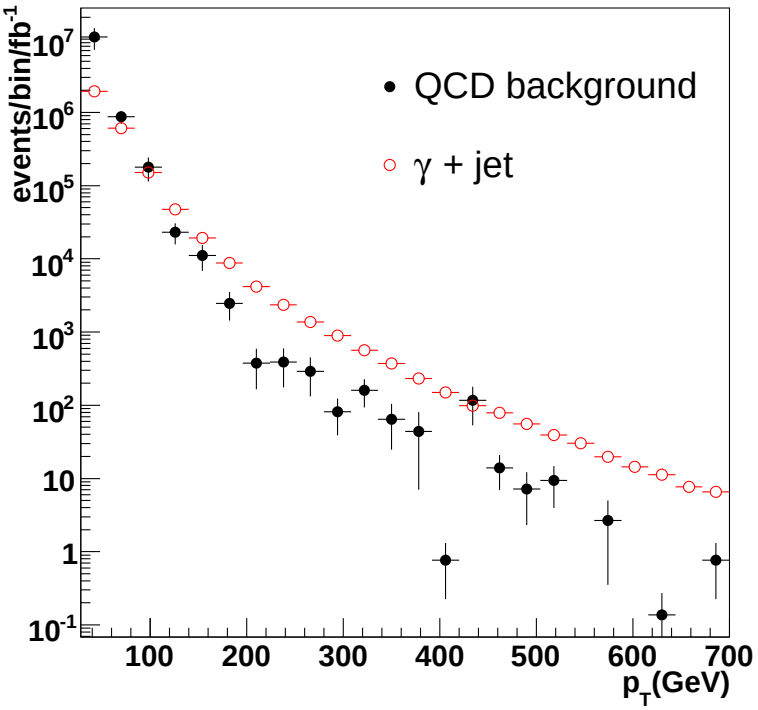

Figure 6.3: *Expected number of events available for the photon and jet balancing including the QCD background. Especially in the region of low transverse momenta the number of background events is comparable to the signal [120].*

The expected number of events for the signal sample as well as for the background are shown in Figure 6.3. Although the background is strongly suppressed by the selection, it remains in the same order of magnitude as the signal and a determination of calibration factors for the region of lower transverse momenta up to 250 GeV is difficult. More details on the influence of the background on the determination of the calibration factors as well as the event selection are given in [120].

The calibration of the absolute jet energy scale using the balancing of a Z boson was not possible at current experiments due to the relatively small number of such events, leading to large statistical uncertainties. However, at the LHC design parameters, a large number of Z bosons will be produced and can be used for jet calibration following the argumentation for the photons. In principle, the decay of the Z boson into a pair of electrons or muons is suitable for such purpose, but the latter has several advantages.

The CMS detector offers an excellent coverage and reconstruction of muons for the precise determination of the kinematics of the Z boson. In addition, only information from the tracker and the muon chambers are considered for the reconstruction, which provides a measurement of the transverse momentum of the balanced Z boson without relying on calorimetric information. Both, a precise understanding of the tracker as well as of the muon system will be achieved with the first data. Furthermore, possible background processes can be suppressed due to the clean signature of this decay leading to a negligible background for such analyses. Although the available statistics is smaller for the Z boson balancing compared to the photon balancing by nearly one order of magnitude as presented in Figure 6.4, the precise reconstruction of the boson kinematics combined with the negligible background make this calibration the candidate of choice for the region of low transverse jet momenta, as discussed in detail in the sections below.

To decrease the statistical uncertainty, also the decay of the Z boson into two electrons will be considered for calibration aspects as soon as the energy scale of the electromagnetic calorimeter is well understood.

6.1.4 Optional Corrections

In addition to the previous discussed required calibration levels, four optional corrections are envisaged in CMS and briefly presented in this section. A more detailed description can be found in [120].

Electromagnetic Energy Fraction

The actual response of the detector can be divided into the responses for various types of particles and therewith in a separate response for the electromagnetic and the hadron calorimeter. To improve the jet resolution, additional correction factors based on the fraction of the jet energy deposited in the electromagnetic calorimeter (EMF) are envisaged to be applied after the calibration

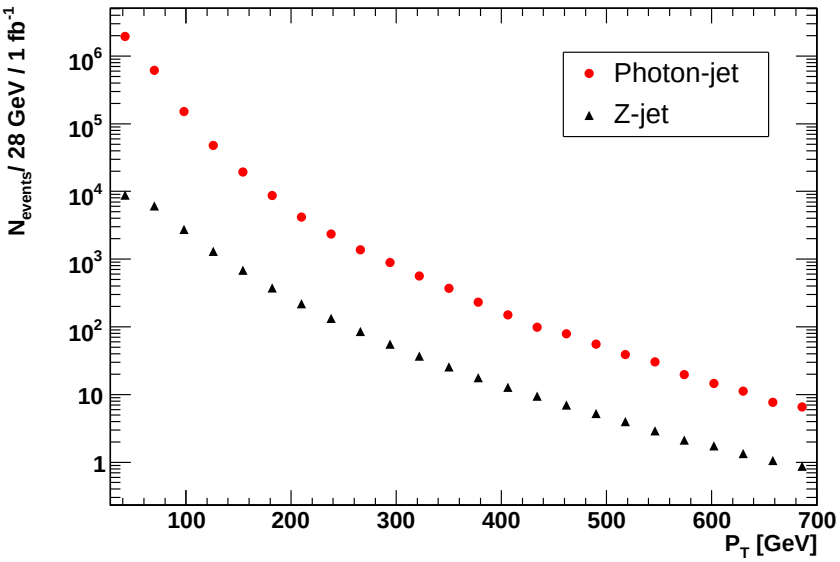

Figure 6.4: *Expected number of events after selection cuts for the photon and Z boson balancing. For the Z boson only the decay into two central muons is considered [120].*

of the absolute jet energy scale. Up to now, Monte Carlo based calibration factors are available, and it has been demonstrated that improvements of the jet resolution between 5% and 10% are feasible.

Jet Flavour

The calorimeter response is different for jets originating from gluons, light or heavy quarks. Compared to the light quark jets, the calorimeter energy response for c and b jets is smaller due to differences in the jet fragmentation and gluon jets are broader as a result of the higher colour-charge of the gluon. Especially analyses which are able to indentify the flavour of a jet in the final state will benefit from such flavour specific correction factors on top of the required calibrations. This correction is referred to as Level 5 in the factorised approach and estimates from Monte Carlo truth information are already available.

Underlying Event

The optional Level 6 calibration is designed to remove the energy offset of the jet coming from the underlying event. Conceptually, the underlying event is the component of the proton-proton interaction that does not originate from the hard parton scattering and is assumed to be luminosity independent and uniformly distributed in an event. Although the underlying event depends on the

details of the hard interaction, it is planned to provide a generic underlying event correction which can be used for all analyses. In the early data, this correction will be determined from the energy per jet area from minimum bias events at low energy after having subtracted the Level 1 offset correction.

Parton Level

Several of the previously discussed calibrations correct jets back to particle jets. The parton level correction (L7) attempts to correct these jets back to the properties of the original parton. This step will also be useful for the combination of available correction factors within one calibration level which may correct back to either particle jets or the parton level. However, it is important to emphasise that this calibration level is model and process dependent and relies strongly on the Monte Carlo generator.

6.2 Z Boson and Jet Balancing

As discussed in the previous chapters, a large number of events containing Z bosons will be available at the LHC and can be used for high precision electroweak measurements, monitoring of the luminosity of the collider and the commissioning of different detector components. Furthermore, they are well suited to evaluate calibrations already existing in CMS and finally to determine correction factors for the absolute jet energy scale.

For the purpose of calibration of the jet energy scale, processes are required in which a Z boson is balanced by exactly one jet of the hard subprocess. As the momentum is conserved in such an event topology, the kinematics of the balanced parton is linked to the Z boson momentum, which can be reconstructed with high precision. Thus, the calibration of the absolute jet energy scale can be achieved by comparing the momentum of the boson with the balanced calorimeter jet.

To be independent from any calorimeter information for the reconstruction of the Z boson, only the decay into muons is considered here. For both, the tracking system as well as the muon chambers, a good calibration will be available even with the first data. In addition, a clean sample of events containing a Z boson and exactly one additional jet is required for the comparison.

In this study, these Z boson plus one jet events are used for two purposes. First, existing calibrations in CMS are investigated using the method of Z boson balancing. In a second step, calibration factors for the absolute jet energy scale are determined. In absence of collision data, this is currently done with simulated Monte Carlo events. However, the analysis is designed in such a way that a calibration from Z boson events can be derived with the first collision data.

6.2.1 Monte Carlo Datasets

The study is performed on data samples taken from the official Monte Carlo production done by the CMS collaboration. The Pythia event generator was used to generate the signal process with a Z boson and one parton in the hard interaction according to the Feynman diagrams presented in Figure 1.6. This event topology is selected by activating the Pythia process parameters

$$\begin{aligned} \texttt{MSUB(15)} = \texttt{1} \quad &: \quad \mathrm{f}_i\bar{\mathrm{f}} \rightarrow \mathrm{g}(\gamma^*/Z) \\ \texttt{MSUB(30)} = \texttt{1} \quad &: \quad \mathrm{f}_i\mathrm{g} \rightarrow \mathrm{f}_i(\gamma^*/Z). \end{aligned}$$

To offer a sufficient number of events with large transverse momenta of the Z boson, the generation has been divided up into 21 bins of the hard transverse momentum $\hat{p_{\mathrm{T}}}$ using selection cuts in Pythia during the event generation. For example, the parameters `CKIN(3)` and `CKIN(4)` restrict the lower and upper range of the generated transverse momentum of the hard interaction. Due to initial and final state radiation, the data samples contain at least one jet and the Z boson which has been chosen to decay only into muons. The latter is requested by `MDME(184,1)=1`.

The production of the data samples has been performed in two steps for practical reasons. For the generation and subsequent detector simulation the CMS software version CMSSW_1_4_6 has been used. The digitisation and reconstruction was done, however, with the newer release CMSSW_1_6_7. Both steps have been performed on grid resources and the data sets are only accessible via grid techniques. The $\hat{p_{\mathrm{T}}}$ ranges as well as the number of fully simulated events and the corresponding cross-sections are presented in Table 6.2. The configuration file for the event generation and simulation is additionally given in Appendix C.1.

6.2.2 Z Boson Event Selection

The first step of the analysis consists of the reconstruction of the kinematics of the Z boson from its decay products, which are the two muons in this case. As discussed in Section 5.4.1 the muon reconstruction offers different options. For this analysis, a precise reconstruction of the momentum and the charge of the muons is required. Therefore, the global muon collection is used, since for the reconstruction the hits in the muon system are combined with information from the tracker. To avoid a dependence on trigger thresholds, border effects of the detector and noise, additional selection criteria are applied on the muon collection.

The CMS detector provides an efficient and precise reconstruction of global muons up to values of the pseudorapidity of $|\eta| < 2.4$, which is restricted by the coverage of the muon system. In order not to be influenced by border effects, tighter cuts than the detector constraints are applied for the final analysis. Requiring the pseudorapidity of the muons to hold $|\eta| < 2.3$ ensures a sufficient distance to the border region of the detector and does not decrease the number of remaining events

$\hat{p_{\mathrm{T}}}$ [GeV]	Number of events	Cross-section [pb]
0 - 15	15000	1.044E+05
15 - 20	25000	3.504E+02
20 - 30	50000	3.267E+02
30 - 50	50000	2.270E+02
50 - 80	50000	9.317E+01
80 - 120	25000	3.148E+01
120 - 170	25000	9.630
170 - 230	25000	2.920
230 - 300	25000	8.852E-01
300 - 380	25000	2.936E-01
380 - 470	25000	1.025E-01
470 - 600	25000	4.242E-02
600 - 800	15000	1.443E-02
800 - 1000	15000	2.859E-03
1000 - 1400	15000	9.400E-04
1400 - 1800	15000	9.536E-05
1800 - 2200	15000	1.232E-05
2200 - 2600	15000	1.839E-06
2600 - 3000	15000	2.881E-07
3000 - 3500	15000	4.764E-08
3500 - inf	15000	4.515E-09

Table 6.1: *Listing of the bins of the hard transverse momentum used for the event generation with Pythia including the number of events and the corresponding cross-section.*

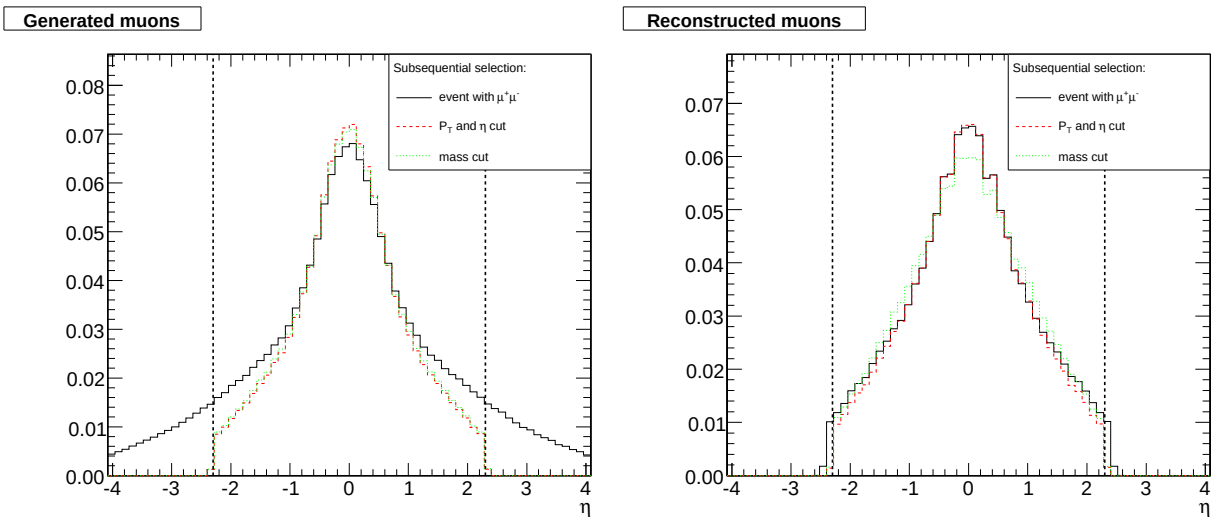

Figure 6.5: *Normalised distribution (within $|\eta| < 2.3$) of the pseudorapidity of the two muons with an invariant mass closest to the mass of the Z boson for both, the generated and the reconstructed global muon collection. A first scenario does not apply further constraints to the muons. For a second scenario only the cut on the transverse momentum and the pseudorapidity is enabled followed by additionally constraining the invariant mass of the di-muon system. The dashed lines indicate the chosen cut of $|\eta| < 2.3$. The different $\hat{p_T}$ bins are combined without weighting for their different cross-section.*

significantly.

In addition, the CMS High Level Trigger System offers a dedicated path for Z boson event candidates with a subsequent decay of the boson into two muons. To be accepted by the HLT, both muons must satisfy $p_T > 7\,\text{GeV}$ [122]. To suppress noise, the cut on the transverse momentum of the muons for measurements of the process $Z \rightarrow \mu^+\mu^-$ is further constrained to 15 GeV for this analysis.

To reject background processes and to select a clean sample of events, the Z boson is reconstructed as follows. Starting with the global muons which hold all selection criteria, the invariant mass of all possible pairs of muons with opposite charge is calculated. The di-muon system with an invariant mass closest to the mass of the Z boson is selected and only events with an invariant mass of the di-muon system closer than 20 GeV to the Z boson mass are considered from here on.

In addition the same cut scenario is investigated using Monte Carlo truth information for comparisons. For this, the event reconstruction and selection discussed above are applied to the muons as they are calculated by the event generator. In the following, this collection is denoted as generated muons and consists of all stable muons, which are indicated by the Pythia flags `pdgid=±13` and `status=1`.

The normalised distribution of the pseudorapidity of the two muons with an invariant mass of the di-muon system closest to the mass of the Z boson is shown in Figure 6.5 for generated and reconstructed muons. As a weighting of events from different $\hat{p_T}$ bins according to their cross-section

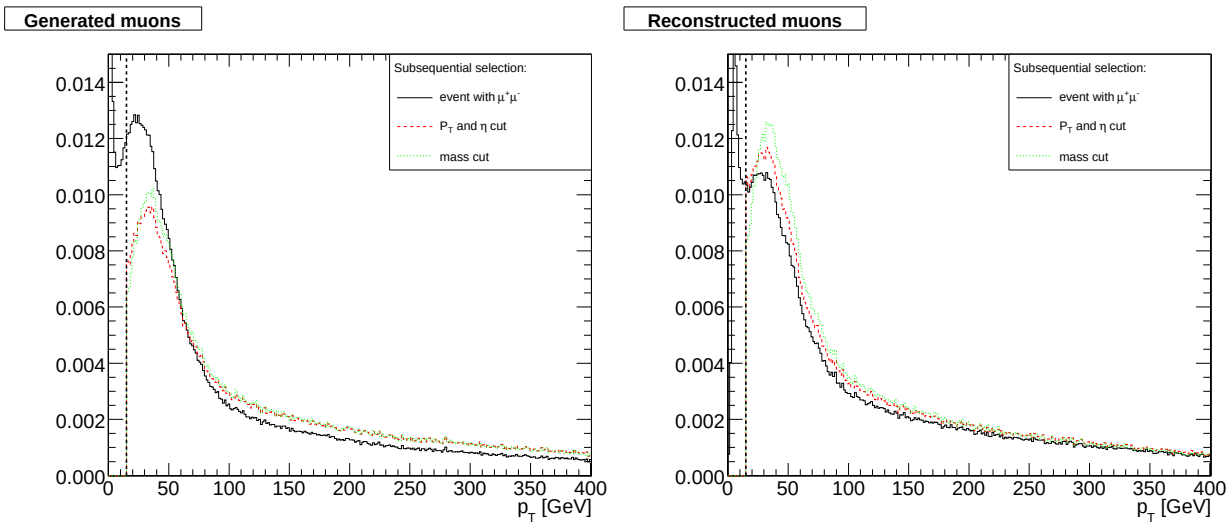

Figure 6.6: *Same as Figure 6.5 but for the transverse momentum of the muons. The dashed line indicates the cut $p_{\mathrm{T}} > 15\,\mathrm{GeV}$.*

would lead to the effect that events from a few bins would completely dominate the distributions, the cross-section difference is not considered in these figures.

If all events with two muons with opposite charge are selected, the main difference between generated and reconstructed muons is related to the detector acceptance. While the generated muons are not limited in η, only muons with $|\eta| < 2.4$ can be reconstructed. It also demonstrates that the tighter cut on the pseudorapidity of the muon of 2.3 (dashed line) affects the number of remaining events only slightly.

Both collections contain a large number of muon pairs with opposite charge and very small p_{T} as shown in Figure 6.6. As the distribution of transverse momentum of the reconstructed muons has a local minimum at about 15 GeV and increases strongly for smaller transverse momenta, this value has been chosen as cut-off.

In Figure 6.7, the invariant mass of the di-muon system is shown including the cut on the transverse momentum and the pseudorapidity. The shape of the distribution is comparable for reconstructed and generated muons and the applied cuts are indicated again by the dashed lines.

The expected number of reconstructed events containing a Z boson and at least one jet after the selection criteria for an integrated luminosity of $1\,\mathrm{fb}^{-1}$ is presented in Figure 6.8. The black bars indicate the statistical uncertainty whereas the red boxes show the uncertainty related to the number of simulated events available. Even for transverse momenta of the Z boson of about 400 GeV a large number of such events will be available at the LHC.

In addition, the background to this process has been investigated in detail in [123] and found to be negligible as shown in Figure 6.9. However, there are small differences with respect to the event

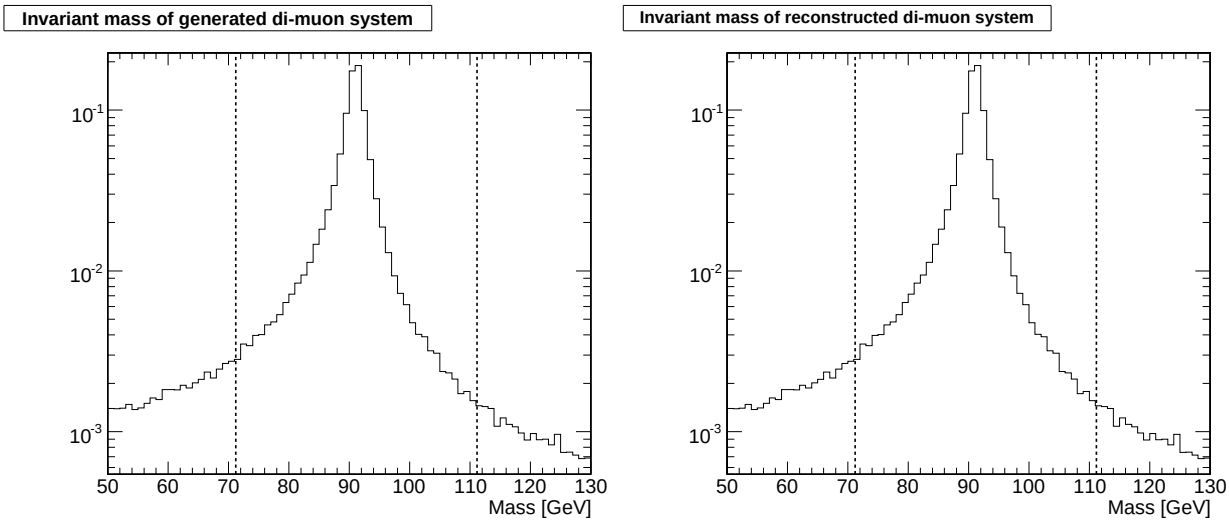

Figure 6.7: *Normalised distribution of the invariant mass M of the di-muon system for both, generated and reconstructed muons including all selection criteria. The dashed lines enclose the region of* $|\mathrm{M} - \mathrm{M_Z}| < 20\,\mathrm{GeV}$.

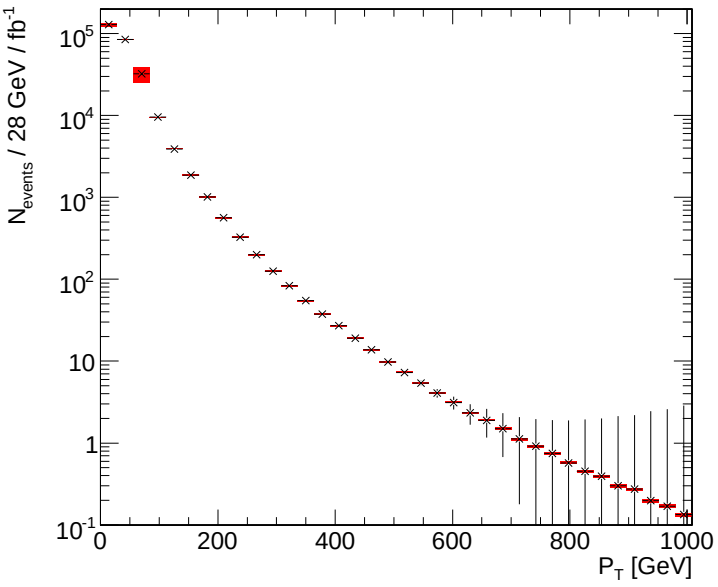

Figure 6.8: *Expected number of reconstructed events containing a Z boson and at least one jet after the selection criteria for an integrated luminosity of 1* fb^{-1}. *The black bars indicate the statistical uncertainty whereas the red boxes show the uncertainty related to the number of simulated events.*

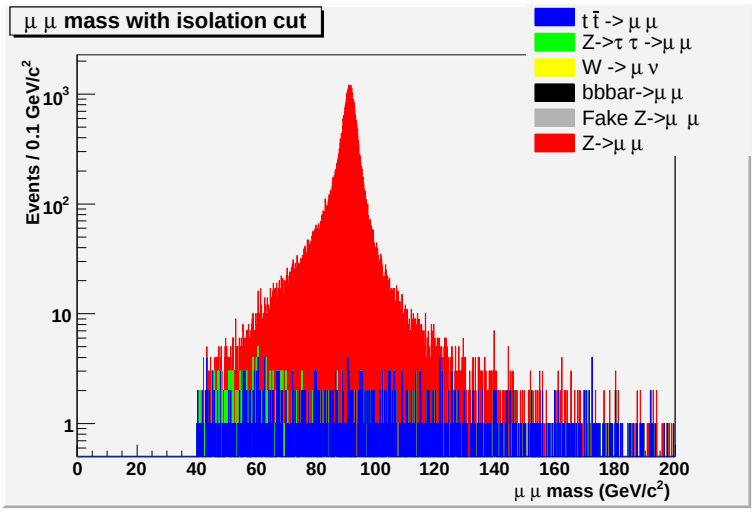

Figure 6.9: *Number of reconstructed Z boson candidates including the expected number of background events for the process $Z \rightarrow \mu^+\mu^-$ [123].*

selection since the cut on the transverse momentum of the muons is 20 GeV. The relaxation of this cut to 15 GeV as used in this study is assumed to not influence the number of background events significantly. In addition, the background study requires that the sum of the transverse momentum of all tracks is less than 3 GeV in a cone of $\Delta R < 0.3$ around the muon. A detailed investigation has shown that the influence of this cut on the remaining number of events for this analysis is less then 0.5% in the region of interest of the transverse momentum.

Name	Value
Generated muons	Pythia stable muons (status = 1)
Reconstructed muons	global muons
Transverse momentum	$p_{\mathrm{T}} > 15\,\mathrm{GeV}$
Pseudorapidity	$\lvert\eta\rvert < 2.3$
Invariant mass M	$\lvert\mathrm{M} - \mathrm{M_Z}\rvert < 20\,\mathrm{GeV}$

Table 6.2: *Summary of the selection and reconstruction criteria for the Z boson.*

6.2.3 Event Selection for Z Boson Jet Balancing

For the calibration of the absolute jet energy scale, only events are considered in which the jet with the highest transverse momentum falls into the control region defined as the central barrel region of the CMS detector with $|\eta| < 1.3$. In addition, to exploit momentum conservation and so to infer from the kinematics of the Z boson on the momentum of the original parton, it is necessary to restrict the analysis to events where the Z boson is balanced by exactly one jet of the hard process. Therefore, a clean selection of the events with respect to these properties is required.

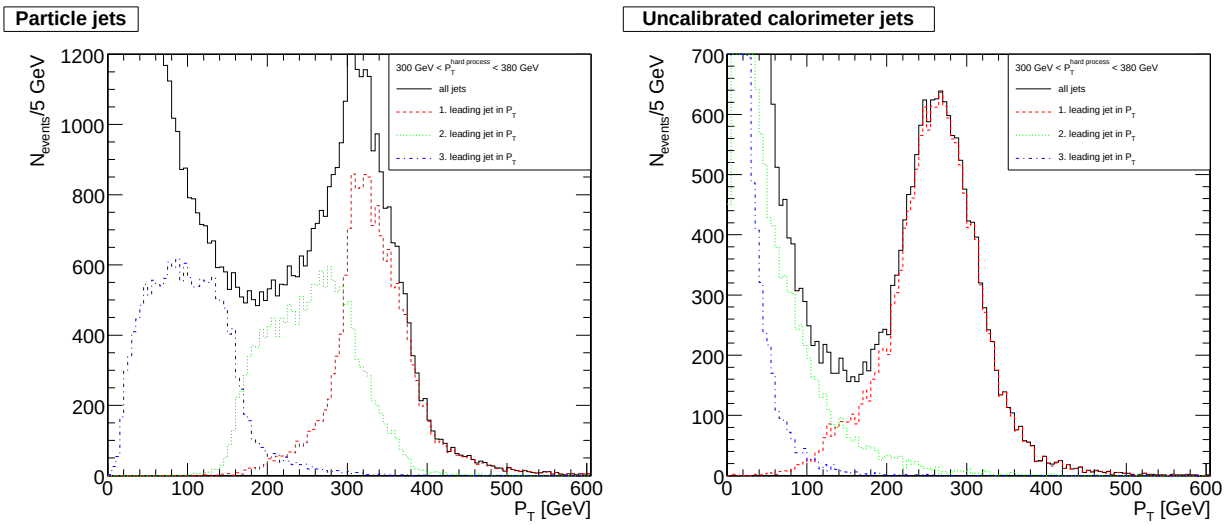

Figure 6.10: *Transverse momentum distribution of particle (left) and calorimeter jets (right). The black curves show the distribution for all jets. In red, only the jet with the highest transverse momentum is considered. The distribution for the second and third leading jet in p_{T} is drawn in green and blue, respectively.*

Figure 6.10 shows the distribution of the transverse momentum of all jets in events with a reconstructed Z boson as well as the corresponding distribution for the first, second and third leading jet in p_{T} for both, particle and calorimeter jets. This comparison is restricted to one bin of the hard transverse momentum, as the effects to be discussed are observable in the other bins as well.

The distribution of this observable looks reasonable for calorimeter jets. The leading jet transverse momentum is distributed around ≈270 GeV, and the second and third leading jet show the expected steep drop-off with increasing p_{T}. However, this is not observed for the transverse momentum of the second and third leading jet in p_{T} of the particle jet collection, which show a pronounced peak instead.

The origin of these artificial jets, which have no correspondence in the calorimeter jet collection, is the definition of the final state for the jet clustering used in CMS. In the current release of the software framework, all particles with status 1 are clustered to particle jets, including the two muons from the decay of the Z boson. As these muons have a large transverse momentum, they act as seed for the jet clustering and lead to additional jets in the particle jet collection which are from here on denoted as fake muon jets.

The muon origin of these additional jets is demonstrated in Figure 6.11, where the minimum of the angular distance ΔR between one of the muons and the closest jet to this muon is drawn versus their ratio of the transverse momentum. The distribution for particle jets shows a significant contribution of entries with very small angular distances and a transverse momentum of the jet comparable to the muon which does not appear for calorimeter jets. As the jet may acquire additional energy from low energetic particles during the clustering, it contains at least the muon

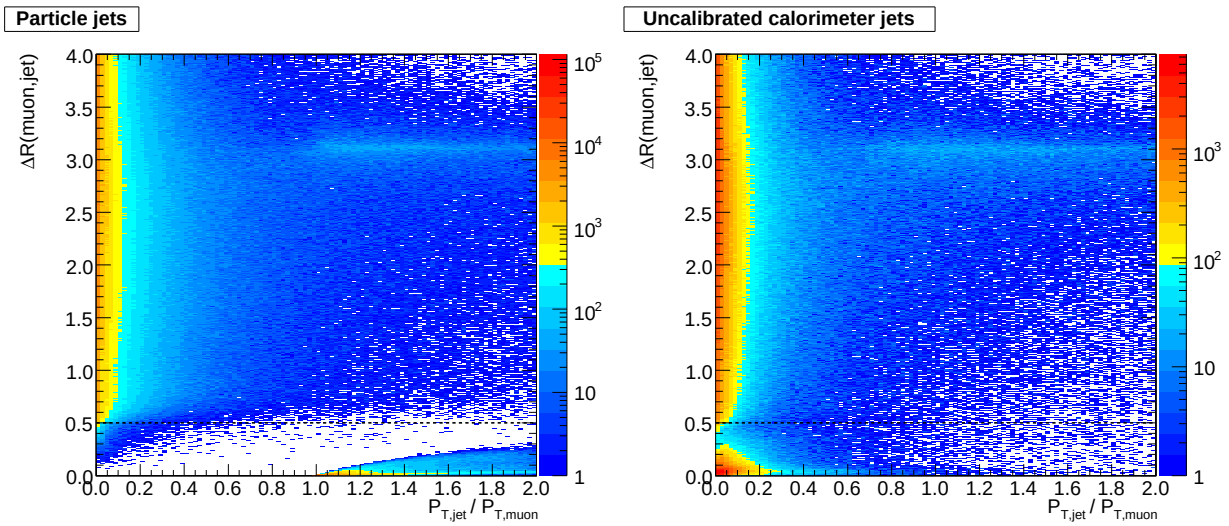

Figure 6.11: *Minimum of the angular distance ΔR between one of the muons and the closest jet to this muon for particle jets (left) and calorimeter jets (right). The distribution for the particle jets shows a significant contribution of entries with very small angular distances and a transverse momentum of the jet comparable to the muon.*

leading to a comparable transverse momentum of the jet and a vanishing angular distance between the muon and the jet. Once additional particles are clustered, the transverse momentum of the jet increases as well as the angular distance to the original muon.

As the presence of the additional fake muon jets introduces ambiguities between the first and second leading jet in p_{T} and hinders the direct comparability between calorimeter and particle jets, they are removed by rejecting jets which are closer to a muon from the decay of the Z boson than $\Delta R < 0.5$. This cut is visualised in Figure 6.11 by the dashed line. Although the removal of the muons before the clustering or the fake muon jets after the clustering may lead to different jets, this effect is negligible for this analysis, as the jet of interest is located in the opposite hemisphere compared to the Z boson and therefore is not influenced by this procedure.

Having removed the fake muon jets from the events, the distributions for particle and calorimeter jets show the expected behaviour as presented in Figure 6.12. In future releases of the CMS software framework, a new definition of the final states which are used for the clustering of particle jets, will be used. In this definition, muons from resonance decays are explicitly excluded and the cut introduced above is no longer necessary.

To avoid a bias in the analysis, only the jet with the highest transverse momentum is considered and no additional cuts on the transverse momentum of the jets are applied. The distribution of the pseudorapidity of this leading jet in p_{T} is drawn in Figure 6.13 for particle and calorimeter jets. The dashed lines enclose the central barrel region of the detector with $|\eta| < 1.3$.

In order to ensure a clean sample of events with the Z boson balanced by exactly one jet of the hard process, additional selection criteria are required. Therefore, events with a second high p_{T}

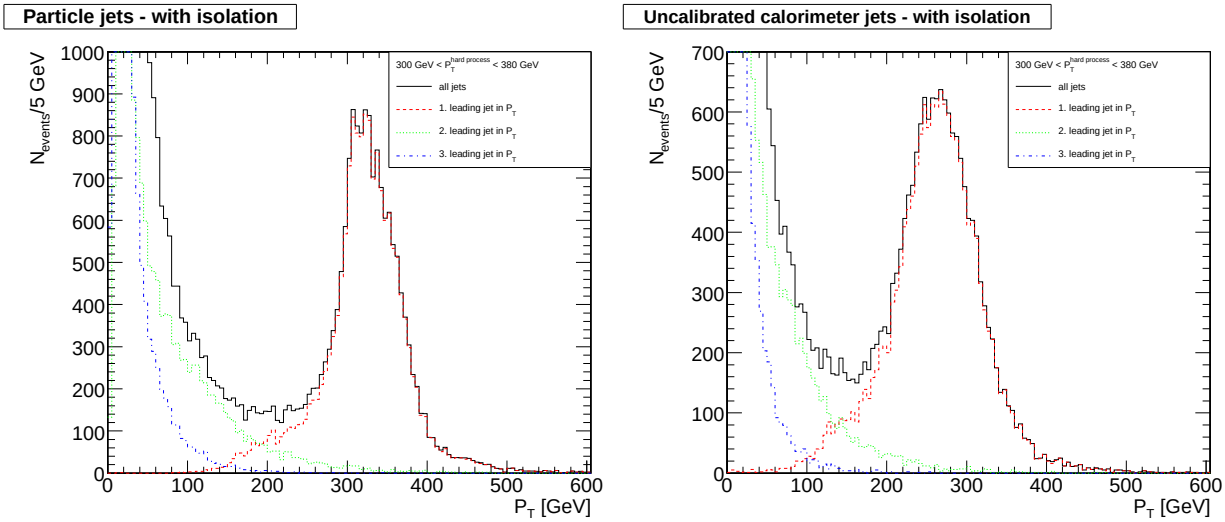

Figure 6.12: *Same as Figure 6.10 but jets which are closer to a muon from the decay of the Z boson than $\Delta R < 0.5$ are rejected.*

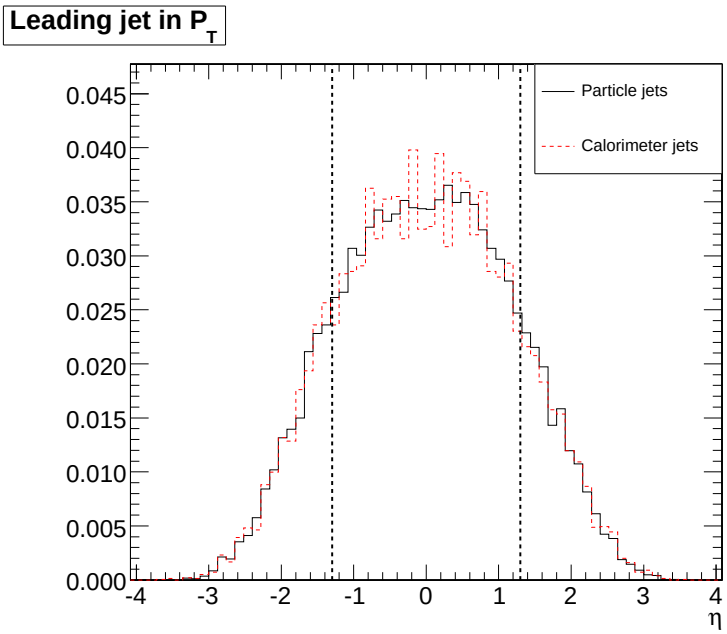

Figure 6.13: *Distribution of the pseudorapidity of the leading jet in p_{T} for particle and calorimeter jets. The dashed lines enclose the central barrel region of the CMS detector with $|\eta| < 1.3$.*

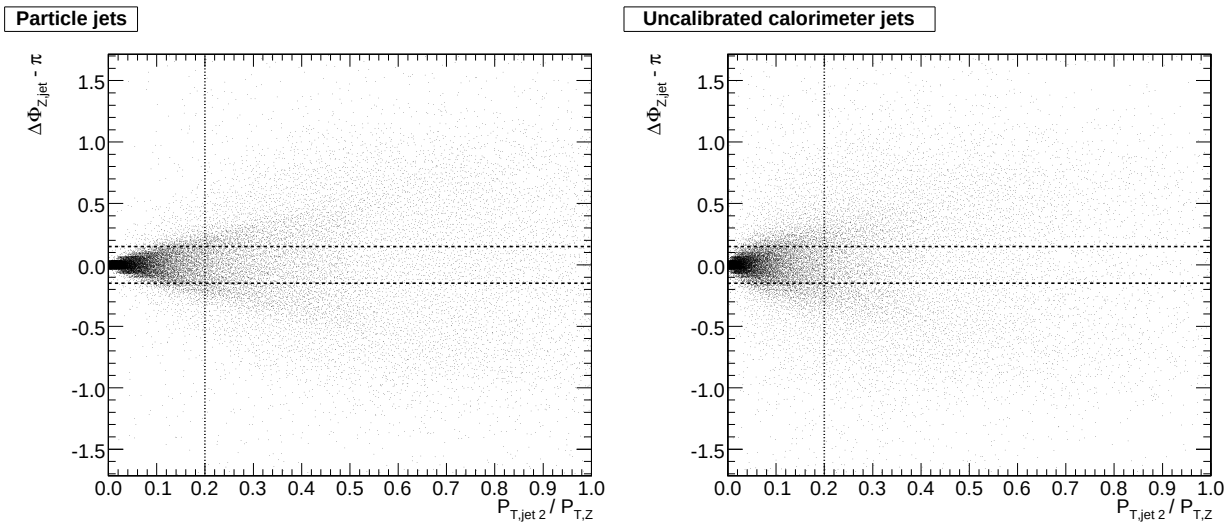

Figure 6.14: *Difference of the azimuth of the leading jet and the Z boson versus the fraction of the momentum of the second leading jet compared to the Z boson for particle (left) and calorimeter jets (right). The dashed lines indicate the cuts.*

jet have to be discarded and the leading jet is required to be well balanced in the azimuthal angle compared to the Z boson.

However, both criteria are correlated and the distribution of the difference of the azimuth of the leading jet and the Z boson versus the fraction of the momentum of the second leading jet compared to the Z boson is drawn in Figure 6.14. Again, events from different $\hat{p_{\mathrm{T}}}$ bins are combined without weighting for the different cross-section.

The distributions of particle and calorimeter jets are comparable, whereas the latter is wider as a result of the detector resolution. To select a pure sample of events where the Z boson is balanced by exactly one jet, only events which are in the area enclosed by the dashed lines are considered. This corresponds to a maximal momentum of the second leading jet of 20% of the momentum of the Z boson. In addition, the Z boson is only considered to balance the leading jet if the difference of their azimuthal angles holds

$$|\Delta\phi(\mathrm{Z}, \mathrm{leadingJet}) - \pi| < 0.15.$$

This choice for the selection cuts is dominated by the optimisation of the purity of the sample and the remaining number of events.

Having applied the event selection discussed above, a large number of reconstructed Z plus one jet events is still remaining for the purpose of this analysis. The available number of events corresponding to an integrated luminosity of 1 fb^{-1} after all selection cuts is presented in Figure 6.15. Here, the jets are reconstructed using the SISCone 0.5 algorithm. Again, the black bars indicate the statistical uncertainty, whereas the red boxes show the uncertainty related to the number of

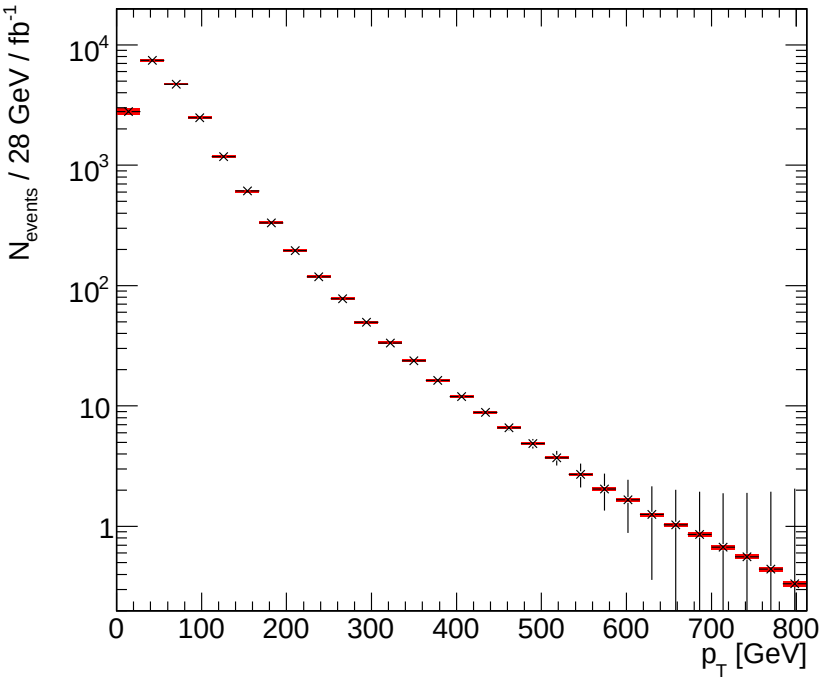

Figure 6.15: *Expected number of reconstructed events containing a Z boson which is balanced by one jet for an integrated luminosity of 1 fb^{-1} after selection. The black bars indicate the statistical uncertainty whereas the red boxes show the uncertainty related to the number of simulated events. The jets are reconstructed using the SISCone 0.5 algorithm.*

simulated events available. The different $\hat{p_{\mathrm{T}}}$ bins are combined by applying weights according to their cross-section. Especially in the region of interest of lower transverse momentum of the Z boson up to 250 GeV, a large number of Z plus one jet events is available and a jet calibration using this process becomes feasible.

Summary of the Event Selection

Reconstruction of the Z boson:

- Transverse momentum of the muons: $p_{\mathrm{T}} > 15\,\mathrm{GeV}$
- Pseudorapidity of the muons: $|\eta| < 2.3$
- Matching of muons with opposite charge and an invariant mass M of the di-muon system closest to the mass of the Z boson
- Only events with $|\mathrm{M} - \mathrm{M_Z}| < 20\,\mathrm{GeV}$

Jet selection:

- Rejection of fake muon jets: $\Delta R(\mathrm{Jet}, \mathrm{Muon}) < 0.5$
- Pseudorapidity of the jet with the highest transverse momentum: $|\eta| < 1.3$
- No cut on transverse momentum

Z plus one jet selection:

- Ratio of the transverse momentum of the second leading jet in p_T and the Z boson: $p_\mathrm{T}^{\mathrm{Jet2}}/p_\mathrm{T}^\mathrm{Z} < 0.2$
- Azimuthal angle between the jet with the highest transverse momentum and the Z boson: $|\Delta\phi(\mathrm{Z}, \mathrm{leadingJet}) - \pi| < 0.15$

6.2.4 Observable

The comparison of the transverse momentum of the jet with the Z boson is based on different bins in order to suppress fluctuations. The observable which is chosen as a measure for the quality of the balancing is the response defined as

$$R(p_\mathrm{T}^\mathrm{Z}) = \frac{p_\mathrm{T}^\mathrm{jet}}{p_\mathrm{T}^\mathrm{Z}}. \tag{6.8}$$

Figure 6.16 shows the distribution of this quantity for generated events with a hard transverse momentum between 120 GeV and 170 GeV. The values of the response are distributed around unity which indicates that this concept of Z boson balancing works fine for particle jets. For values of the response of about 0.2 a slight increase of the number of events is observable. This is related to events in which the Z boson is not balanced by exactly one jet and increases with a relaxing of the selection criteria. For the comparison of the results for different bins, the peak position of the response of a dedicated bin is taken from here on as the measure for the quality of the balancing. This peak is determined using iterative Gaussian fits as described in the following:

1. Calculate the mean and the RMS of the histogram
2. These values are taken as start parameters for a first Gaussian fit. The resulting function is drawn with a dashed line in Figure 6.16.
3. The mean and the width from the first fit are taken as start parameters for the final Gaussian fit (solid line in Figure 6.16) and its mean is referred to as the peak of the corresponding distribution. For this fit, only the 1σ area around the peak of the first fit is considered.

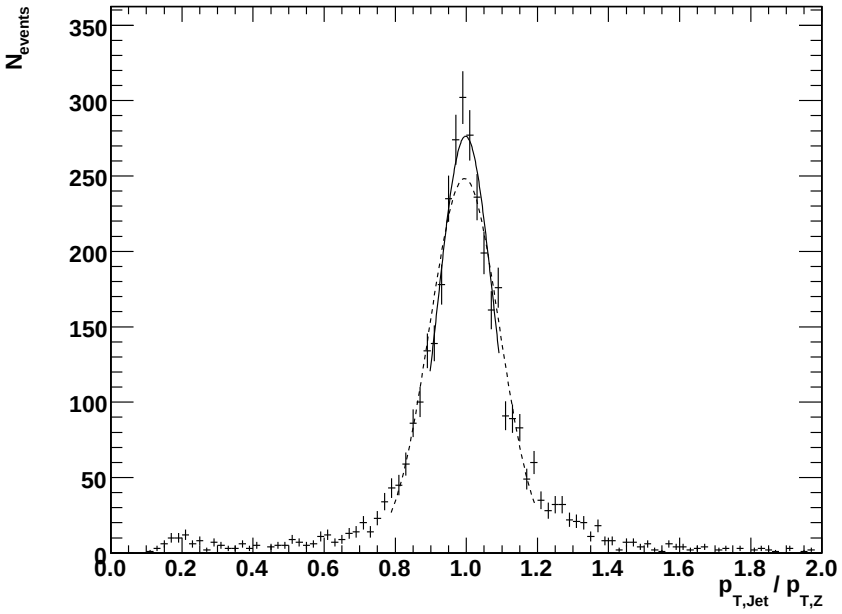

Figure 6.16: *Distribution of the response for particle jets for events with a hard transverse momentum between 120 GeV and 170 GeV. The measure for the quality of the balancing is the peak of this distribution, which is determined using two Gaussian fits. Starting with the mean and the RMS of the histogram, a first Gaussian fit (dashed line) provides the start parameters for the final fit (solid line). The mean of the second fit is referred to as the peak of this distribution.*

This peak of the response is used to evaluate existing jet calibrations in CMS and to derive correction factors based on the principle of Z boson and jet balancing. To be able to evaluate the existing calibrations up to large values of the transverse momentum of the Z boson, a segmentation into bins of the hard transverse momentum as they are used for the generation and simulation process of the datasets is suitable. However, for the determination of correction factors from Z boson balancing, a splitting into bins of the transverse momentum of the Z boson is required.

6.3 Evaluation of Existing Calibrations

For an investigation of the quality of existing calibrations with respect to the Z boson balancing up to highest $p_{\mathrm{T}}^{\mathrm{Z}}$, the steeply falling spectrum of the transverse momentum of the Z boson has to be taken into account. To have a sufficient number of events available over the whole region of transverse momentum, the study is performed in bins of the hard transverse momentum equivalent to the simulated datasets. For each bin, the peak value is determined by the procedure described in the previous section and plotted versus the centre of the bin of the hard transverse momentum. In the following, the Level 2 and Level 3 corrections available in CMS are investigated with the Z boson balancing. Both are derived from a QCD dijet sample using Monte Carlo truth information.

Name	Jet definition	Parameter
Iterative cone 0.5	Iterative cone	R = 0.5
Iterative cone 0.7	Iterative cone	R = 0.7
Midpoint cone 0.5	Midpoint cone	R = 0.5
Midpoint cone 0.7	Midpoint cone	R = 0.7
SISCone 0.5	SISCone	R = 0.5
SISCone 0.7	SISCone	R = 0.7
k_T 0.4	k_T	D = 0.4
k_T 0.6	k_T	D = 0.6

Table 6.3: *Listing of the different jet clustering algorithms used in CMS including their name, the used algorithm and the jet size parameter. The actual jet definition is specified in Appendix B.*

6.3.1 Particle Jets

To demonstrate the consistency of this method, the balancing of the transverse momentum of the jet by the Z boson is investigated for particle jets. The Z boson is reconstructed based on the kinematics of the generated muons and compared to the transverse momentum of the balanced particle jet. As motivated above, the different bins of the hard transverse momentum are analysed independently and the peak of the response is determined for each bin. The different cross-sections are not taken into account, which allows a comparison of the transverse momentum of the boson with the balanced particle jet up to highest transverse momentum for all algorithms currently used in CMS as listed in Table 6.3.

The result of the comparison is shown in Figure 6.17. In this scenario, the response of a perfect balancing would be equal to unity and is indicated in the figure by the solid line. The eight algorithms investigated show only a small discrepancy from the prefect balancing. For the region of small transverse momenta the deviation is about $\pm 2\%$ and becomes negligible for larger transverse momenta. However, there is a slight difference between the algorithms which is correlated to the size of the jet parameter R and D used for the clustering. From here on, this parameter is generally denoted as R. Whereas the lower set of parameters which is 0.5 for the cone-based and 0.4 for the k_T algorithm tend to underestimate slightly the transverse momentum of the particle jet, this reverses for larger R. In general, the kinematics of the Z boson enables a precise determination of the transverse momentum of the balanced particle jet.

6.3.2 Uncalibrated Calorimeter Jets

The situation changes for uncalibrated calorimeter jets as shown on the left hand side of Figure 6.18. As a result of the non-linear response of the calorimeter and the absorber material, the transverse momentum of the jet is systematically underestimated compared to the Z boson. For the region of small transverse momenta, the reconstructed energy of the jet is less than 60% of the cor-

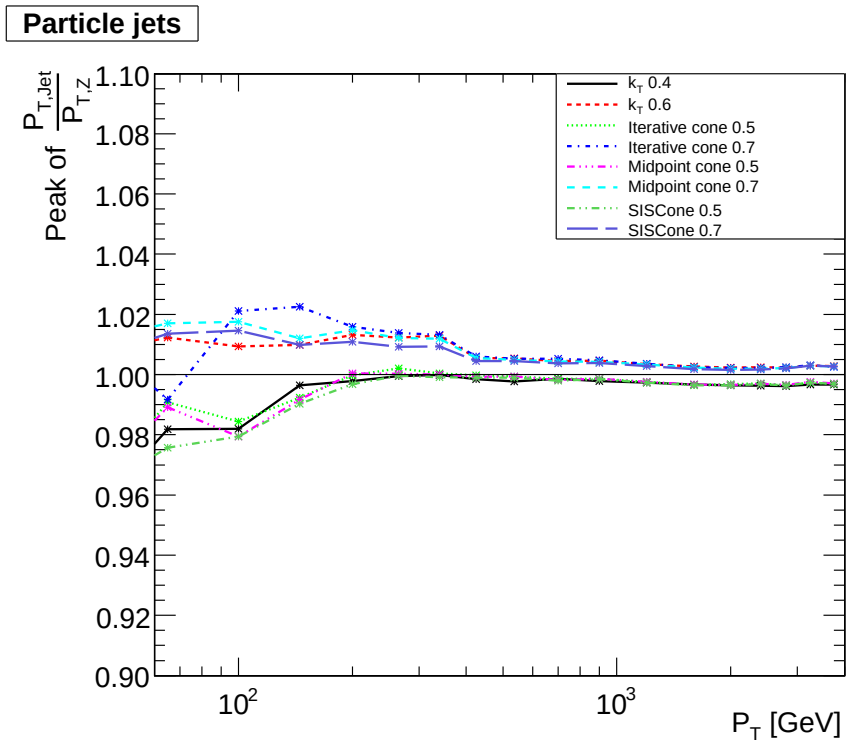

Figure 6.17: *Peak of the ratio of the transverse momentum of a particle jet and the balanced Z boson determined by a Gaussian fit versus the p_{T} for all jet algorithms available in CMS. For the region of small transverse momenta the deviation is about $\pm 2\%$ and becomes negligible for larger transverse momenta.*

responding momentum of the Z boson. With increasing p_{T}, this effect decreases and for the bins of the largest transverse momenta, the energy underestimation is only about 10%. As it is difficult to resolve differences between the eight algorithms on the absolute scale of the response, the plot on the right hand side of Figure 6.18 presents the curves divided by the values of the $\mathrm{k_T}$ 0.4 algorithm and resolves the same trend as already observed for the particle jets. The underestimation of the transverse momentum for larger jet size parameters is slightly lower compared to smaller values of R and vanishes with increasing transverse momentum.

The level of uncalibrated calorimeter jets is the starting point for the factorised multi-level calibration used by CMS. As discussed, the first calibration to be applied is the correction for the offset coming from pile-up, electronic noise and the jet energy loss due to the calorimeter thresholds. Since this Level 1 calibration was not available on the timescale of this study it has not been applied.

Hence, the first correction of the uncalibrated calorimeter jets is the Level 2 correction intended to flatten the detector response as a function of the pseudorapidity with respect to the control region and is derived from Monte Carlo truth information based on a QCD dijet sample. The effect of this calibration on the Z boson balancing is shown in Figure 6.19. On the left hand side the absolute

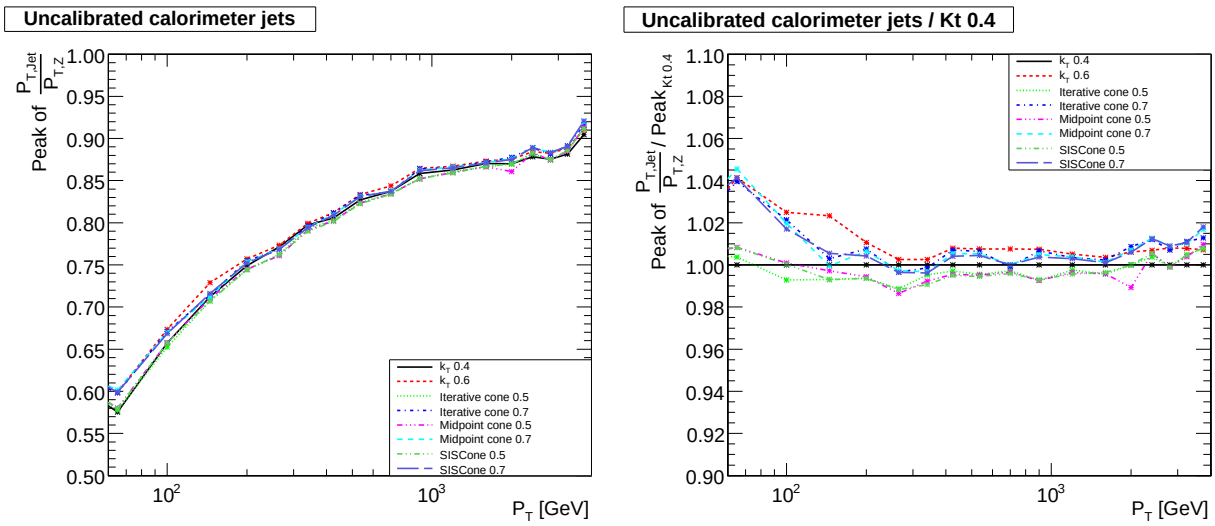

Figure 6.18: *Peak of the ratio of the transverse momentum of an uncalibrated calorimeter jet and the balanced Z boson determined by a Gaussian fit versus the p_T (left). On the right hand side, the distributions are divided by the values of the k_T 0.4 algorithm for a better visibility of variations due to the different algorithms.*

value of the peak of the response is presented, whereas the figure on the right hand side shows the curves divided by the values of the k_T 0.4 algorithm. As the balancing is only investigated in the central barrel region of the CMS detector, which is the control region for the L2 correction, the response is not influenced by this correction, as expected. This allows to start with the L2 corrected jets to derive correction factors for the absolute jet energy scale independent of angular variables.

6.3.3 Monte Carlo based QCD Dijet Corrections

The default correction of the absolute jet energy scale in CMS starts from the L2 calibrated calorimeter jets and is currently based on Monte Carlo truth information of a QCD dijet sample. The result of this $L3^{dijet}$ calibrated jets for the Z boson balancing are presented in Figure 6.20. For the region of small transverse momentum, the energy is overestimated by 4%-13% for the different jet algorithms, and jets with larger values of R seem to be more affected. This discrepancy from perfect balancing decreases for larger values of the transverse momentum and changes to a slight underestimation for the bins of highest $\hat{p_T}$.

One effect, which may lead to the observed overestimation for smaller transverse momentum is related to the different subprocesses leading to a different fraction of jets originating from light quarks or gluons. This effect is discussed in more detail in the next section. However, the large deviation as observed for the uncalibrated calorimeter jets is significantly reduced by applying the $L3^{dijet}$ correction factors.

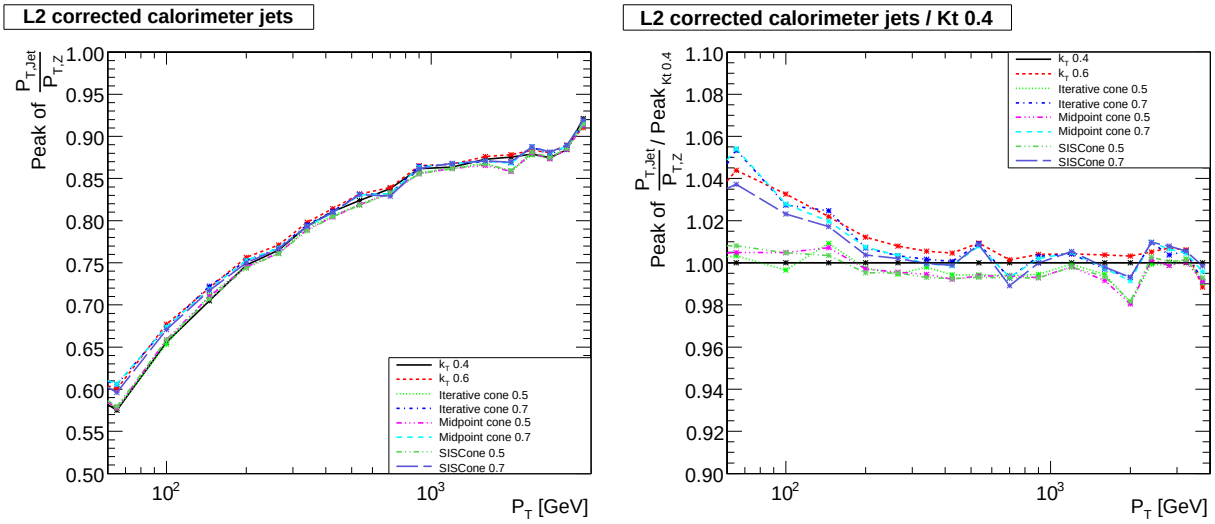

Figure 6.19: *Same as Figure 6.18 but for Level 2 corrected calorimeter jets.*

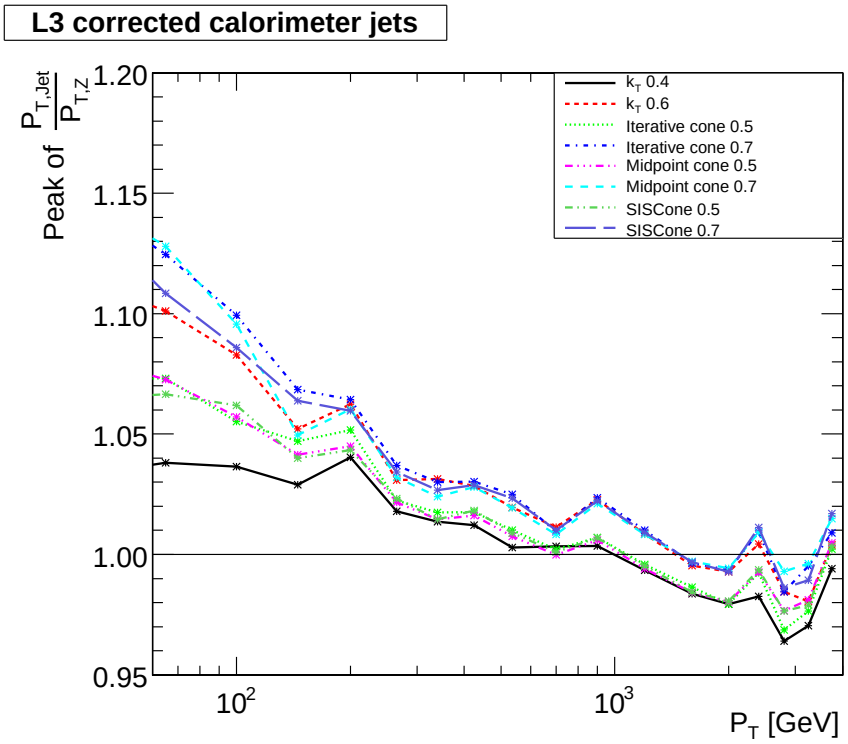

Figure 6.20: *Peak of the ratio of the transverse momentum of calibrated calorimeter jets and the balanced Z boson determined by a Gaussian fit versus p_{T}. The L3 calibration is derived from Monte Carlo truth information of a QCD dijet sample.*

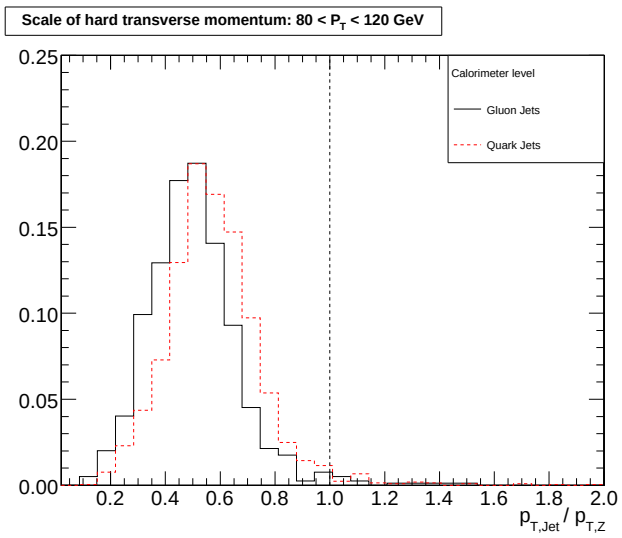

Figure 6.21: *Response for quark and gluon initiated uncalibrated calorimeter jets for one bin of the hard transverse momentum. The response of gluon jets is slightly smaller compared to quark jets.*

6.3.4 Dependence on the Quark-Gluon Fraction

As each calibration is derived from a dedicated physics process, the fraction of light quark and gluon jets will vary, which has an influence on the calibration. Gluon initiated jets radiate more than light quark jets and are therefore slightly wider. This leads to more particles not being clustered into the jet by the algorithm depending on the jet size parameter, which is commonly referred to as out-of-cone effect and results in a lower response for gluon jets. In Figure 6.21 the difference of the response for quark and gluon initiated uncalibrated calorimeter jets is presented for the Z boson balancing for one bin of the hard transverse momentum. The response of gluon initiated jets is slightly smaller compared to quark jets.

The fractions of quark and gluon jets for the QCD dijet and the Z boson sample are presented in Figure 6.22 and Figure 6.23, respectively. Whereas the gluon contribution with more than 60% is dominant for the region of small transverse momentum in the QCD sample, the jets balancing the momentum of the Z boson originate in only about 30% from a gluon up to highest p_{T}.

Thus, the L3$^{\mathrm{dijet}}$ correction factors are determined on a gluon dominated sample which leads to an overestimation of the transverse energy for quark jets. As this jet flavour is dominant in the Z boson jet balancing sample, it can explain the observed overestimation in Figure 6.20. With increasing p_{T}, the fraction of gluon jets in the QCD dijet sample decreases and therewith the overestimation compared to the momentum of the Z boson. This indicates the need for the Level 4 flavour specific calibration. In addition, this difference has to be taken into account for the combination of different Level 3 corrections, as it is planned by the CMS collaboration.

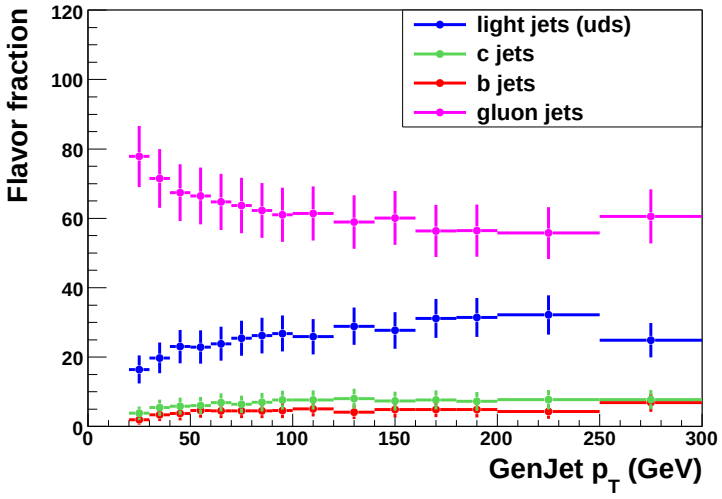

Figure 6.22: *Fraction of jet flavours in a QCD dijet sample. In the region of small transverse momentum the sample is dominated by gluon jets [120].*

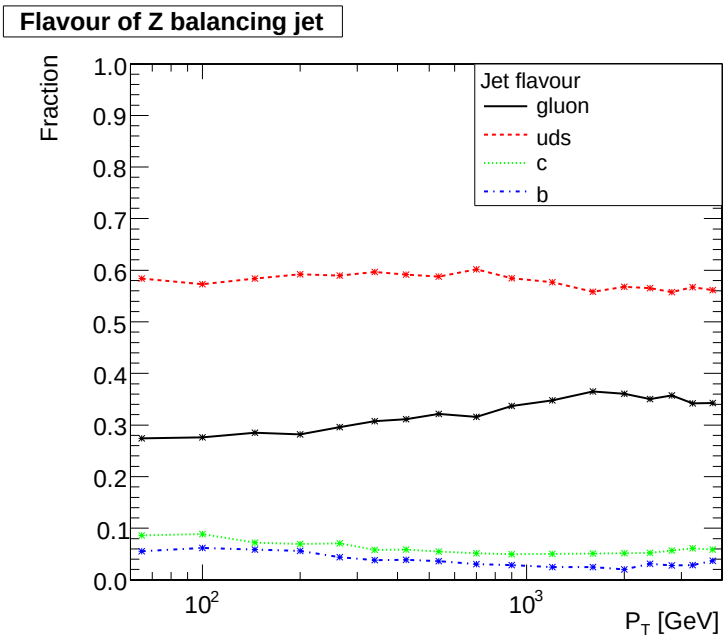

Figure 6.23: *Fraction of the jet flavour in the Z plus one jet sample. Only about 30% of the jets originate from a gluon.*

Name	Jet definition
Midpoint cone	Midpoint cone
Iterative cone	Iterative cone
SISCone	SISCone (LH)
k_T	k_T
Cambridge/Aachen	Cambridge/Aachen

Table 6.4: *Listing of the different jet clustering algorithms used for the study presented in this section. The jet size parameter is varied between 0.3 and 1.0 in steps of 0.1. The definition of the jet clustering algorithms are presented in Appendix B.*

6.4 Influence of the Jet Size Parameter

The quality of the balancing of the Z boson and the jet depends on the actual choice of the jet size parameter as already observed in Section 6.3.1. Smaller jet size parameters result in smaller transverse momenta of the clustered jets. To give an estimate on the influence of R on the quality of the balancing, it is varied between 0.3 and 1.0 in steps of 0.1 for three scenarios of the underlying event.

Besides the Tune DWT, which is the default one of CMS, the Tune S0 is investigated as alternative. Both tunes are derived from data and provide a good description of the Tevatron data. However, an extrapolation to the design parameters of the LHC shows large discrepancies between the tunes and indicates that a precise measurement of the underlying event is mandatory at the LHC.

The data samples which have been used for the study presented in this section have been generated within this thesis on grid resources. For all three underlying event scenarios, the events have been generated in bins of the hard transverse momentum according to Table 6.2.

For technical reasons, the jet definition of the SISCone algorithm is changed for this part of the study and the Cambridge/Aachen algorithm is additionally considered. The actual naming scheme and the corresponding jet definition is given in Table 6.4 and is only valid in this section.

The results of the balancing for the variation of the jet size parameter is shown in Figure 6.24 and Figure 6.25 for fully hadronised Pythia events based on both tunes for the different algorithms. All five algorithms exhibit a rather similar behaviour that small jet sizes on average under- and larger jet sizes overbalance the transverse momentum of the Z boson. Whereas this effect remains below 2% for transverse energies above ≈500 GeV a strong dependence on the actual choice of R for the balancing can be observed for small values of p_T.

Also the results for the alternative Tune S0 show a comparable behaviour. For smaller values of R, both tunes lead to nearly the same results for the balancing. However, for large jet size parameters a slight loss in energy for the Tune S0 is exhibited compared to the Tune DWT.

In order to give an estimate on the offset to the jet energy coming from the underlying event,

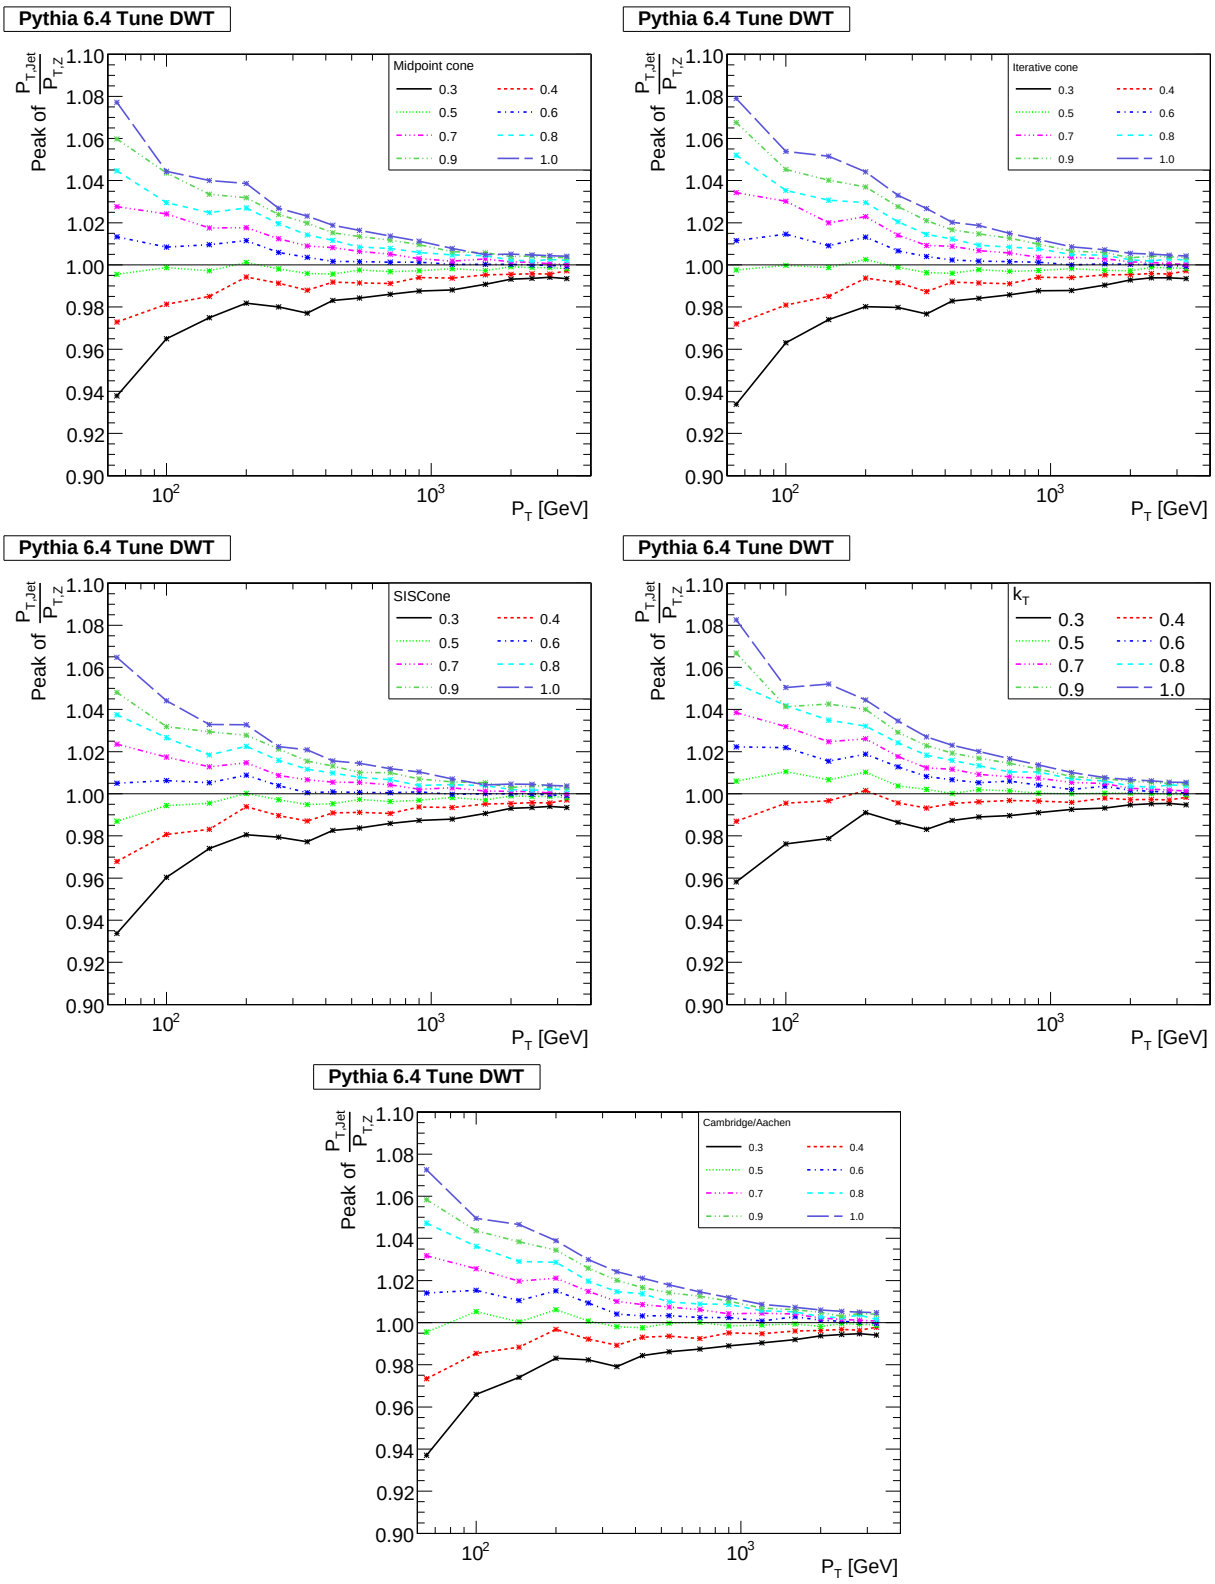

Figure 6.24: *Peak of the ratio of the transverse momentum of a particle jet and the balanced Z boson determined by a Gaussian fit versus the p_T for fully hadronised Pythia Tune DWT events. The jet size is increased from 0.3 to 1.0 in steps of 0.1 for the Iterative cone (upper left), the Midpoint cone (upper right), the SISCone (middle left), the k_T (middle right) and the Cambridge/Aachen algorithm (bottom).*

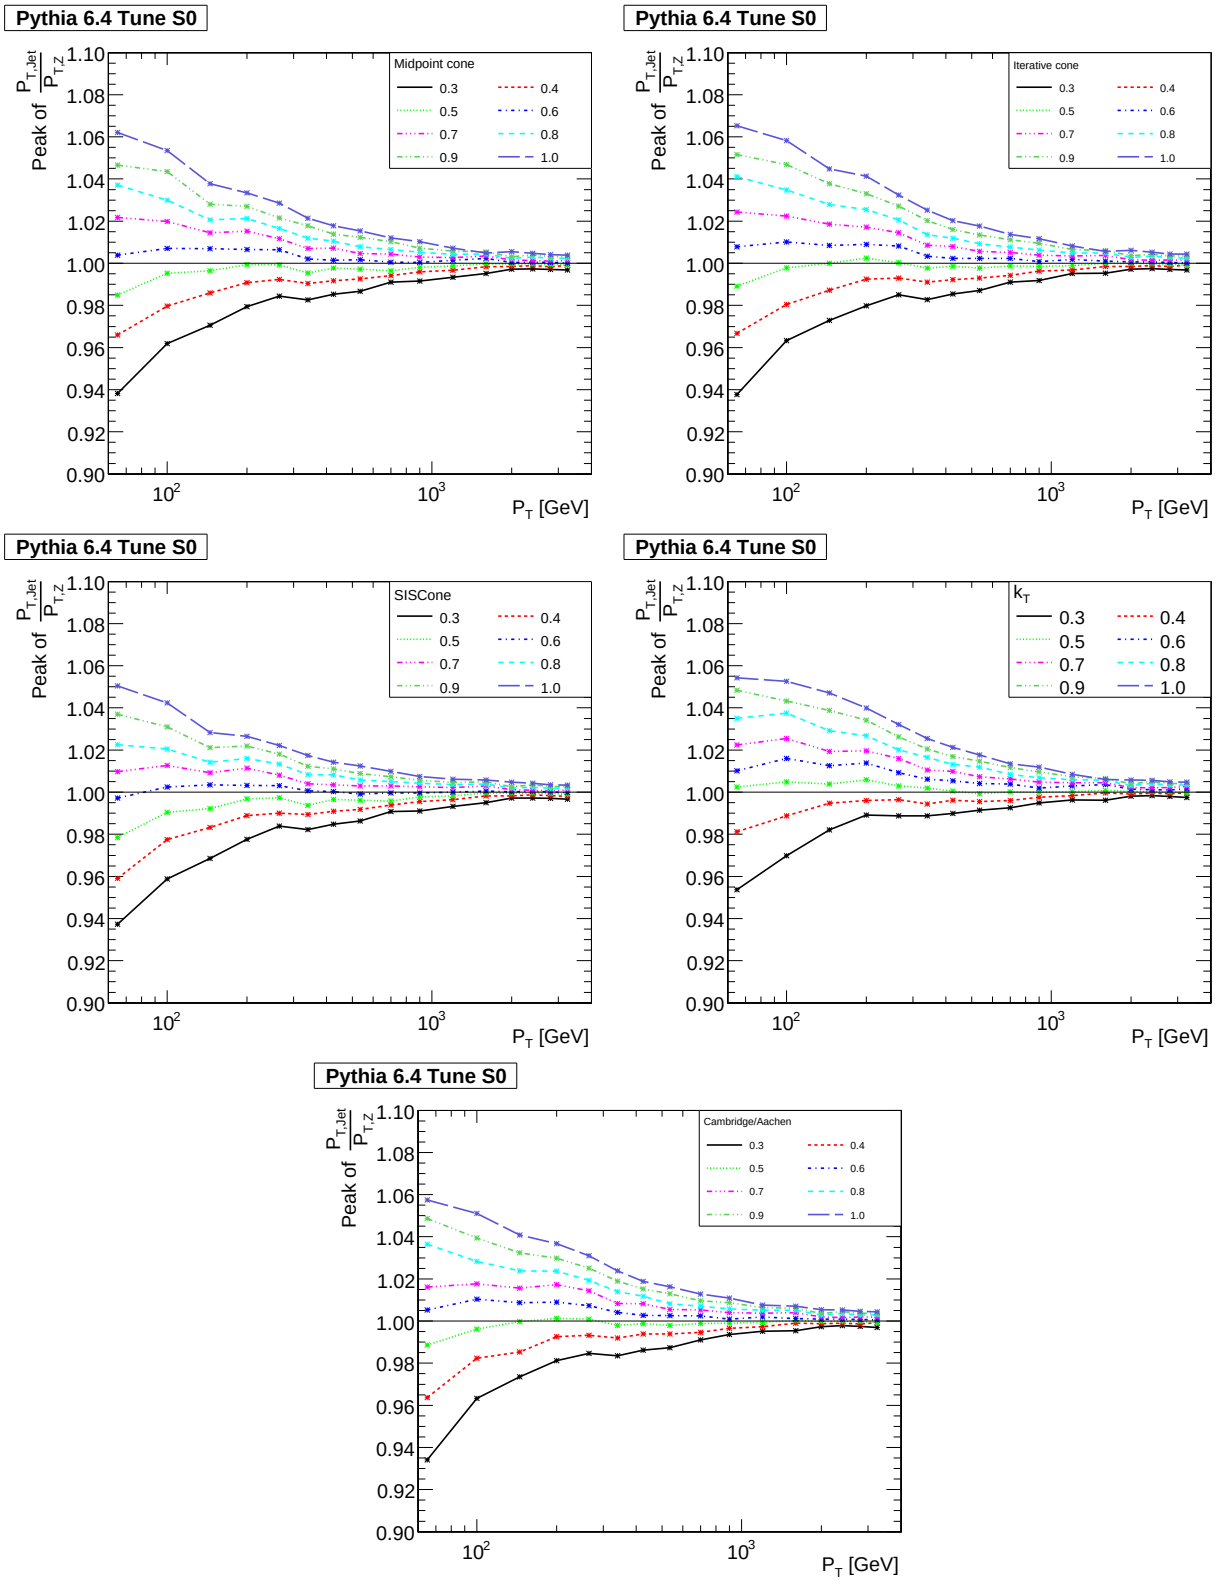

Figure 6.25: *Same as Figure 6.24 but for Pythia Tune S0.*

the peak of the response for fully hadronised Pythia Tune DWT events without Multiple Interactions (MPI) is shown in Figure 6.26. Disabling multiple interactions leads to a systematic underestimation of the transverse momentum of the jet compared to the Z boson. This effect decreases for larger R parameters but remains visible which indicates that the jet algorithm does not accumulate the whole energy of the parton into the jet. Enabling multiple interactions, this feature is compensated by acquiring additional energy from the underlying event into the jet. However, without MPI even the largest employed jet sizes hardly suffice to collect all energy to balance the transverse momentum of the Z boson.

Although the underlying event is more than only multiple interactions, its large influence on the balancing of the Z boson can be demonstrated by disabling MPI. Therefore, the final choice of the preferred parameter of the jet size for the balancing depends heavily on the actual underlying event occurring at the LHC which indicates the need for its precise measurement.

6.5 Jet Calibration based on Z Boson Balancing

Up to now, the evaluation of the different calibrations and effects was based on bins of the hard transverse momentum in order to have a sufficiently large number of events up to highest transverse momenta. Although this binning was advantageous for the evaluation of existing scenarios, it is not for the actual determination of calibration factors using the Z boson balancing method. For this, the quantity of interest is the response in dependence on the transverse momentum of the Z boson, which is regarded as true value of the jet energy:

$$p_{\mathrm{T}}^{\mathrm{jet}} = R(p_{\mathrm{T}}^{\mathrm{Z}}) \cdot p_{\mathrm{T}}^{\mathrm{Z}}. \tag{6.9}$$

However, this relation holds only on average and is evaluated in bins of the transverse momentum of the Z boson. In the current implementation, the binning is chosen similar to the previously used $\hat{p_{\mathrm{T}}}$ bins and is summarised in Table 6.5.

The whole dataset is divided into subsamples of different ranges of the hard transverse momentum. To merge the events, the different cross-sections of the subsamples $\sigma_{\mathrm{subsample}}$ have to be considered. The corresponding weight factor w for a dedicated integrated luminosity $\mathcal{L}$ calculates as

$$w = \mathcal{L} \cdot \frac{\sigma_{\mathrm{subsample}}}{N} \tag{6.10}$$

where N is the number of events in the corresponding subsample.

After having combined all events, the statistical uncertainty on the number of events N_i in bin i is assumed to be purely statistical and therefore calculated by $1/\sqrt{N_i}$. Again, the measure for the quality of the balancing is the peak of the response determined by the procedure described in Section 6.2.4. If the uncertainty on the peak of the response is assumed to be purely statistical,

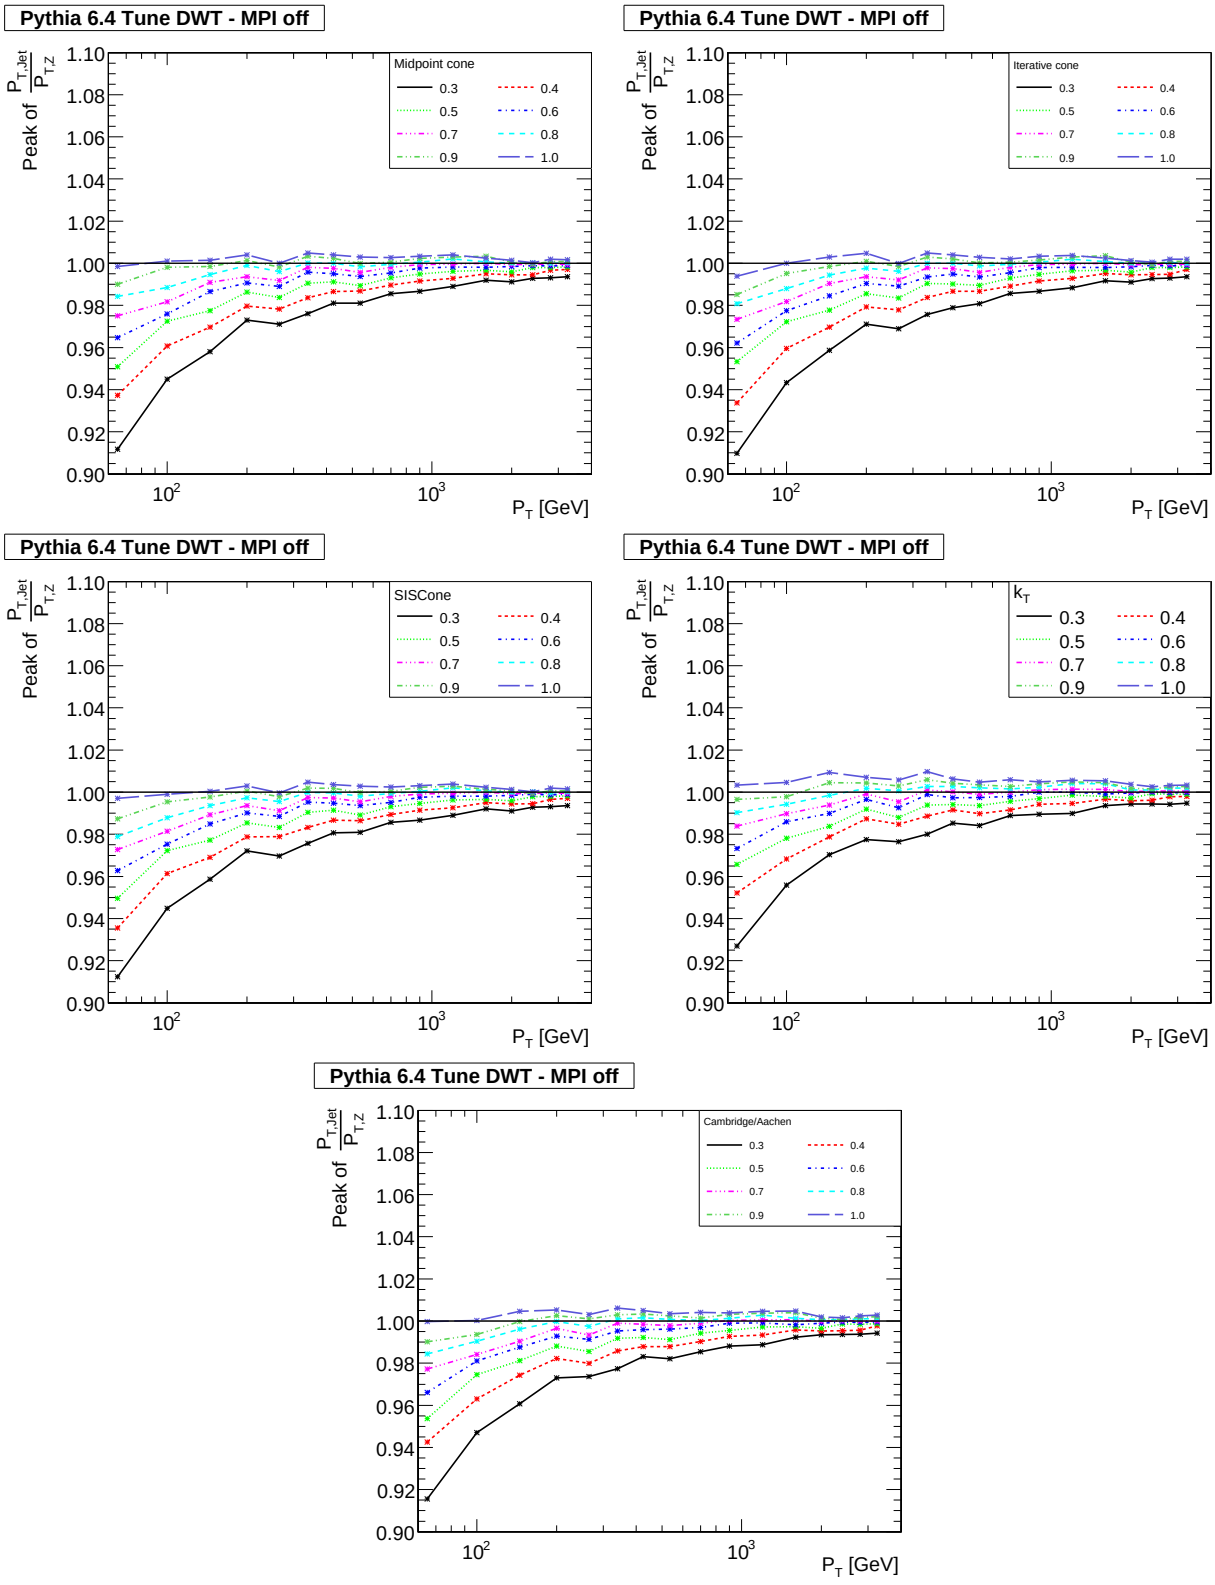

Figure 6.26: *Same as Figure 6.24 but with multiple parton interactions switched off.*

Bins of the transverse momentum of the Z boson
30 - 50 GeV
50 - 80 GeV
80 - 120 GeV
120 - 170 GeV
170 - 230 GeV
230 - 300 GeV
300 - 380 GeV
380 - 470 GeV
470 - 600 GeV

Table 6.5: *Used binning of the transverse momentum of the Z boson for the determination of the correction factors of the absolute jet energy scale.*

it can be identified with the error on the mean value of the Gaussian fit. In Figure 6.27 the peak of the response is shown versus the mean of the transverse momentum of the Z boson for each bin including the statistical uncertainty for the expected number of events corresponding to an integrated luminosity of $1000\,\mathrm{pb}^{-1}$.

The transverse momentum of the jet is only about 50% of the momentum of the Z boson for $p_{\mathrm{T}}^{\mathrm{Z}} \approx 40\,\mathrm{GeV}$. Unfortunately, a determination of the response for lower values of the Z boson is difficult as only few events pass the selection criteria. The statistical uncertainty on the response is negligible up to a transverse momentum of the Z boson of up to 350 GeV and indicates that for the assumed integrated luminosity the balancing can be used up to $p_{\mathrm{T}}^{\mathrm{Z}} \approx 450\,\mathrm{GeV}$.

In Figure 6.28, the response is shown for the number of events corresponding an integrated luminosity of 10 and $100\,\mathrm{pb}^{-1}$, respectively. Even for this reduced number of events, correction factors derived from Z boson and jet balancing can be deduced up to a transverse momentum of the Z boson of 250 GeV and 350 GeV, which is the region of interest for this type of calibration as discussed previously.

6.5.1 Influence of Event Selection on the Response

To ensure a clean sample of events in which the Z boson is exactly balanced by one jet of the hard process, only event topologies are considered which have a second jet with a transverse momentum small compared to the Z boson and a well balanced leading jet with respect to the azimuthal angle. The influence of this selection on the response is estimated by varying the cuts as follows:

- Fraction of the transverse momentum of the second leading compared to the Z boson:

$$\mathrm{F}_{\mathrm{Z,Jet2}} = \frac{p_{\mathrm{T}}^{\mathrm{Jet2}}}{p_{\mathrm{T}}^{\mathrm{Z}}} < \{0.1, 0.2, 0.3\}$$

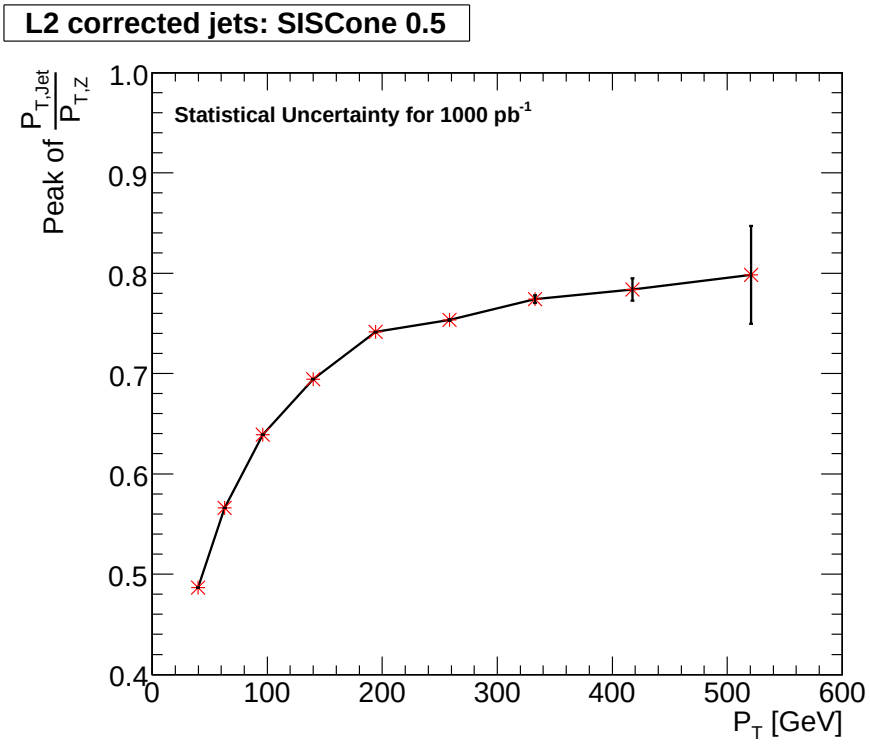

Figure 6.27: *Peak of the ratio of the transverse momentum of L2 calibrated SISCone 0.5 jets and the balanced Z boson versus $p_{\mathrm{T}}^{\mathrm{Z}}$. The peak is determined by a Gaussian fit and the error bars indicate the statistical uncertainty on the Gauss mean for a number of events corresponding to an integrated luminosity of 1000 pb^{-1}.*

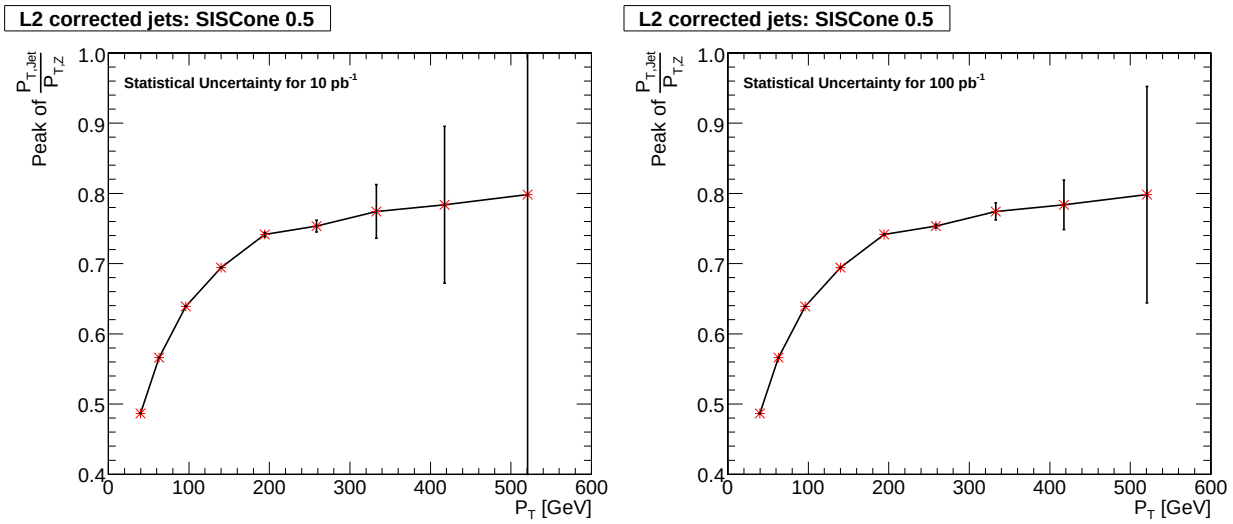

Figure 6.28: *Same as Figure 6.27 but for an expected statistics corresponding to an integrated luminosity of 10 pb^{-1} (left) and 100 pb^{-1} (right).*

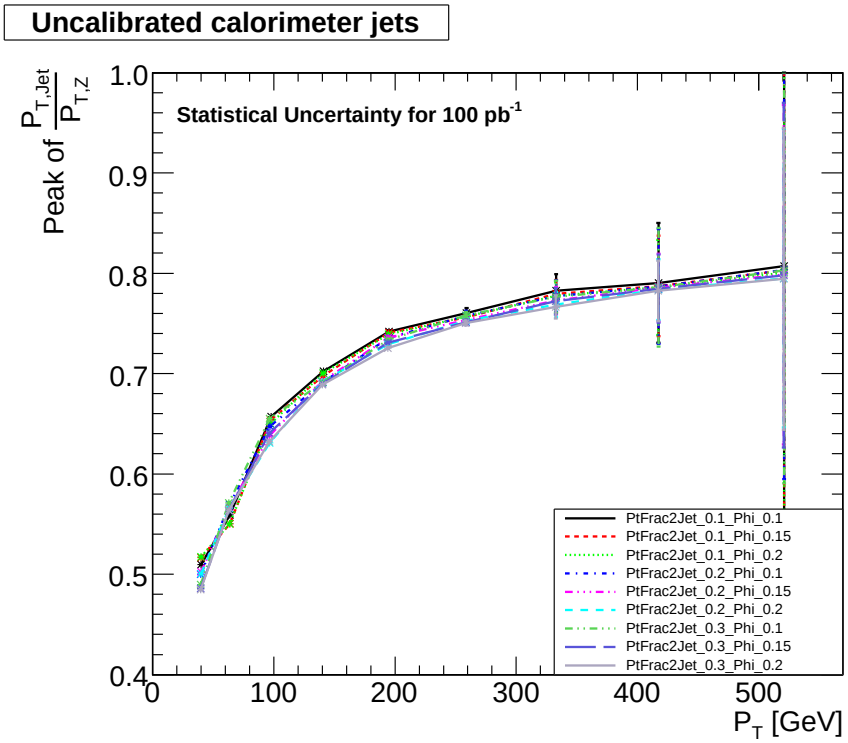

Figure 6.29: *Influence of a variation of the selection cuts on the response for a statistical uncertainty of 100 pb^{-1}. The cut on the fraction of the transverse momentum of the second leading jet compared to the Z boson (PtFrac2Jet) is varied between 0.1 and 0.3 in steps of 0.1. The azimuthal angle between the leading jet in p_{T} and the Z boson (Phi) has to be smaller than 0.1, 0.15 or 2.0.*

- Azimuthal angle between leading jet in p_{T} and the Z boson:

$$\phi_{\mathrm{Z,Jet}} = |\Delta\phi(\mathrm{Z}, \mathrm{leadingJet}) - \pi| < \{0.1, 0.15, 0.2\}$$

The resulting nine curves for the response are shown in Figure 6.29. The shape and the offset of the response is only slightly affected by the cut variation which indicates that the choice of the cuts for this study ($\mathrm{F_{Z,Jet2}} < 0.2$ and $\phi_{\mathrm{Z,Jet}} < 0.15$) does not introduce large uncertainties. This behaviour is expected as the peak of the response is used as measure for the balancing. Possible events with two or more jets of the hard process are located at very low values of the response and so do not influence the position of the peak.

6.5.2 Determination of Correction Factors

Unfortunately, the response determined above is a function of the transverse momentum of the Z boson, see Equation 6.9. However, the final correction factor will be applied on calorimeter jets and thus has to be a function of the measured jet momentum for obvious reasons.

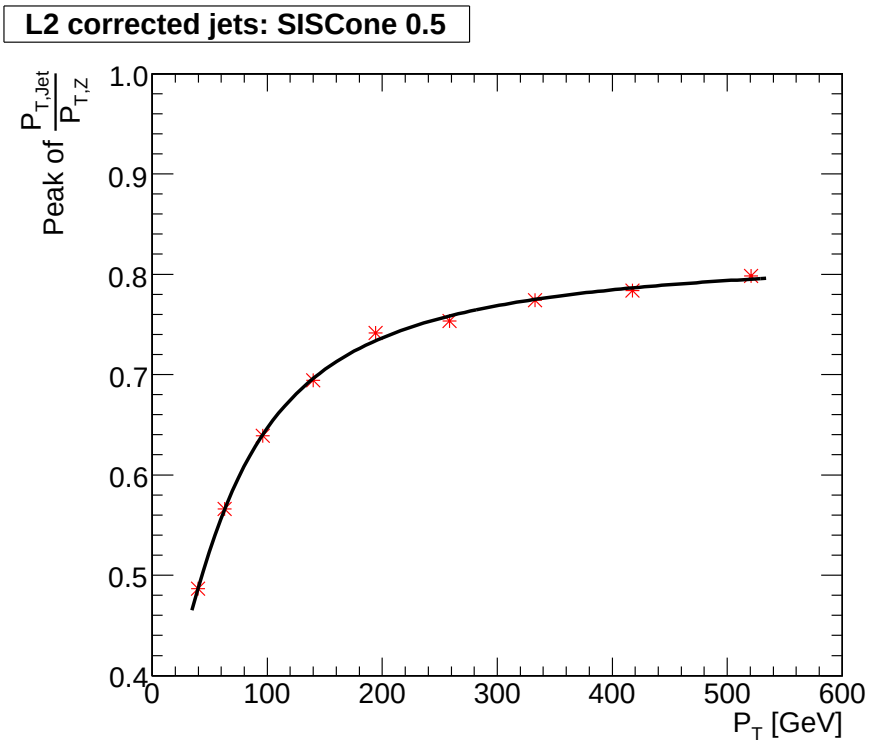

Figure 6.30: *Peak of the ratio of the transverse momentum of L2 calibrated SISCone 0.5 jets and the balanced Z boson versus $p_{\mathrm{T}}^{\mathrm{Z}}$ as well as the corresponding fit with the parameters b_0 = 0.817432, b_1 = 73.3576, b_2 = 7.99148, b_3 = 208.101 and b_4 = -1.4516.*

Therefore, the response as a function of the transverse momentum of the Z boson is fitted with the smoothing function

$$R(p_{\mathrm{T}}^{\mathrm{Z}}) = b_0 - \frac{b_1}{[\log(p_{\mathrm{T}}^{\mathrm{Z}})]^{b_2} + b_3} + \frac{b_4}{p_{\mathrm{T}}^{\mathrm{Z}}} \tag{6.11}$$

which is also used for the dijet based calibration. The second term of this function reflects the natural logarithmic dependence of the calorimeter response on the particles' momentum. The last term models an artificial rise at transverse momenta below 20 GeV in the QCD sample which does not appear in the Z sample. Temporarily, it is kept but the influence of this term on the fit of the Z boson sample will be investigated in detail in near future. In addition, the statistical uncertainties on the response are not taken into account in the following. The response including the fit is presented in Figure 6.30 and shows that the chosen function is well designed to describe the distribution. Once collision data are available, a first set of fit parameters will be determined based on Monte Carlo simulations and will be used as start values for the actual fit of the measured response.

The parameterised smooth function, which is fitted to the response is an estimation of the true detector response and can be used to derive the actual correction factors as a function of the transverse momentum of the calorimeter jet. Therefore, $R(p_{\mathrm{T}}^{\mathrm{Z}})$ has to be inverted in order to determine

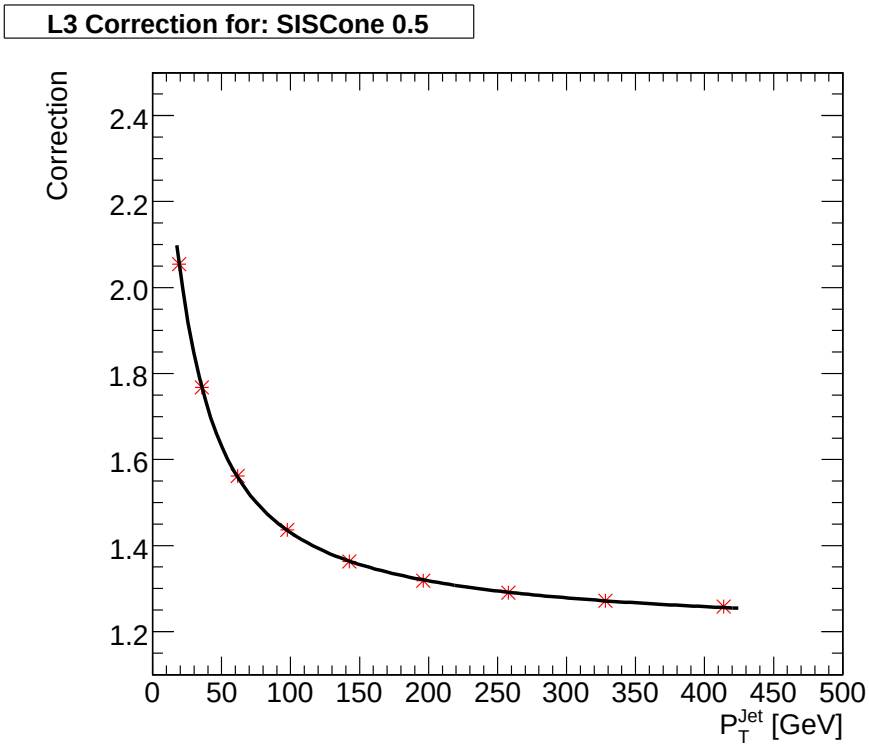

Figure 6.31: *Correction factors for SISCone 0.5 jets determined for different $p_{\mathrm{T}}^{\mathrm{calo}}$ by a numerical Iterative inversion of the response curve of Figure 6.30. The correction factors are approximated by a fit with the parameters $a_0 = 1.19871$, $a_1 = 12.4366$, $a_2 = 5.41948$ and $a_3 = 10.7148$. This fit function represents the Level 3 calibration based on the Z boson balancing.*

the correction factors which are defined as

$$C(p_{\mathrm{T}}^{\mathrm{calo}}) = \frac{p_{\mathrm{T}}^{\mathrm{Z}}}{p_{\mathrm{T}}^{\mathrm{calo}}}. \tag{6.12}$$

In principle, the relation of Equation 6.9 can be inverted numerically to extract the correction factors. The detailed description of the procedure used is presented in Appendix D. However, there are some general remarks concerning this inversion.

The measure for the quality of the balancing used up to now is the peak of the response in a bin of the transverse momentum of the Z boson. As a result of the non-linear response of the calorimeter as well as resolution effects, this peak value is an approximation and would only become the true response for small bin sizes in which the spectrum of $p_{\mathrm{T}}^{\mathrm{Z}}$ is roughly flat. This is not true for the chosen bin sizes and the steeply falling spectrum of the transverse momentum of the Z boson, which introduces an unavoidable bias into the analysis. To reduce this bias to a minimum, the binning will be refined with an increasing number of events.

The correction factors are calculated for different values of $p_{\mathrm{T}}^{\mathrm{calo}}$ and approximated by the function

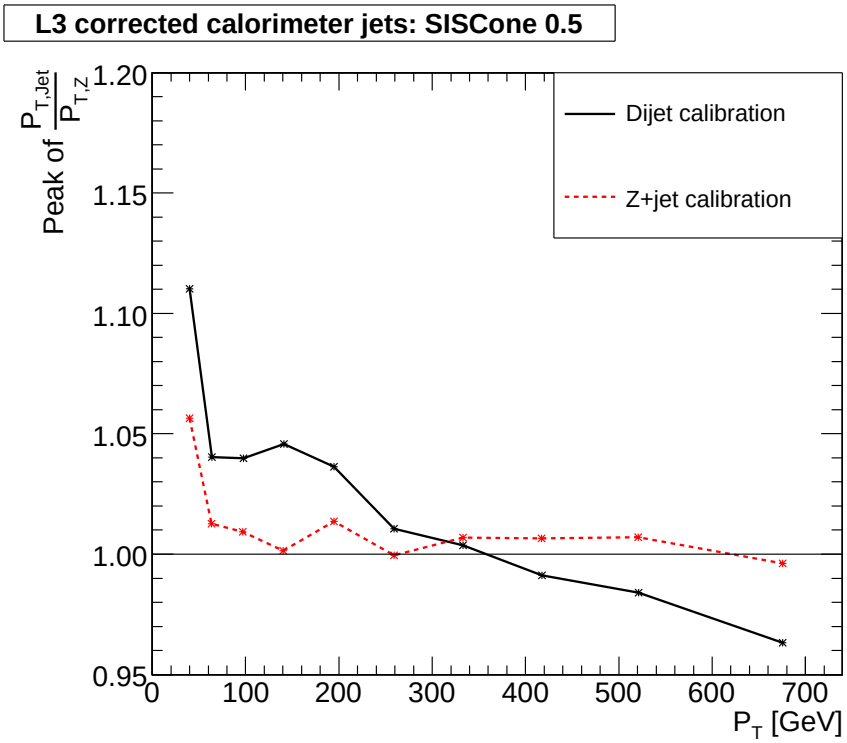

Figure 6.32: *Peak of the ratio of the transverse momentum of L3 calibrated SISCone 0.5 jets and the balanced Z boson versus $p_{\mathrm{T}}^{\mathrm{Z}}$ for the QCD dijet based correction and the calibration based on Z boson balancing. The deviation from unity for the latter is smaller than 2% for the region of $p_{\mathrm{T}}^{\mathrm{Z}} > 65\,\mathrm{GeV}$ demonstrating the self-consistency of the method.*

$$C(p_{\mathrm{T}}^{\mathrm{CaloJet}}) = a_0 + \frac{a_1}{\left[\log(p_{\mathrm{T}}^{\mathrm{CaloJet}})\right]^{a_2} + a_3}. \tag{6.13}$$

Both, the correction factors as well as the fit are shown in Figure 6.31.

A first validation of the derived correction factors is performed on the same data set which is used for the determination of the response. The result for both, the calibration determined from Z boson balancing as well as the dijet calibration are presented in Figure 6.32. The deviation from unity is smaller than 2% for the region of $p_{\mathrm{T}}^{\mathrm{Z}} > 65\,\mathrm{GeV}$ which demonstrates the self-consistency of Z balancing based calibration. For comparison, also the result for the $\mathrm{L3}^{\mathrm{dijet}}$ corrected jets is presented showing the typical overestimation of the jet energy for smaller transverse momenta.

Although the calibration is determined and tested on the same data set, the bias is assumed to be small due to the large number of fully simulated events available and the usage of the fitted peak of the response as a measure for the balancing.

One possibility to combine calibration factors from different methods is to determine their differences using Monte Carlo samples. The difference between the QCD dijet and the Z boson

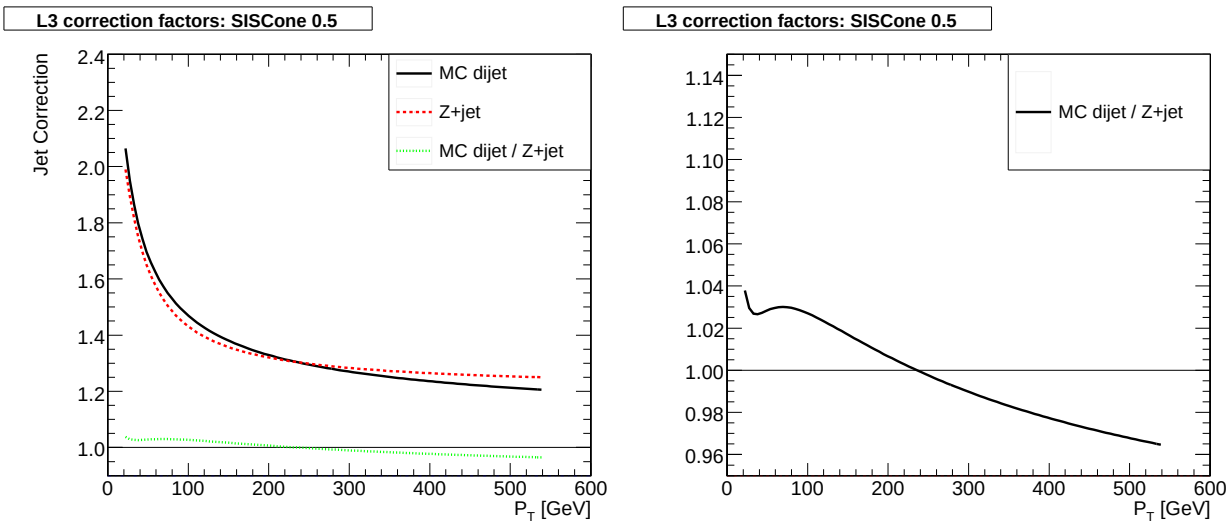

Figure 6.33: *Comparison of the Level 3 correction factors for the SISCone 0.5 jet clustering algorithm derived from a QCD dijet sample and by the method of Z boson balancing.*

balancing based calibration is shown in Figure 6.33. As expected, the $\mathrm{L3}^{\mathrm{dijet}}$ correction factors are larger than those for the Z balancing in the region of small transverse momentum. This changes for large transverse momenta where the Z boson balancing factors are larger. Their ratio is also presented and can be used to combine the different calibrations once collision data is available. In addition, this ratio can be used to calibrate dijet events with data-driven Z based calibrations.

The principles of the determination of a calibration based on the balancing of a Z boson have been exemplified for the SISCone 0.5 algorithm up to now including the successful test of the self-consistency of the method. Based on these results, Z plus jet based calibration factors are determined for all eight jet algorithms used by the CMS collaboration. The closure test for the different calibrations is presented in Figure 6.34 and the corresponding parameters for the correction function are summarised in Table 6.6. The dedicated fit of the response as well as the derived calibration factors of all algorithms are given in Appendix A.

Again, the deviation from unity is smaller than 2% for the region of $p_{\mathrm{T}}^{\mathrm{Z}} > 65\,\mathrm{GeV}$ and larger discrepancies are only observable for very low transverse momenta, especially for the $\mathrm{k_T}$ 0.4 algorithm. The latter is a result of the inversion which was not possible for the region of smallest transverse momentum due to technical limitations of the actual implementation of the inversion procedure.

In this chapter, the concept of jet calibration used by the CMS collaboration was introduced focussing on the correction of the absolute jet energy scale. The selection of a clean sample of events in which the Z boson is balanced by exactly one jet of the hard process is possible and the available number of events will be sufficient for calibration aspects including both, the test of existing and the determination of Z based calibration factors. Especially in the region of lower

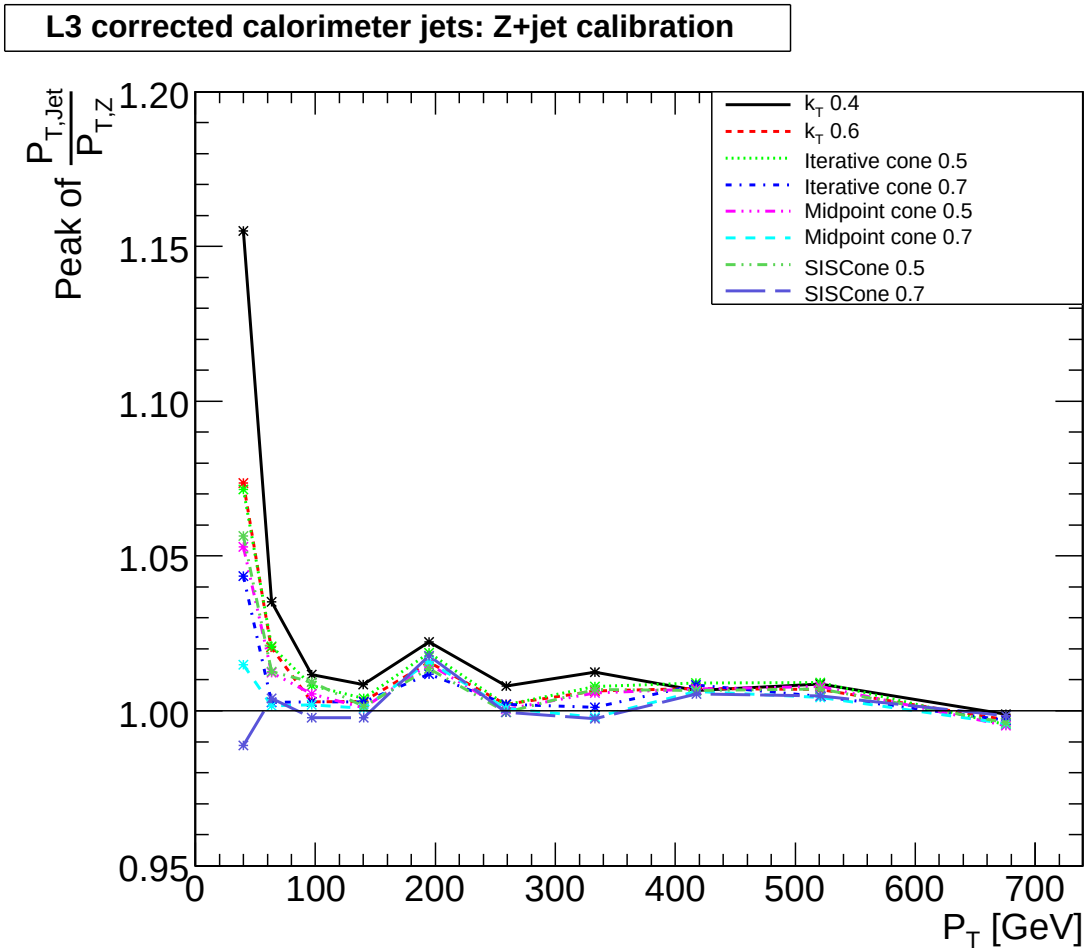

Figure 6.34: *Peak of the ratio of the transverse momentum of L3 calibrated jets and the balanced Z boson versus $p_{\mathrm{T}}^{\mathrm{Z}}$ for the Z boson balancing based calibration. The deviation from unity is smaller than 2% for the region of $p_{\mathrm{T}}^{\mathrm{Z}} > 65\,\mathrm{GeV}$ and for all algorithms demonstrating the self-consistency of the method.*

Algorithm	a_0	a_1	a_2	a_3
$\mathrm{k_T}$ 0.4	1.16382	6.01222	4.3297	2.5721
$\mathrm{k_T}$ 0.6	1.18159	9.37209	5.08716	7.43299
Iterative cone 0.5	1.18853	9.49331	5.01489	6.80117
Iterative cone 0.7	1.18200	9.23337	5.00439	7.82561
Midpoint cone 0.5	1.19847	12.5520	5.44065	10.6173
Midpoint cone 0.7	1.17391	8.37926	4.78199	6.96187
SISCone 0.5	1.19871	12.4366	5.41948	10.7148
SISCone 0.7	1.18810	13.1943	5.46065	13.1091

Table 6.6: *Listing of the jet calibration factors from Z boson balancing for all algorithms used in CMS. These calibrations based on the function $C(p_{\mathrm{T}}^{\mathrm{CaloJet}}) = a_0 + \frac{a_1}{\left[\log(p_{\mathrm{T}}^{\mathrm{CaloJet}})\right]^{a_2}+a_3}$ are now available within the CMS framework.*

transverse momentum, where the photon-plus-jet method suffers from a large QCD background, the Z based calibration will be the candidate of choice in CMS. The determination of the calibration factors from Z boson balancing is possible for the region of $65\,\text{GeV} < p_{\text{T}}^{\text{Z}} < 350\,\text{GeV}$ even for an expected number of events corresponding to an integrated luminosity of $100\,\text{pb}^{-1}$.

In absence of collision data, calibration factors based on the Z boson balancing were determined from Monte Carlo simulations in this study and their self-consistency is demonstrated. These calibrations are now officially available within the CMS framework and can be used for analyses.

Currently, this new calibration is investigated on independent data sets and the first results look promising. In addition, a detailed study of the differences between the response for gluon and quark initiated jets would be desirable and is partly investigated in an on-going diploma thesis.

To provide a Z boson based calibration up to smaller regions of p_{T}^{Z}, an optimisation of the current binning is required. One possibility is to adapt the actual bins according to the available number of events. This goes ahead with the preparation of the analysis for the first collision data to deduce the first data-driven calibrations in CMS.

Conclusion and Outlook

The Large Hadron Collider which will start operation this year is designed to probe the Standard Model of particle physics in an energy range not accessible so far. With a centre-of-mass energy of 14 TeV and a design luminosity of 10^{34} cm^{-2}s^{-1} high precision measurements will further constrain the predictions of the Standard Model and possibly yield hints to new physics. At the LHC, processes of quantum chromodynamics have by far the largest cross-section leading to a large number of those events, which will be used for the commissioning of the detector and measurements of processes of the Standard Model at an energy range not accessible in current experiments. A precise understanding of these events is required as they are a dominant background for most physics analyses.

Besides the physics goals of the LHC, safe storage and processing of the enormous amount of collision data are challenging. The four experiments have decided to cope with these requirements using grid technologies. However, the necessary services do not fully load server machines at smaller grid sites like the IEKP Tier 3 centre. Here, operating system virtualisation enables the consolidation of different services on few server machines. Different virtualisation techniques have been investigated with respect to their suitability and the whole grid infrastructure of the IEKP and several cluster services have been deployed on virtual machines. As a result, most of the institute's service infrastructure is now running on only two server machines. Through disk mirroring, the different services can be migrated from one host to another in case of maintenance or hardware problems. The whole system becomes more flexible and stable and is easier to maintain.

The concept of virtualisation is extendable to the field of batch queuing systems to overcome intrinsic limitations. If a cluster is shared between different groups, multiple operating systems can be offered to different user groups without losing the advantages of load balancing using virtualisation techniques. Within the scope of this thesis, a prototype of a virtualised batch system has been developed and tested on the production cluster of the IEKP. During this evaluation period, the behaviour and particularities of the system have been investigated. The possibility to provide

operating systems individually adapted for the different user groups is a desirable feature. Also maintenance and deployment of new operating systems is simplified and becomes more efficient as shorter service interruptions are required. However, in the current implementation some key functionalities of the batch system and the scheduler are bypassed by an additional program leading to problems. Therefore, the principle of horizontal partitioning should be implemented directly into existing batch systems which is investigated in additional works based on the results of this study.

A reliable and precise jet calibration is indispensable for all physics analyses at the LHC dealing with jets either as signal or background processes. The CMS collaboration envisages a factorised multi-level jet calibration which has been introduced briefly focussing on the Level 3 correction of the absolute jet energy scale. Currently, a L3 correction based on Monte Carlo truth information is available which has been deduced from a sample of simulated QCD dijet events. Once collision data are available, data-driven methods will replace this correction. At previous experiments, photon events have been used to determine the absolute jet energy scale from collision data which are affected by a large and irreducible background in the region of small transverse momenta. Especially in this region, Z boson events are the most promising candidate for calibration aspects and their suitability for this purpose at the LHC has been demonstrated in this study.

For calibration aspects, processes are required in which the Z boson is balanced by exactly one jet of the hard subprocess. As the momentum is conserved in such an event topology, the kinematics of the balanced parton is linked to the Z boson momentum which can be reconstructed with high precision. In this study, only the decay of the Z boson into two muons is considered. Thus, the determination of the kinematics of the Z boson does not rely on any calorimetric information. It was demonstrated that a large number of such events will be available after selection.

As a measure for the quality of the balancing, the ratio of the transverse momentum of the jet and the Z boson was used and the self-consistency of this method was demonstrated with particle jets. The eight investigated algorithms exhibit only a small discrepancy from the perfect balancing. For the region of small transverse momenta the deviation is only about $\pm 2\%$ and becomes negligible for larger transverse momenta.

This situation changes for uncalibrated calorimeter jets where the transverse momentum of the jet is systematically underestimated compared to the Z boson as a result of the non-linear response of the calorimeter and the absorber material. For the region of small transverse momentum, the reconstructed energy of the jet is less than 60% of the corresponding momentum of the Z boson. With increasing transverse momentum, this effect decreases and for the largest transverse momenta, the energy underestimation is only about 10%.

The investigation of the QCD dijet Monte Carlo truth based correction of the absolute jet energy scale has shown an overestimation of the energy of the balanced jet after correction. This overestimation, which depends on the algorithm used for the jet reconstruction, is between 4% and 13% in the region of small transverse momenta and decreases with increasing p_{T}. It has been demonstrated that this effect is related to the different fraction of quark and gluon jets in the investigated

samples. Generally, the response for gluon jets is slightly smaller than for quark jets. As the dijet calibration is deduced on a gluon jet dominated sample in the region of lower transverse momentum, its application on the Z boson sample in which only about 30% of the balanced jets originate from a gluon leads to the observed overestimation.

In addition, the quality of the Z boson and jet balancing depends on the jet size parameter and the underlying event model. To give an estimate on this effect, three different scenarios for the underlying event have been probed with a variation of the jet size parameter. A slight difference between the Pythia Tunes DWT and S0 could be observed. Disabling multiple interactions leads to a systematic underestimation of the transverse momentum of the jet compared to the Z boson. Although the underlying event is more than only multiple interactions, the results demonstrate the large influence of the underlying event model on the balancing. Therefore, its precise measurement at the LHC energy scale is essential to opt for a preferred jet size parameter for this analysis.

Taking purely statistical uncertainties into account, a precise determination of the response from Z boson balancing can be achieved for the region of $65\,\text{GeV} < p_\text{T}^\text{Z} < 350\,\text{GeV}$ even for an expected number of events corresponding to an integrated luminosity of only $100\,\text{pb}^{-1}$ and the influence of the selection criteria on the response is negligible. The final correction factors are deduced from a fit of the response by an iterative numerical inversion for all algorithms supported by the CMS collaboration. A first validation of the derived correction factors has shown that the deviation from perfect balancing is smaller than 2% for the region of $p_\text{T}^\text{Z} > 65\,\text{GeV}$ demonstrating the self-consistency of Z balancing based calibration. These calibration factors are now officially available within the CMS framework and can be used for analyses.

Currently, this new calibration is investigated on independent data sets and first results look promising. In addition, a detailed study of the differences between the response for gluon and quark initiated jets would be desirable. To provide Z boson based calibrations up to smaller regions of the transverse momentum of the Z boson requires an optimisation of the current binning which may depend on the available number of events.

With the correction factors deduced during the scope of this thesis, a calibration of the absolute jet energy scale derived from Z boson events has become available for the first time in CMS. As the Z boson is reconstructed from muons it does not rely on any calorimetric information and the background is negligible. While the current correction factors have been determined from Monte Carlo simulations, the analysis is ready to provide the first collision data-driven calibration with such LHC events. Especially in the region of low transverse momenta, this method will yield a precise determination of the absolute jet energy scale and thus essentially contribute to the reduction of the dominant systematical uncertainty of most CMS analyses.

Response and Calibration Factors

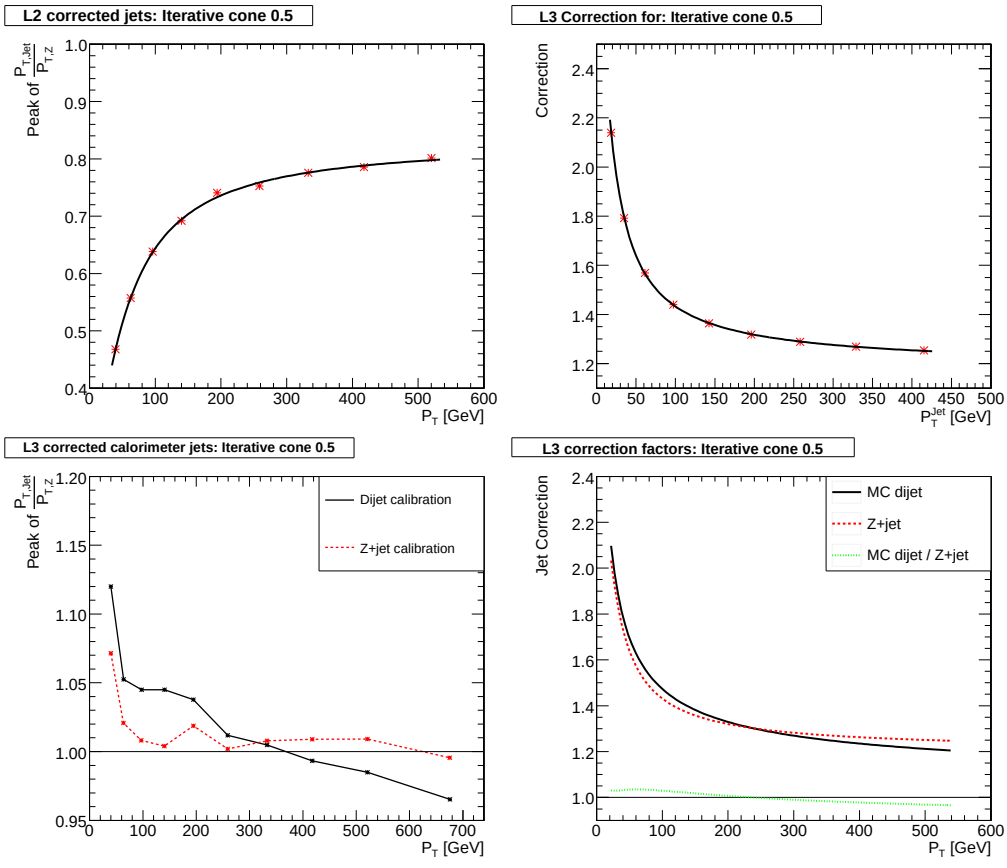

Figure A.1: *Peak of the ratio of the transverse momentum of L2 calibrated Iterative cone 0.5 jets and the balanced Z boson versus $p_{\mathrm{T}}^{\mathrm{Z}}$ as well as the corresponding fit (upper left), the correction factors determined for different $p_{\mathrm{T}}^{\mathrm{calo}}$ by a numerical iterative inversion of the response curve (upper right), the demonstration of self-consistency of the calibration from Z boson balancing (lower left) and the comparison with the correction derived from a QCD dijet sample (lower right).*

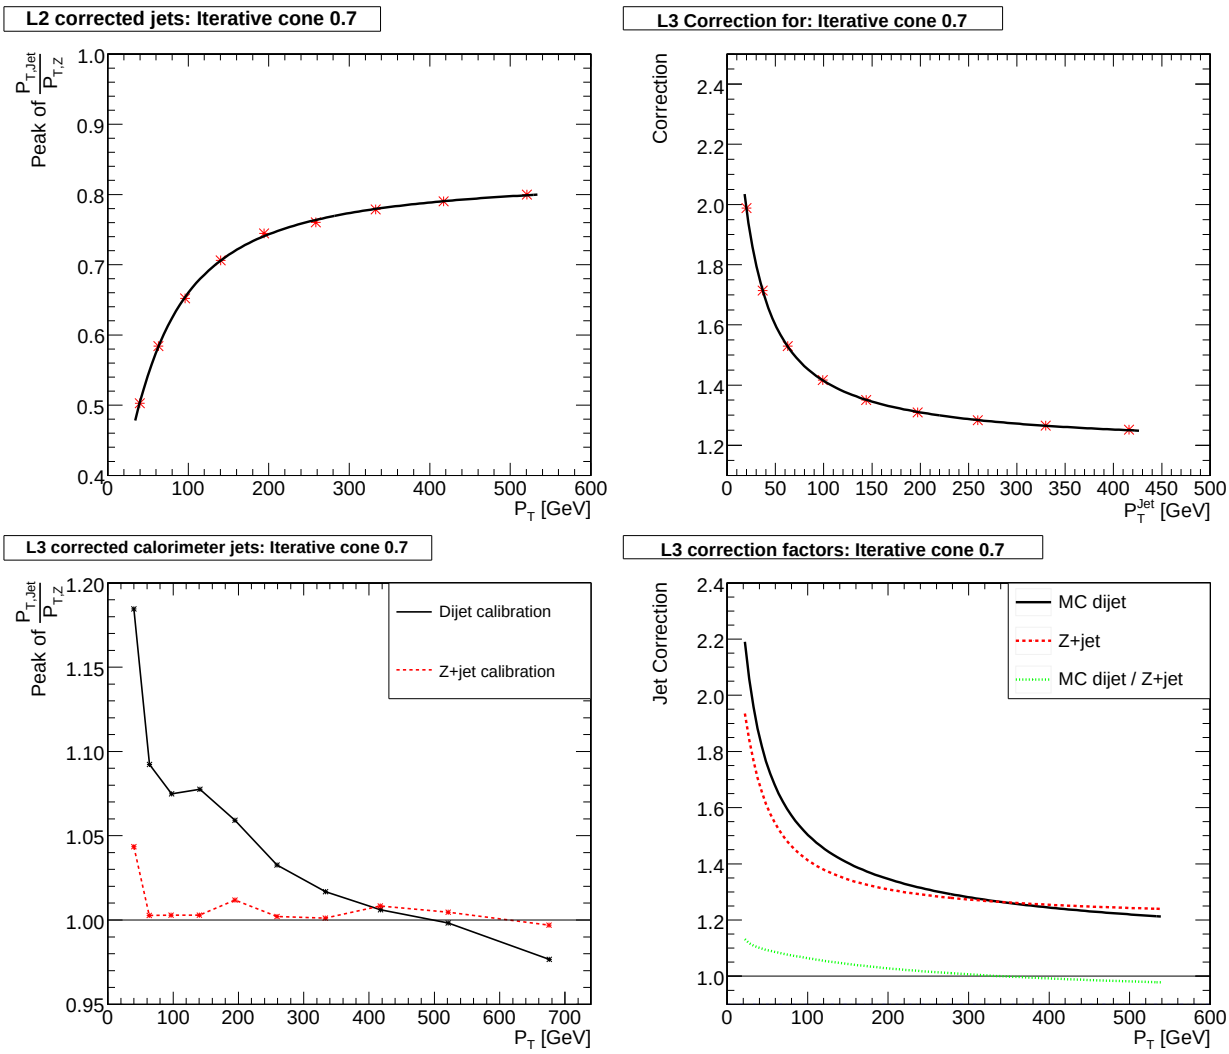

Figure A.2: *Same as Figure A.1 but for the Iterative cone 0.7 jet algorithm.*

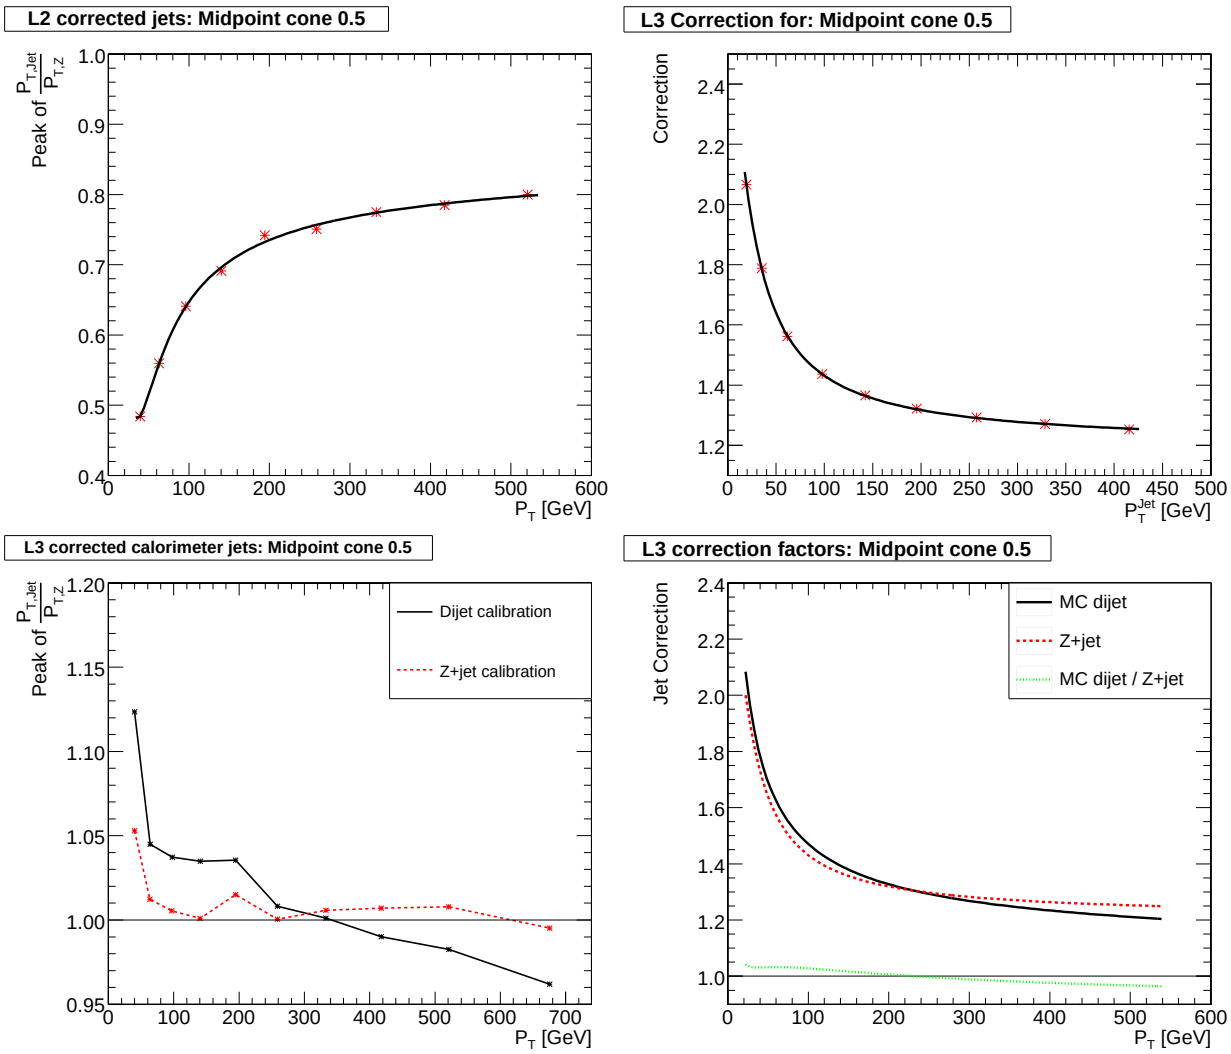

Figure A.3: *Same as Figure A.1 but for the Midpoint cone 0.5 jet algorithm.*

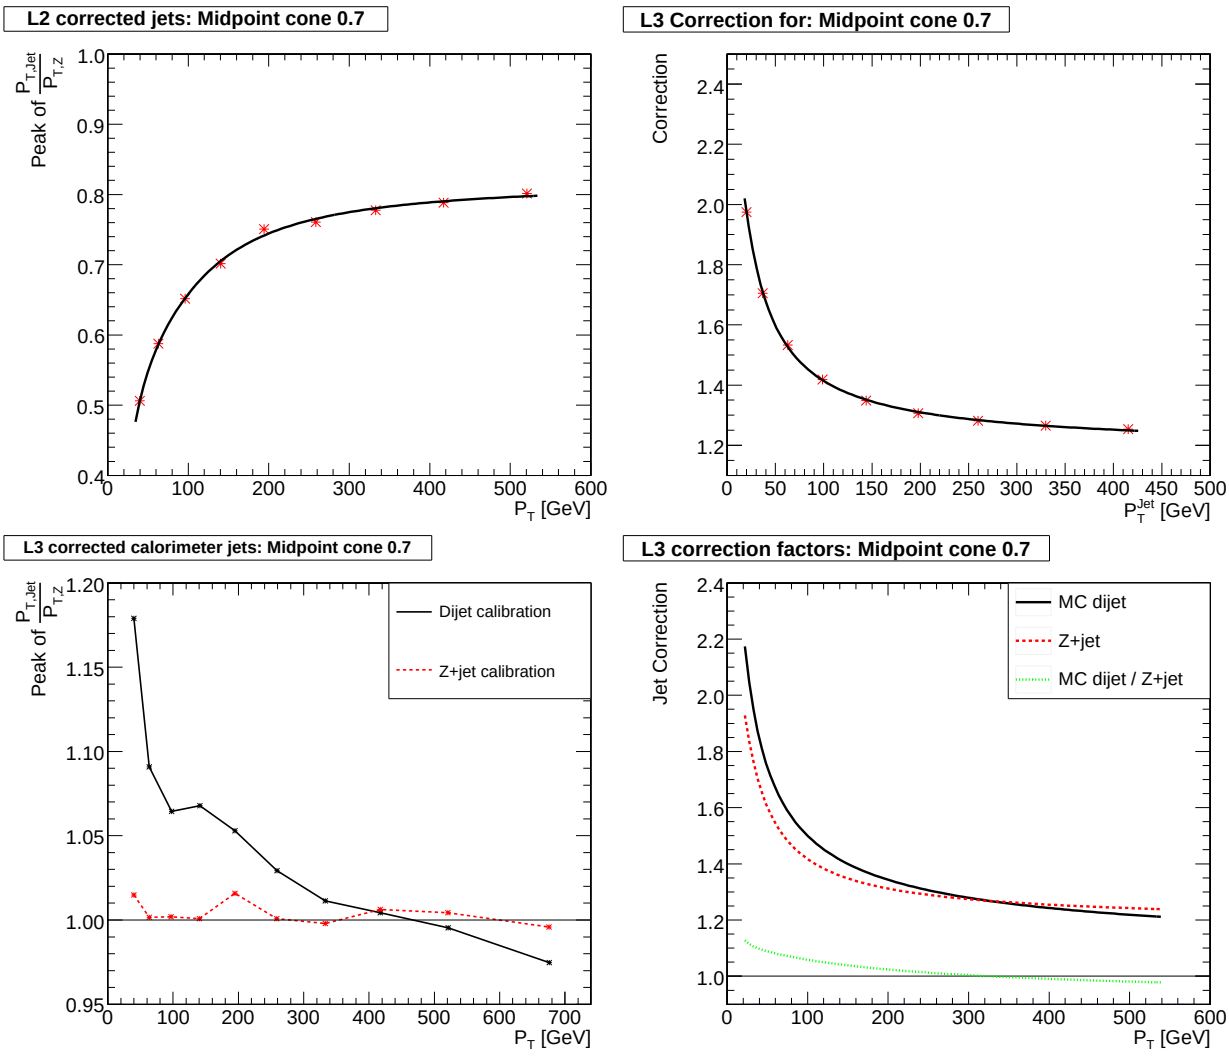

Figure A.4: *Same as Figure A.1 but for the Midpoint cone 0.7 jet algorithm.*

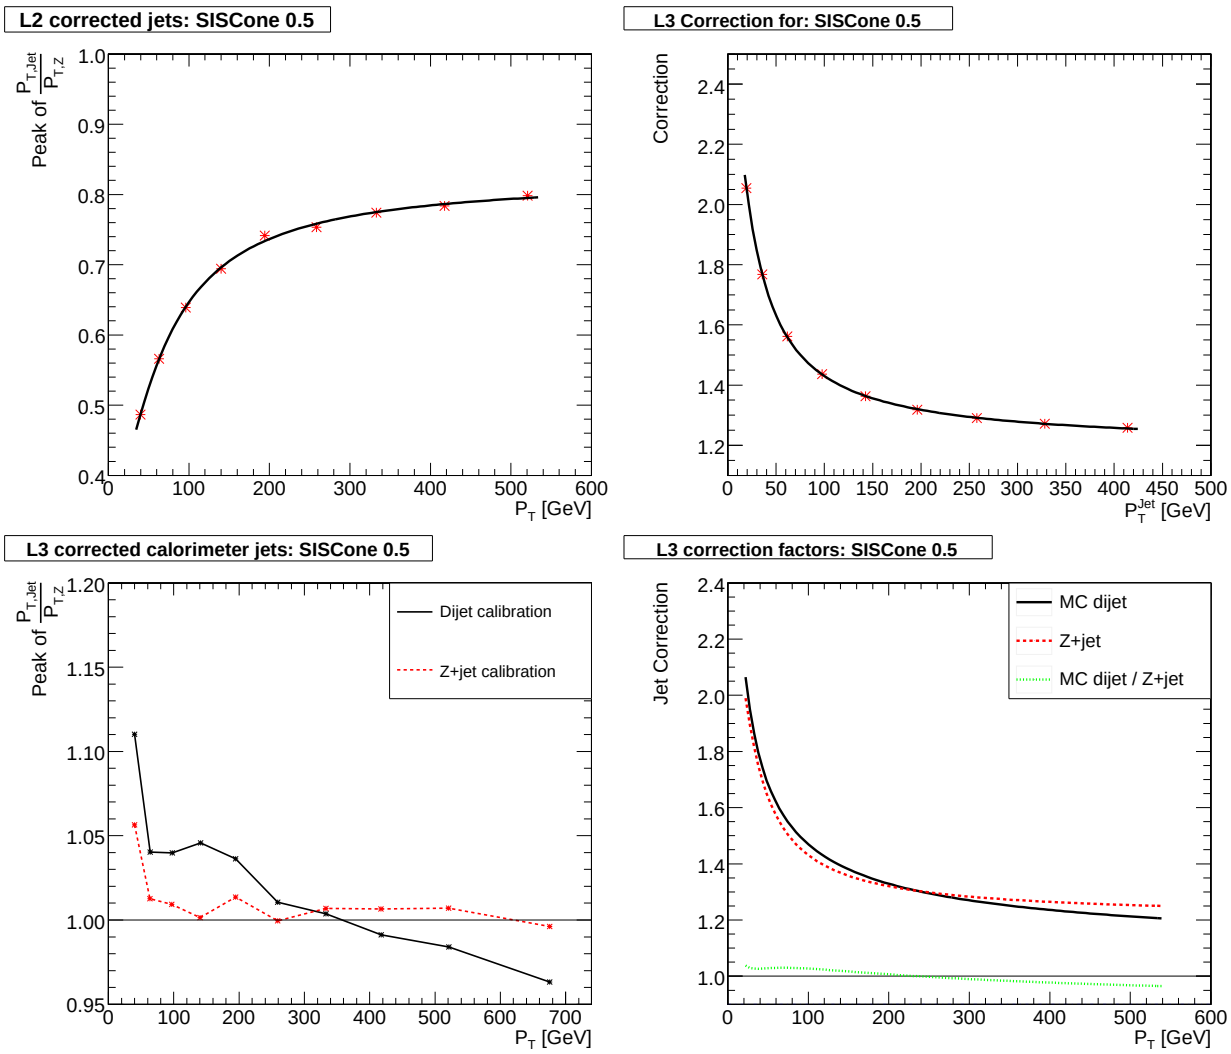

Figure A.5: *Same as Figure A.1 but for the SISCone 0.5 jet algorithm.*

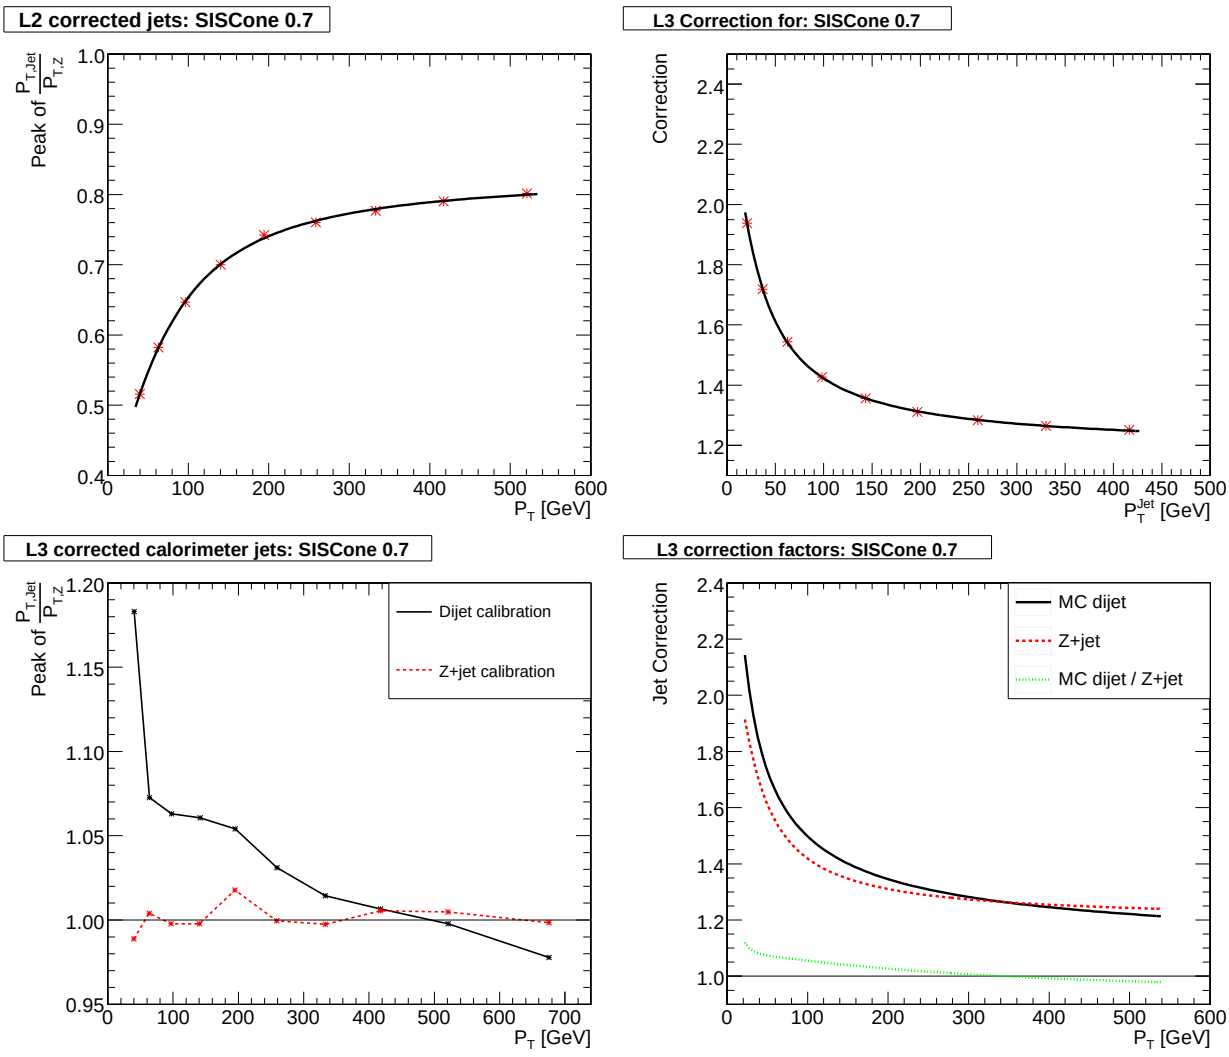

Figure A.6: *Same as Figure A.1 but for the SISCone 0.7 jet algorithm.*

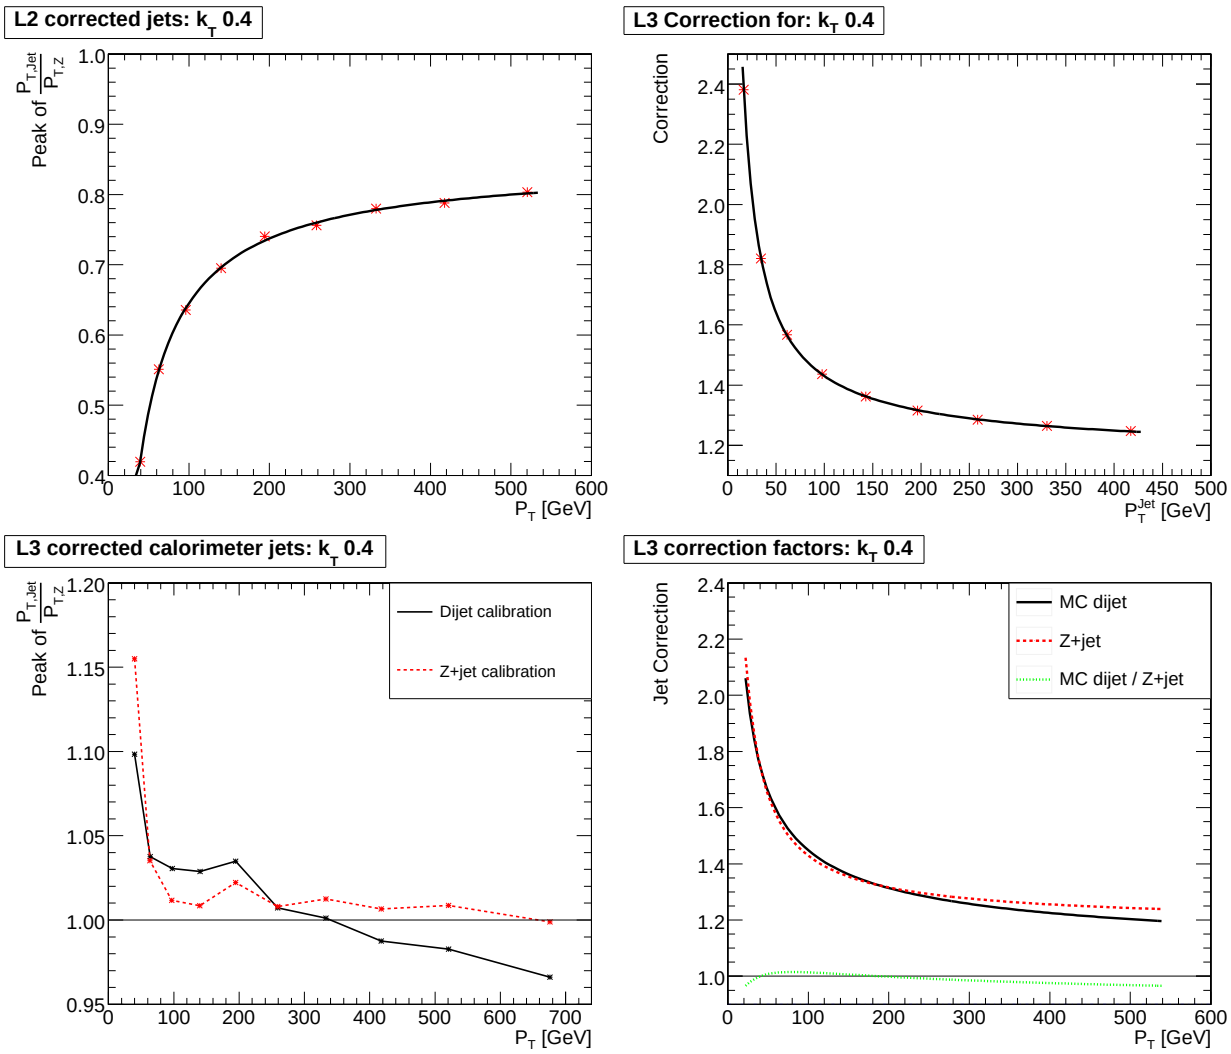

Figure A.7: *Same as Figure A.1 but for the k_T 0.4 jet algorithm.*

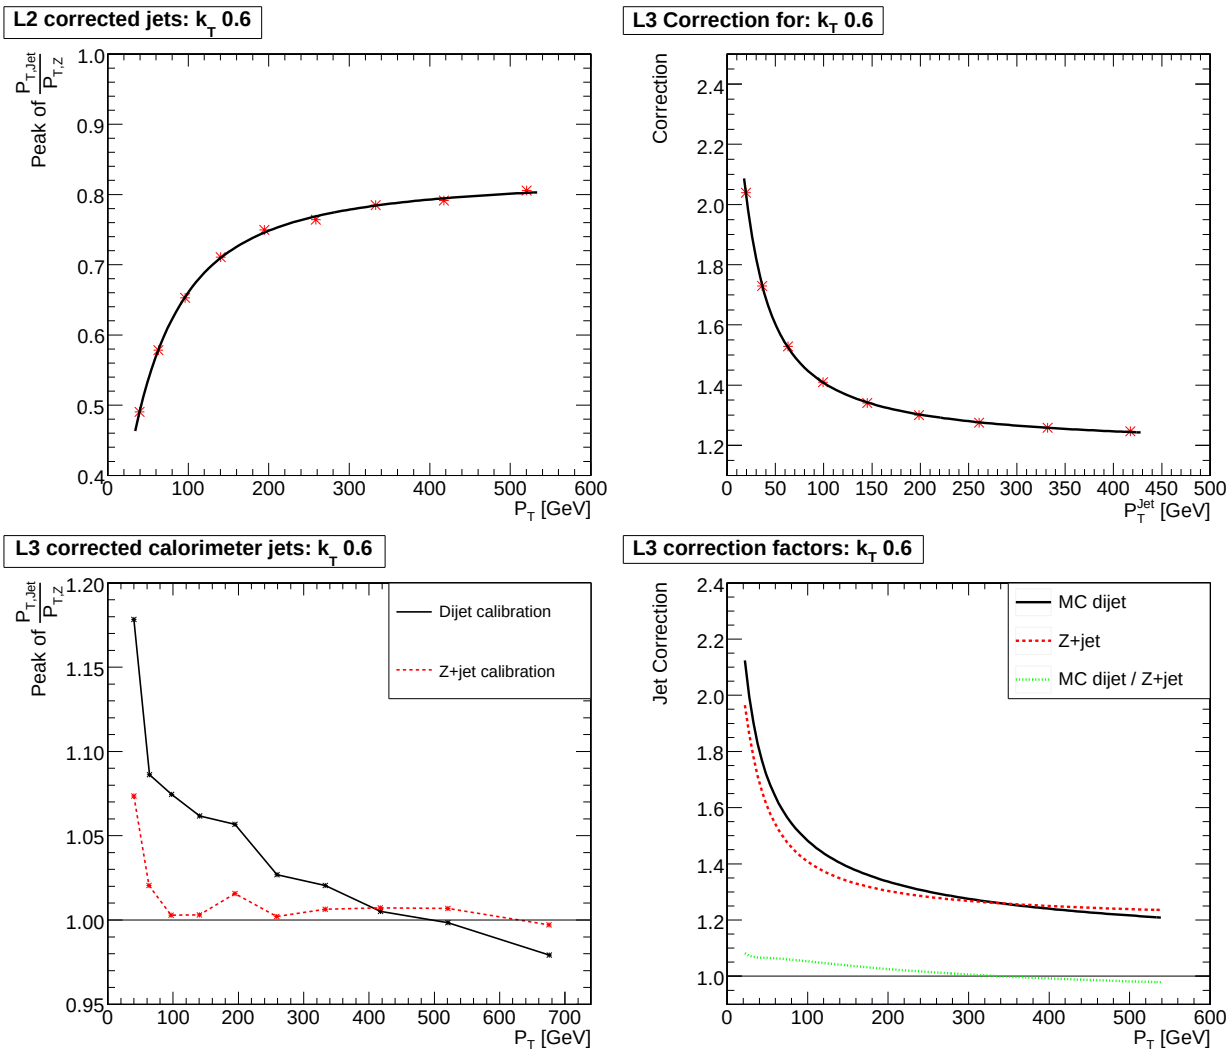

Figure A.8: *Same as Figure A.1 but for the k_T 0.6 jet algorithm.*

Definition of Jet Algorithms

Iterative Cone Algorithm

```
debugLevel    = 0
seedThreshold = 1.0
maxIterations = 100
```

The cone is defined to be stable if $\Delta\eta < 0.001$ and $\Delta\phi < 0.001$. Recombination Scheme: E_{T}

Midpoint Cone Algorithm

```
debugLevel       = 0
seedThreshold    = 1.0
coneAreaFraction = 1.0
maxPairSize      = 2
maxIterations    = 100
overlapThreshold = 0.75
```

As recombination scheme four-vector addition is used (E scheme).

SISCone Algorithm

```
coneOverlapThreshold = 0.5
splitMergeScale      = "pttilde"
maxPasses            = 0
protojetPtMin        = 0.
caching              = false
```

The cone is defined to be stable if $\Delta y < 0.001$ and $\Delta\phi < 0.001$. Recombination scheme: E scheme

SISCone Algorithm (LH)

```
coneOverlapThreshold = 0.75
splitMergeScale      = "pttilde"
maxPasses            = 0
protojetPtMin        = 0.
caching              = false
JetPtMin             = 1.
```

The cone is defined to be stable if $\Delta y < 0.001$ and $\Delta\phi < 0.001$. Recombination scheme: E scheme

k_T Algorithm

In CMSSW, the implementation of the FastJet package is used.

```
JetFinder           = "kt_algorithm"
dcut                = -1
njets               = -1
Strategy            = "Best"
UE_Subtraction      = "no"
Ghost_EtaMax        = 6
Active_Area_Repeats = 5
GhostArea           = 0.01
```

Recombination scheme: E scheme

Cambridge/Aachen Algorithm

```
JetFinder           = "cambridge_algorithm"
Dcut                = -1
njets               = -1
Strategy            = "Best"
UE_Subtraction      = "no"
Ghost_EtaMax        = 6
Active_Area_Repeats = 5
GhostArea           = 0.01
```

Recombination scheme: E scheme

Monte Carlo Datasets

C.1 Event Generation and Simulation

The Pythia event generator is interfaced to the CMS software framework CMSSW. In the following, a typical CMSSW configuration file is presented which has been used for the Monte Carlo event generation using Pythia and the subsequent detector simulation for the data samples used in this study.

In the first step, 1000 events of the process

$$\begin{aligned} \texttt{MSUB(15)} = 1 \quad &: \quad \mathrm{f}_i\bar{\mathrm{f}} \rightarrow \mathrm{g}(\gamma^*/Z) \\ \texttt{MSUB(30)} = 1 \quad &: \quad \mathrm{f}_i\mathrm{g} \rightarrow \mathrm{f}_i(\gamma^*/Z). \end{aligned}$$

with a subsequent decay of the Z boson into two muons are generated.

The parameters for the underlying events are set to the CMS default Pythia Tune DWT which is described in the next section and the hard transverse momentum is restricted to the interval of $300\,\mathrm{GeV} < \hat{p_\mathrm{T}} < 380\,\mathrm{GeV}$. After the generation, a full detector simulation of the events is performed. The resulting events are stored in the ROOT-file format and are taken as input for the digitisation and reconstruction which is implemented as an additional step for technical reasons.

```
process Sim  = {

  untracked PSet maxEvents = {untracked int32 input = 1000}

  untracked PSet configurationMetadata = {
```

```
  untracked string version = "$Revision: 1.1 $"
  untracked string name = "$Source: /cvs_server/repositories/CMSSW/  \
                           CMSSW/Configuration/CSA07Production/data/ \
                           CSA07ZmumuJets_Pt_300_380_GEN_SIM.cfg,v $"
  untracked string annotation = "ZmumuJets pt hat 300-380"
}


include "FWCore/MessageService/data/MessageLogger.cfi"
replace MessageLogger.cout.threshold = "ERROR"
replace MessageLogger.cerr.default.limit = 10

service = RandomNumberGeneratorService {

  untracked uint32 sourceSeed = 123456789
  PSet moduleSeeds = {
    untracked uint32 VtxSmeared = 98765432
    untracked uint32 g4SimHits  = 11
    untracked uint32 mix        = 12345
  }
}


source = PythiaSource {
untracked int32 pythiaPylistVerbosity = 0
untracked bool pythiaHepMCVerbosity = false
untracked int32 maxEventsToPrint = 0

untracked double crossSection = 0.2936
untracked double filterEfficiency = 1.0


PSet PythiaParameters = {

  # Vector of ParameterSet names to be read, in this order
  vstring parameterSets = {
    "pythiaUESettings",
    "processParameters"
}

include "Configuration/Generator/data/PythiaUESettings.cfi"
vstring processParameters = {
```

```
    'MSEL=0        ',  // enable selection of subprocesses
    'MSUB(15)=1    ',  // ff~ -> g(gamma* / Z)
    'MSUB(30)=1    ',  // fg  -> f(gamma* / Z)

    'MDME(174,1)=0 ',  //dd~
    'MDME(175,1)=0 ',  //uu~
    'MDME(176,1)=0 ',  //ss~
    'MDME(177,1)=0 ',  //cc~
    'MDME(178,1)=0 ',  //bb~
    'MDME(179,1)=0 ',  //tt~
    'MDME(182,1)=0 ',  //ee
    'MDME(183,1)=0 ',  //nunu
    'MDME(184,1)=1 ',  //mumu
    'MDME(185,1)=0 ',  //nunu
    'MDME(186,1)=0 ',  //tautau
    'MDME(187,1)=0 ',   //nunu

    'CKIN(3)=300.  ! minimum pt hat of the hard interactions',
    'CKIN(4)=380.  ! maximum pt hat of the hard interactions'
}


include "Configuration/StandardSequences/data/Simulation.cff"
include "Configuration/StandardSequences/data/VtxSmearedGauss.cff"
include "Configuration/StandardSequences/data/MixingNoPileUp.cff"


path p1 = { psim} # simulation

# Event output
include "Configuration/EventContent/data/EventContent.cff"
module GEN-SIM = PoolOutputModule
  {
    using FEVTSIMEventContent
    untracked string fileName = "ZmumuJets_300_380.root"
    untracked PSet dataset ={
     untracked string dataTier = "GEN-SIM"
  }
}

endpath outpath = {GEN-SIM}
```

```
  schedule = {p1,outpath}
}
```

The configuration files for the dedicated steps are stored in the DBS and are accessible via a web page[1].

C.2 Underlying Event Tunes

Different scenarios of the underlying event are investigated in this study. The set of Pythia parameters specifying the CMS default Tune DWT as well as the additional Tune S0 are listed below including a short description. More details on the different parameter are given in [106].

C.2.1 Pythia Tune DWT

```
MSTJ(11)=3      ! choice of the fragmentation function
MSTJ(22)=2      ! cut off on decay length for particle decay
PARJ(71)=10 .   ! for which ctau 10 mm
MSTP(2)=1       ! which order running alphaS
MSTP(33)=0      ! no K factors in hard cross sections
MSTP(51)=7      ! structure function chosen
MSTP(81)=1      ! multiple parton interactions (1 is Pythia default)
MSTP(82)=4      ! defines the multi-parton model
MSTU(21)=1      ! check on possible errors during program execution
PARP(82)=1.9409 ! pt cutoff for multiparton interactions
PARP(89)=1960.  ! sqrts for which PARP82 is set
PARP(83)=0.5    ! multiple interactions: matter distribution parameter
PARP(84)=0.4    ! multiple interactions: matter distribution parameter
PARP(90)=0.16   ! multiple interactions: rescaling power
PARP(67)=2.5    ! amount of initial-state radiation
PARP(85)=1.0    ! gluon prod. mechanism in MI
PARP(86)=1.0    ! gluon prod. mechanism in MI
PARP(62)=1.25   ! cut off pt value for initial state radiation
PARP(64)=0.2    ! evolution scale for initial state radiation
MSTP(91)=1      ! treatment of bean-remnants
PARP(91)=2.1    ! kt distribution
```

[1] DBS web page: http://cmsdbs.cern.ch/DBS2_discovery/

```
PARP(93)=15.0   ! cut off for pt of partons inside the proton
```

The parameter settings for the scenario of the Tune DWT without multiple interaction is similar to the listing above except the choice of the parameter `MSTP(81)` which is changed to

```
MSTP(81)=0          ! multiple interactions OFF=0, ON=1
```

C.2.2 Pythia Tune S0

```
MSTJ(11)=3          ! Choice of the fragmentation function
MSTJ(22)=2          ! cut off on decay length for particle decay
PARJ(71)=10 .       ! for which ctau 10 mm
MSTP(2)=1           ! which order running alphaS
MSTP(33)=0          ! no K factors in hard cross sections
MSTP(51) =7         ! PDF set
MSTP(52) =1         ! PDF set internal (=1) or pdflib (=2)
PARP(64) =1.0000    ! ISR renormalization scale prefactor
MSTP(70) =2         ! ISR IR regularization scheme
MSTP(72) =0         ! ISR scheme for FSR off ISR
PARP(71) =4.0000    ! FSR Q2max factor for non-s-channel procs
PARJ(81) =0.1400    ! FSR Lambda_QCD scale
MSTP(81) =21        ! UE model
PARP(82) =1.8500    ! UE IR cutoff at reference ecm
PARP(89) =1800.0000 ! UE IR cutoff reference ecm
PARP(90) =0.1600    ! UE IR cutoff ecm scaling power
MSTP(82) =5         ! UE hadron transverse mass distribution
PARP(83) =1.6000    ! UE mass distribution parameter
MSTP(88) =0         ! BR composite scheme
MSTP(89) =1         ! BR colour scheme
PARP(79) =2.0000    ! BR composite x enhancement
PARP(80) =0.0100    ! BR breakup suppression
PARP(93) =5.0000    ! BR primordial kT UV cutoff
MSTP(95) =6         ! FSI colour (re-)connection model
PARP(78) =0.2000    ! FSI colour reconnection strength
```

C.3 Data Processing on the Grid

The program grid-control, developed at the IEKP, is designed to create, submit, monitor and finally retrieve a large number of grid jobs. It is written in the programming language Python and is

Name of the dataset
/ZmumuJets_Pt_380_470/CMSSW_1_6_7-CSA07-1192835755/RECO
/ZmumuJets_Pt_0_15/CMSSW_1_6_7-CSA07-1192836125/RECO
/ZmumuJets_Pt_80_120/CMSSW_1_6_7-CSA07-1192836866/RECO
/ZmumuJets_Pt_30_50/CMSSW_1_6_7-CSA07-1192836603/RECO
/ZmumuJets_Pt_50_80/CMSSW_1_6_7-CSA07-1192836708/RECO
/ZmumuJets_Pt_2600_3000/CMSSW_1_6_7-CSA07-1192836072/RECO
/ZmumuJets_Pt_1000_1400/CMSSW_1_6_7-CSA07-1192836178/RECO
/ZmumuJets_Pt_3500_-1/CMSSW_1_6_7-CSA07-1192836656/RECO
/ZmumuJets_Pt_15_20/CMSSW_1_6_7-CSA07-1192836337/RECO
/ZmumuJets_Pt_300_380/CMSSW_1_6_7-CSA07-1192835967/RECO
/ZmumuJets_Pt_800_1000/CMSSW_1_6_7-CSA07-1192836814/RECO
/ZmumuJets_Pt_1400_1800/CMSSW_1_6_7-CSA07-1192836284/RECO
/ZmumuJets_Pt_170_230/CMSSW_1_6_7-CSA07-1192836020/RECO
/ZmumuJets_Pt_2200_2600/CMSSW_1_6_7-CSA07-1192836496/RECO
/ZmumuJets_Pt_600_800/CMSSW_1_6_7-CSA07-1192836761/RECO
/ZmumuJets_Pt_120_170/CMSSW_1_6_7-CSA07-1192836232/RECO
/ZmumuJets_Pt_3000_3500/CMSSW_1_6_7-CSA07-1192836549/RECO
/ZmumuJets_Pt_1800_2200/CMSSW_1_6_7-CSA07-1192836390/RECO
/ZmumuJets_Pt_230_300/CMSSW_1_6_7-CSA07-1192835703/RECO
/ZmumuJets_Pt_470_600/CMSSW_1_6_7-CSA07-1192835808/RECO
/ZmumuJets_Pt_20_30/CMSSW_1_6_7-CSA07-1192836443/RECO

Table C.1: *Listing of the official fully simulated CMS Pythia datasets used for this study.*

divided into modules offering the required functionality. This includes modules for the grid job itself but also the support for the particularities of jobs utilising the CMS software framework.

Once a CMSSW job is developed and tested locally, grid-control collects all required libraries, creates the jdl files and takes care of submission, monitoring and retrieval. An interface to the DBS system is also available for jobs which require input datasets from grid storage.

The actual parameters for the jobs are provided by a configuration file and are discussed exemplarily in the following with the configuration used for the analysis of the ZmumuJets dataset for this study.

This dataset is divided into 21 bins of the hard transverse momentum and each bin consists of several blocks located at different grid sites. A list of the names of the datasets is given in Table C.1.

The following configuration file tells grid-control to collect all additional libraries from the locally installed release `CMSSW_1_6_9` of the CMS software framework. The libraries are sent to the grid together with the script executing cmsRun with the configuration file `exampleFlat.cfg`.

Parameters like seeds, input datasets, number of events to process in the job and more are auto-

matically replaced in the CMSSW configuration file by grid-control. Through its DBS and DLS interface, the program splits the analysis in jobs of 4000 events and creates as many jobs as needed to process the whole dataset. In each job, a root file is produced which is transferred to the grid storage element of the IEKP for later offline analyses.

```
[DEFAULT]
; Specification of working directory
dir             = /home/vbuege/Projects/


[global]
; Specification of the execution mode
module          = CMSSW     ; CMSSW or UserMod

; Location of the configuration and output directories of the grid job
workdir         = /storage/5/vbuege/GridControleJobs/ZmumuJetsPt/ \
                  ZJetCali_All/CMSSW_1_6_9/Final/ZmumuJets_Pt_120_170


[jobs]
;jobs            = 2    ; Comment out to run over whole data set
wall time       = 10    ; minimum proxy lifetime, in hours
in flight       = 40    ; maximum number of submitted jobs

; Requirements on the grid resources
cpu time        = 10    ; in hours
memory          = 512   ; in MB


[grid]
proxy           = VomsProxy
wms             = GliteWMS   ; or Glite or LCG (depends on the used RB)
vo              = cms        ; name of the Virtual Organisation
sites           = -blah      ; Blacklist for sites


; Replace configuration files for a different Resource Broker
[lcg]
;config-vo      = %(dir)s/glite_wl_ui.conf

[glite]
```

```
;config-vo      = %(dir)s/glite_wl_ui.conf

[glite-wms]
;config         = %(dir)s/glite_wms.conf


[CMSSW]
; comment out the variable [jobs] jobs in order to
; run over all events of the dataset

project area    = %(dir)s/CMSSW_1_6_9
scram arch      = slc4_ia32_gcc345
config file     = %(project area)s/src/SkimAndFlat/exampleFlat.cfg
dataset         = /ZmumuJets_Pt_120_170/CMSSW_1_6_7-CSA07-1192836232/RECO
dbsapi          = DBSApiv2
events per job  = 4000
area files      = -.* -config lib module */data *.xml *.sql *.cf[if]
se output files = exampleSkim_flat.root
se path         = gsiftp://ekp-lcg-se.physik.uni-karlsruhe.de/ \
                  /wlcg/data/users/cms/vbuege/store5/Skim/FromGrid
                  /ZmumuJetsPt/ZJetCali_All/CMSSW_1_6_9/Final \
                  /ZmumuJets_Pt_120_170/rootfiles
gzip output     = True
seeds           = 1234 345767 23231 12312
```

Numerical Inversion of the Response

The determination of the correction factors as a function of the transverse momentum of the calorimeter jet is realised by a numerical iterative inversion of the fitted response R. The latter is a function of the mean of the transverse momentum of the Z boson $< p_{\mathrm{T}}^{\mathrm{Z}} >$.

In the following, the procedure which is used to calculate the correction factor for a dedicated $p_{\mathrm{T}}^{\mathrm{calo}}$ is presented. Starting with the mean of the transverse momentum of the Z boson in one bin, the corresponding transverse momentum of the calorimeter jet is calculated by evaluating

$$p_{\mathrm{T}}^{\mathrm{calo}} = R(< p_{\mathrm{T}}^{\mathrm{Z}} >) \cdot < p_{\mathrm{T}}^{\mathrm{Z}} > . \tag{D.1}$$

The actual inversion of this equation is done by the following iteration:

$$p_{\mathrm{T,0}}^{\mathrm{Z}} = < p_{\mathrm{T}}^{\mathrm{Z}} > \tag{D.2}$$

$$p_{\mathrm{T,i+1}}^{\mathrm{Z}} = \frac{p_{\mathrm{T}}^{\mathrm{calo}}}{R(p_{\mathrm{T,i}}^{\mathrm{Z}})} \tag{D.3}$$

The iteration ends after 50 cycles or when the calculated value of $p_{\mathrm{T,i}}^{\mathrm{Z}}$ becomes smaller than 1. This restriction is necessary due to the presence of the logarithm in Equation 6.12.

The actual C++ implementation of this procedure is given below where the fit of the response is called `RespFit`. The mean of the transverse momentum of a dedicated bin is `MeanZPt` and `C` denotes the derived correction factor.

```
double CaloPt = MeanZPt* RespFit->Eval(MeanZPt);
double ZPt;
Invert(RespFit,0,2000,CaloPt,ZPt);
double C = ZPt / CaloPt;
```

```
void Invert(TF1* RespFit, double Min, double Max,
            double CaloPt, double& ZPt) {

  int i=0;
  double tmp,r;
  ZPt = CaloPt;
  for(i=0;i<50;i++)
   {
     if (ZPt<Min)
       r = RespFit->Eval(Min);
     if (ZPt>Max)
       r = RespFit->Eval(Max);
     if (ZPt>=Min || ZPt<=Max)
       r = RespFit->Eval(ZPt);
     tmp = CaloPt/r;

// Constrained due to the presence of log(pt) in RespFit
     if (tmp<1)
       {
         ZPt = 0;
         break;
       }
     else
       ZPt = tmp;
   }
}
```

List of Figures

List of Tables

Bibliography

[1] D. J. Griffiths, *Introduction to elementary particles*. Wiley, 1987.

[2] P. Schmüser, *Feynman-Graphen und Eichtheorien für Experimentalphysiker*. Springer, 2. neubearbeitete Auflage ed., 1995.

[3] W.-M. Yao *et al.*, Review of Particle Physics, *Journal of Physics G* **33** (2006).

[4] B. Povh, *Teilchen und Kerne*. Springer, 1993.

[5] The Coordinated Theoretical-Experimental Project on QCD - CTEQ http://www.phys.psu.edu/ cteq/.

[6] MRST Parton Distributions http://durpdg.dur.ac.uk/hepdata/mrs.html.

[7] Durham University - Online Plotting and Calculation http://durpdg.dur.ac.uk/hepdata/pdf3.html.

[8] The LEP Electroweak Working Group http://lepewwg.web.cern.ch/LEPEWWG/.

[9] S. Catani et al., *QCD*. CERN-TH/2000-131 [arXiv:hep-ph/0005025v1], 2000.

[10] S. Alekhin, *The NNLO predictions for the rates of the W/Z production in pp and pbar p collisions*. arXiv:hep-ph/0307219v1, 2003.

[11] A. D. Martin, R. G. Roberts, W. J. Stirling, and R. S. Thorne, *Parton Distributions and the LHC: W and Z Production*. Eur.Phys.J. C14 (2000) 133-145 [arXiv:hep-ph/9907231v1], 1999.

[12] *Proceedings of the Workshop on Standard Model Physics (and more) at the LHC (CERN 2000-004)*, 2000.

[13] The International Linear Collider Project Homepage
http://www.linearcollider.org.

[14] Fermi National Accelerator Facility
http://www.fnal.gov/.

[15] The Collider Detector at Fermilab Experiment (CDF) - Homepage
http://www-cdf.fnal.gov/.

[16] The DØExperiment
http://www-d0.fnal.gov/.

[17] K. Desler and D. Edwards, *Accelerator Physics of Colliders*.
http://pdg.lbl.gov/2002/accelrpp.pdf, 2001.

[18] The CMS Collaboration, *CMS Physics Technical Design Report Volume I: Detector Performance and Software*. CERN/LHCC 2006-001, 2006.

[19] CERN Communication Group, *CERN LHC Guide faq*.
http://cdsweb.cern.ch/record/1092437/files/CERN-Brochure-2008-001-Eng.pdf, 2008.

[20] ALICE - A Large Ion Collider Experiment at CERN LHC
http://aliceinfo.cern.ch/.

[21] The Large Hadron Collider beauty experiment
http://lhcb.web.cern.ch/lhcb/.

[22] The ATLAS Experiment
http://www.atlas.ch.

[23] CMS Outreach
http://cmsinfo.cern.ch/outreach.

[24] The CMS Collaboration, *The CMS Experiment at the CERN LHC*.
http://cmsdoc.cern.ch/cms/Publications/detpaper/finaldraft.pdf, 2008.

[25] The CMS Collaboration, *The Tracker Project Technical Design Report*.
CERN/LHCC 98-006, 1998.

[26] The CMS Collaboration, *Addendum to the CMS Tracker TDR*. CERN/LHCC 200-016, 2000.

[27] The CMS Collaboration, *The CMS electromagnetic calorimeter project: Technical Design Report*. CERN-LHCC-97-033, 1997.

[28] The CMS Collaboration, *The CMS hadron calorimeter project: Technical Design Report*. CERN-LHCC-97-031, 1997.

[29] The CMS Collaboration, *CMS MUON Technical Design Report*. CERN-LHCC-97-032, 1997.

[30] The CMS Collaboration, *CMS TriDAS project: Technical Design Report*. CERN-LHCC-2000-038, 1997.

[31] Updated CMS Computing Resource Requests 2008 - 2010 and Resources Requirements for 2011 and 2012
http://lcg.web.cern.ch/LCG/MB/revised_resources/aug07/CMS-resources-070824-final.pdf.

[32] The LCG TDR Editorial Board, Chair: J. Knobloch, *LHC Computing Grid Technical Design Report*. CERN - ISBN 92-9083-253-3, 2005.

[33] LHC Computing Grid - List of Tier1 Centres (July 2007)
http://lcg.web.cern.ch/LCG/C-RRB/Tier-1/ListTier1Centres_16jul07.pdf.

[34] Standard Performance Evaluation Corporation
http://www.spec.org/.

[35] The CMS Collaboration, *CMS Computing Technical Design Report*. CERN-LHCC-2005-023, ISBN 92-9083-252-5, 2005.

[36] LHC Computing Grid - List of Tier2 Centres (April 2008)
http://lcg.web.cern.ch/LCG/C-RRB/Tier-2/ListTier2Centres04Apr08.pdf.

[37] Memorandum of Understanding for Collaboration in the Deployment and Exploitation of the Worldwide LHC Computing Grid
http://lcg.web.cern.ch/LCG/C-RRB/MoU/WLCGMoU.pdf, April 2008.

[38] Enabling Grids for E-sciencE Homepage
http://public.eu-egee.org/.

[39] R. Housley, W. Polk, W. Ford, and D. Solo, *Internet X.509 Public Key Infrastructure - Certificate and Certificate Revocation List (CRL) Profile*. The Internet Society - http://tools.ietf.org/html/rfc3280, 2002.

[40] The EUGridPMA - coordinating grid authentication in e-Science
http://www.eugridpma.org/.

[41] DFN-CA
https://www.pki.dfn.de/.

[42] GridKa-CA Karlsruhe
http://grid.fzk.de/ca/.

[43] EGEE User's Guide - VOMS CORE SERVICES
https://edms.cern.ch/file/571991/1/voms-guide.pdf.

[44] The EGEE Information System
https://twiki.cern.ch/twiki//bin/genpdf/EGEE/InformationSystem.

[45] The Berkeley Database Information Index
https://twiki.cern.ch/twiki//bin/genpdf/EGEE/BDII.

[46] GLUE schema Homepage
http://glueschema.forge.cnaf.infn.it/.

[47] LFC General Description
https://twiki.cern.ch/twiki/bin/view/LCG/LfcGeneralDescription.

[48] EGEE User's Guide - WMS SERVICE
https://edms.cern.ch/file/572489/1/WMS-guide.pdf.

[49] The Globus Toolkit: Overview of the Grid Security Infrastructure
http://www.globus.org/security/overview.html.

[50] GLITE 3.1 USER GUIDE
https://edms.cern.ch/file/722398//gLite-3-UserGuide.pdf.

[51] dCache, the Overview
http://www.dcache.org/manuals/dcache-whitepaper-light.pdf.

[52] B. Tierney, R. Aydt, D. Gunter, W. Smith, M. Swany, V. Taylor, and R. Wolski, *A Grid Monitoring Architecture*. GGF Performance Working Group - http://www-didc.lbl.gov/GGF-PERF/GMA-WG/papers/GWD-GP-16-2.pdf, 2000.

[53] Open Grid Forum
http://www.gridforum.org/.

[54] CMS Dataset Bookkeeping System
https://twiki.cern.ch/twiki/bin/view/CMS/DBS-TDR.

[55] CMS Data Location Service
https://twiki.cern.ch/twiki/bin/view/CMS/DLS.

[56] PhEDEx - CMS Data Transfers
http://cmsdoc.cern.ch/cms/aprom/phedex/prod/Info::Main.

[57] CMS Remote Analysis Builder
http://cmsdoc.cern.ch/cms/ccs/wm/www/Crab/.

[58] The Alpha Magnetic Spectrometer Experiment (AMS) - Homepage
http://ams.cern.ch/.

[59] Cluster Resources Homepage
http://www.clusterresources.com/.

[60] V. Büge, U. Felzmann, C. Jung, H. Stadie, and A. Vest, *Integrating the IEKP Linux cluster as a Tier-2/3 prototype centre into the LHC Computing Grid*. IEKP–KA/2006–3, 2006.

[61] Intel Homepage
http://www.intel.com.

[62] SUSE Linux Enterprise
http://www.novell.com/linux/.

[63] Lustre Cluster File System
http://www.lustre.org/.

[64] Simple Linux Resource Management System
https://computing.llnl.gov/linux/slurm/.

[65] Scientific Linux - Homepage
https://www.scientificlinux.org/.

[66] Red Hat Enterprise Linux
http://www.redhat.com/rhel/.

[67] Research project Billing the Grid: Development and implementation of incentive mechanisms for efficient trading with grid-resources
http://www.iw.uni-karlsruhe.de/Billing_Grid/indexen.html.

[68] R. Adair, R. Bayles, L. Comeau, and R. Creasy, *A virtual machine system for the 360/40*. IBM Cambridge Scientific Center report 320-2007, Cambridge, MA, May, 1966.

[69] VMware White Paper
Understanding Full Virtualization, Paravirtualization, and Hardware Assist, 2007
http://www.vmware.com/files/pdf/VMware_paravirtualization.pdf.

[70] G. Bengel, C. Baun, M. Kunze, and K.-U. Stucky, *Masterkurs Parallele und Verteilte Systeme*. Vieweg+Teubner, ISBN-13 978-3834803948, 2008.

[71] M. T. Jones, *Virtual Linux - An overview on virtualization methods, architectures and implementations*. IBM - http://www.ibm.com/developerworks/library/l-linuxvirt/, 2006.

[72] Bochs IA-32 Emulator Project
http://bochs.sourceforge.net/.

[73] QEMU Homepage
http://fabrice.bellard.free.fr/qemu/.

[74] Parallels Virtuozzo Containers Homepage
http://www.parallels.com/en/products/virtuozzo/.

[75] OpenVZ Hompepage
http://openvz.org/.

[76] Microsoft Virtual PC Homepage
http://www.microsoft.com/windows/products/winfamily/virtualpc/default.mspx.

[77] VMware Homepage
http://www.vmware.com/.

[78] Xen Homepage
http://xen.org/.

[79] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, *Xen and the Art of Virtualization*. Proceedings of SOSP'03, October 19-22, 2003.

[80] VMware White Paper
Virtualization: Architectural Considerations and other evaluation criteria, 2005
http://www.vmware.com/pdf/virtualization_considerations.pdf.

[81] Citrix XenServer
http://www.citrixxenserver.com/Pages/default.aspx.

[82] The User-mode Linux Kernel Home Page
http://user-mode-linux.sourceforge.net/.

[83] Intel Virtualization Technology
http://www.intel.com/technology/virtualization/index.htm.

[84] AMD-V Virtualisation Technology
http://www.amd.com/us-en/Processors/
ProductInformation/0,,30_118_8796_14287,00.html.

[85] K. Adams and O. Agesen, *A Comparison of Software and Hardware Techniques for x86 Virtualization*. Proceedings of ASPLOS'06, October 21-25, 2006,, 2006.

[86] The Trac Project Homepage
http://trac.edgewall.org/.

[87] Squid-cache.org Homepage
http://www.squid-cache.org/.

[88] Debian Homepage
http://www.debian.org/.

[89] The SAMGrid Project - Homepage
http://projects.fnal.gov/samgrid/project/project.html.

[90] P. Reisner and L. Ellenberg, *DRBD v8 - Replicated Storage with Shared Disk Semantics*. Proceedings of the 12th International Linux System Technology Conference October 11-14, 2005 at University of Hamburg, Germany, 2005.

[91] V. Büge, Y. Kemp, M. Kunze, and G. Quast, *Application of Virtualisation Techniques at a University Grid Center*. e-Science'06 - Second IEEE International Conference on e-Science and Grid Computing, ISBN 0-7695-2734-5, 2006.

[92] GridKA School 2006 Hompage
http://www.fzk.de/gks06.

[93] The D-Grid initiative
http://www.dgrid.de/.

[94] GridKA School 2007 Hompage
http://www.fzk.de/gks07.

[95] T. Sjöstrand, L. Lönnblad, and S. Mrenna, *PYTHIA 6.2 Physics and Manual*. arXiv:hep-ph/0108264, 2001.

[96] S. Agostinelli et al., *G4 - a simulation toolkit*. Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment, Volume 506, Issue 3, 1 July 2003, Pages 250-303, 2003.

[97] GalProp Hompage
http://galprop.stanford.edu/web_galprop/galprop_home.html.

[98] N. Patwardhan, E. Siever, and S. Spainhour, *Perl in a nutshell*. O'Reilly, 2. auflage ed., 2003.

[99] V. Büge, Y. Kemp, M. Kunze, O. Oberst, and G. Quast, *Virtualizing a Batch Queuing System at a University Grid Center*. Springer, Lecture Notes in Computer Science, ISBN 978-3-540-49860-5, 2006.

[100] Platform LSF Product Homepage
http://www.platform.com/Products/platform-lsf.

[101] The EGEE-II Collaboration, *An EGEE Comparative Study: Grids and Clouds - Evolution or Revolution?* https://edms.cern.ch/file/925013/3/EGEE-Grid-Cloud.pdf, 2008.

[102] CMSSW Application Framework
https://twiki.cern.ch/twiki/bin/view/CMS/WorkBookCMSSWFramework.

[103] G.Corcella, I. G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M. H. Seymour, and B. R. Webber, *HERWIG 6.5: an event generator for Hadron Emission Reactions With Interfering Gluons*. arXiv:hep-ph/0011363v3, 2002.

[104] M. L. Mangano, M. Moretti, F. Piccinini, R. Pittau, and A. D. Polosa, *ALPGEN, a generator for hard multiparton processes in hadronic collisions*. arXiv:hep-ph/0206293v2, 2003.

[105] J. Alwall, P. Demin, S. de Visscher, R. Frederix, M. Herquet, F. Maltoni, T. Plehn, D. L. Rainwater, and T. Stelzer, *MadGraph/MadEvent v4: The New Web Generation*. arXiv:0706.2334v1, 2007.

[106] T. Sjostrand, S. Mrenna, and P. Skands, *Pythia 6.4 physics and manual*, *JHEP* **0605** (2006) 026.

[107] M. Heinrich, *The Influence of Hadronisation and Underlying Event on the Inclusive Jet Cross-Section at the LHC*. IEKP-KA/2008-11, 2007.

[108] G. Bruno, P. T. Cox, S. Lacaprara, N. Neumeister, B. V. de Vyver, S. Villa, and R. Wilkinson, *Local Reconstruction in the Muon Detectors*. CMS NOTE 2002/043, 2002.

[109] M. Mulders, *Muon Reconstruction and Identification at CMS*. CMS CR-2006/101, 2006.

[110] P. Paolucci, *The CMS Muon system*. CMS NOTE 2006/006, 2006.

[111] C. Buttar, J. D'Hondt, M. Kramer, G. Salam, and M. Wobisch et. al, *Standard Model Handles and Candles Working Group: Tools and Jets Summary Report*. arXiv:0803.0678, 2007.

[112] A. Bhatti, F. Chlebana, G. Dissertori, C. Dragoiu, R. Harris, M. Jha, K. Kousouris, P. Kurt, Z. Qi, F. Ratnikov, P. Schieferdecker, H. Topakli, N. Varelas, and M. Zielinski, *Performance of the SISCone Jet Clustering Algorithm*. CMS AN-2008/002, 2008.

[113] G. C. Blazey, J. R. Dittmann, S. D. Ellis, V. D. Elvira, K. Frame, S. Grinstein, R. Hirosky, R. Piegaia, H. Schellman, R. Snihur, V. Sorin, and D. Zeppenfeld, *Run II Jet Physics: Proceedings of the Run II QCD and Weak Boson Physics Workshop*. arXiv:hep-ex/0005012, 2000.

[114] A. Heister, O. Kodolova, V. Konoplianikov, S. Petrushanko, J. Rohlf, C. Tully, and A. Ulyanov, *Measurement of Jets with the CMS Detector at the LHC*. CMS NOTE 2006/036, 2006.

[115] G. P. Salam and G. Soyez, *A practical Seedless Infrared-Safe Cone jet algorithm*. JHEP 05 (2007) 086 [arXiv:0704.0292 [hep-ph]], 2007.

[116] J. M. Butterworth, J. P. Couchman, B. E. Cox, and B. M. Waugh, *KtJet: A C++ implementation of the Kt clustering algorithm*. arXiv:hep-ph/0210022v1, 2002.

[117] M. Cacciari and G. P. Salam, *Dispelling the N^3 myth for the Kt jet-finder*. Phys.Lett. B641 (2006) 57-61 [arXiv:hep-ph/0512210v2], 2006.

[118] L. Apanasevich, A. Bhatti, F. Chlebana, C. Dragoiu, R. Harris, M. Jha, K. Kousouris, P. Kurt, Z. Qi, F. Ratnikov, P. Schieferdecker, A. Smoron, H. Topakli, N. Varelas, and M. Zielinski, *Performance of Jet Algorithms in CMS*. CMS AN-2008/001, 2008.

[119] ROOT - An Object-Oriented Data Analysis Framework http://root.cern.ch/.

[120] A. Bhatti, V. Büge, J. Cammin, and R. C. et al., *Plans for Jet Energy Corrections at CMS*. CMS AN-2007/055, to be published.

[121] R. Harris and K. Kousouris, *MC Truth L2 & L3 Factorized Jet Corrections at CMS*. CMS AN-2008/003, 2008.

[122] The CMS Collaboration, *CMS Physics Technical Design Report Volume 2: Physics Performance*. CERN/LHCC 2006-021, 2006.

[123] N. Adam, J. Alcaraz, N. Amapane, and R. Bellan et al., *Towards a measurement of the inclusive $W \to \mu\nu$ and $Z \to \mu^+\mu^-$ cross sections in pp collisions at $\sqrt{s} = 14$ TeV*. CMS AN-2007/031, 2008.

Danksagung

An erster Stelle möchte ich mich bei Herrn Professor Dr. Günter Quast für die hervorragende Betreuung und sein Engagement bedanken. Er ließ mir stets Raum zur Entwicklung eigener Ideen und nahm sich immer die Zeit, offene Fragen zu diskutieren. Beides hat viel zum tieferen Verständnis der Physik und ihrer Didaktik und damit zum Gelingen dieser Arbeit beigetragen.

Ebenso herzlich danke ich Herrn Professor Dr. Michael Feindt für die Übernahme des Korreferats sowie für seine hilfreichen Kommentare zu dieser Doktorarbeit.

Meine Dankbarkeit gilt Herrn Dr. habil. Marcel Kunze für seine Betreuung und die zahlreichen Diskussionen im Bereich Grid-Computing und Virtualisierung sowie dem Forschungszentrum Karlsruhe für die Finanzierung dieses Projekts.

In besonderer Weise danke ich zudem Herrn Dr. Klaus Rabbertz für die enge und vertrauensvolle Zusammenarbeit. Er hatte stets ein offenes Ohr für mich und stand mir bei Fragen immer beratend zur Seite. Auch sein gründliches Korrekturlektorat und die damit verbundenen konstruktiven Anmerkungen haben zur Verbesserung dieser Arbeit beigetragen.

Den Mitgliedern der CMS Gruppe, Dr. Ulrich Felzmann, Christoph Hackstein, Michael Heinrich, Dr. Christopher Jung, Dr. Yves Kemp, Benjamin Klein, Viktor Mauch, Oliver Oberst, Andreas Oehler, Danilo Piparo, Christophe Saout, Armin Scheurer, Fred Stober und Manuel Zeise gilt mein Dank für die gute und freundschaftliche Zusammenarbeit, die angenehme Arbeitsatmosphäre, die vielfältigen Diskussionen und das Korrekturlesen meiner Ausarbeitungen. Dieser schließt die Mitglieder des EKP-Admin-Teams mit ein, ohne deren unermüdlichen Einsatz die Durchführung dieser Arbeit kaum möglich gewesen wäre. Ebenso möchte ich mich bei den Mitarbeitern des Instituts für Wissenschaftliches Rechnen sowie den Sekretariaten beider Institute bedanken.

Furthermore, I express my gratitude to the QCD, Electroweak and Jet Energy Correction working groups of CMS for the fruitful discussions. This cooperative and constructive collaboration contributed to improve this analysis.

Die angenehme Atmosphäre am Institut und in der Arbeitsgruppe hat dazu beigetragen, dass ich mich in den vergangenen Jahren sehr wohl gefühlt habe. Neben der Zusammenarbeit schließt dies auch die vielen Freundschaften mit ein, die sich in dieser Zeit entwickelt haben und mir sehr viel bedeuten. Hierfür allen ein herzliches Dankeschön!

Auch möchte ich die Gelegenheit nutzen, um mich an dieser Stelle bei meinen Freunden für ihre Unterstützung in den letzten Jahren zu bedanken. Vielen Dank für die vielen schönen Stunden und Eure Hilfe!

Mein ganz besonderer Dank gilt meinen Eltern und meiner Frau, die mich in jeder Hinsicht so unglaublich unterstützt, motiviert und gefördert haben. Danke, dass Ihr mich zur richtigen Zeit auf andere Gedanken gebracht habt, so dass ich immer wenn es notwendig war einen freien Kopf behalten konnte. Euer Rückhalt war und ist für mich ein unschätzbares Gut!