

Jet Momentum Resolution for the CMS Experiment and Distributed Data Caching Strategies

Christoph Alexander Heidecker

Zur Erlangung des akademischen Grades eines

DOKTORS DER NATURWISSENSCHAFTEN
(Dr. rer. nat.)

von der KIT-Fakultät für Physik des
Karlsruher Instituts für Technologie (KIT)

angenommene

DISSERTATION

von

MSc. Christoph Alexander Heidecker
aus Stuttgart

Tag der mündlichen Prüfung: 10.07.2020

Referent: Prof. Dr. Günter Quast

Korreferent: Priv.-Doz. Dr. Klaus Rabbertz

Betreuender Mitarbeiter: Dr. Manuel Giffels



This document is licensed under a Creative Commons
Attribution-ShareAlike 4.0 International License (CC BY-SA 4.0)
<https://creativecommons.org/licenses/by-sa/4.0/deed.en>

Contents

1	Introduction	5
2	Particle physics at the terascale	9
2.1	The Standard Model of Particle Physics	9
2.1.1	Quantum Chromodynamics and jets	11
2.2	Monte-Carlo simulation	13
2.3	Z bosons at hadron colliders	15
3	The experimental environment	19
3.1	Accelerating particles at the Large Hadron Collider	19
3.2	Detecting particles with the Compact Muon Solenoid	21
3.2.1	Coordinate system of the detector	22
3.2.2	Pile-up interactions	24
3.2.3	The tracking system	26
3.2.4	The electromagnetic calorimeter	27
3.2.5	The hadronic calorimeter	28
3.2.6	The superconducting solenoid and the muon system	29
3.2.7	The Trigger and Data Acquisition System	29
4	Jet measurement and calibration at CMS	31
4.1	Reconstruction of physical objects by the Particle-Flow algorithm	33
4.2	Jet clustering	34
4.3	Pile-up mitigation	36
4.4	Balancing methods used for jet calibration	37
4.4.1	p_T -balance method	38
4.4.2	Missing E_T Projection Fraction method	38
4.5	Factorized calibration approach for jet momentum scale at CMS	39
4.5.1	Determination of pile-up offset corrections	40
4.5.2	Determination of simulated detector response corrections	41

4.5.3	Determination of relative residual corrections	41
4.5.4	Determination of absolute residual corrections	42
4.5.5	Determination of flavor corrections	43
4.6	Estimation of the jet momentum resolution at CMS	43
4.6.1	Extracting the jet momentum resolution using dijet events . .	44
4.6.2	Extracting the jet momentum resolution using γ + jets events	45
5	Preparations for calibrating jets using Z + jets events	49
5.1	Data samples	50
5.2	Data processing workflow	51
5.3	Event selection	53
5.3.1	Data certification	54
5.3.2	Pile-up reweighting	54
5.3.3	Selection criteria for leptons and reconstructed Z boson . . .	55
5.3.4	Selection criteria for jets	57
5.3.5	Missing transverse momentum corrections	58
5.4	Kinematic studies of Z+jet events	59
5.4.1	Characteristics of muons and Z bosons	59
5.4.2	Jet characteristics	61
5.4.3	E_T^{miss} characteristics	63
5.5	Study of the jet response using Z + jet events	63
5.5.1	Comparison of p_T -balance and Missing E_T Projection Fraction (MPF) method	64
5.5.2	Study of second jet activity	66
5.5.3	Time dependency of the jet response	67
6	Determination of absolute residual corrections using Z + jet events	71
6.1	Extrapolation to zero second jet activity	71
6.2	Jet energy scale corrections extracted from Z + jet events	73
6.3	CMS combination of absolute jet energy corrections by global fit . .	73
6.4	Validation of the absolute residual jet energy corrections	74
7	Extraction of jet momentum resolution from Z + jet events	77
7.1	Proposed extraction procedure for Z + jets events	77
7.2	Extraction of response variance	78
7.3	Resolution estimation using quadratic subtraction	82
7.4	Extrapolation to zero second jet activity	83
7.5	Validation with Monte-Carlo (MC) closure test	84
7.6	MC correction by jet momentum resolution scale factors	86

8	High-throughput computing in high energy physics	91
8.1	The Worldwide LHC Computing Grid	93
8.2	Dynamic Tier 3 landscape	94
8.3	Processing end-user analyses	96
8.4	Data throughput optimization via data locality	96
8.4.1	Data locality within High Energy Physics (HEP) computing	97
8.4.2	Existing data locality approaches for end-user analyses	98
9	Accelerating end-user analyses with coordinated distributed caching	101
9.1	Distributed coordinated caching	102
9.1.1	Data selection and placement	103
9.1.2	Coordination of jobs to distributed caches	104
9.1.3	Optimization of overall data throughput	105
9.2	Realization of distributed coordinated caching	107
9.2.1	Prerequisites of the implementation	107
9.2.2	Conceptual design of the coordination service	109
9.2.3	Scalability of data and job coordination	112
9.3	Benchmarking of prototype systems	113
9.3.1	Small-scale testing setup	114
9.3.2	Benchmarking data throughput and coordination efficiency	115
9.3.3	Benefits for data-intensive HEP analysis workflows	118
9.3.4	Applicability for shared computing resources	120
9.3.5	Dedicated high throughput computing resources	124
10	Conclusion and outlook	131
A	Appendix	135
A.1	Technical information for analysis	135
A.1.1	Data sets	135
A.1.2	Data certification	138
A.1.3	Muon Identification	138
A.1.4	Jet identification	138
A.1.5	Characteristics of muons, Z boson, and jets for the MadGraph data set	140
A.1.6	Manually removed or tuned bins during the jet momentum resolution extraction	144
A.1.7	Additional plots for jet momentum resolution MC closure test	148
A.1.8	Additional plots for derived jet momentum scale factors	150
A.2	Computing	152
A.2.1	Draft for a general coordination service	152

Contents

List of Figures	155
List of Tables	165
Bibliography	167
Acronyms	177

Chapter 1

Introduction

Human curiosity has long been a driving factor in the search for the basic principles of nature. It makes us explore the world we live in, from the galaxies far away to the innermost parts of matter.

In the past, the search for the basic building blocks of matter led to the realization of huge experiments. It culminates in the construction of the Large Hadron Collider (LHC), the largest and most powerful particle accelerator and collider at the time of writing this thesis. With a center-of-mass energy of 13 TeV within proton-proton collisions, it outperforms previous experiments. Such collisions mimic the conditions of the early universe with its gigantic energy density. This makes it possible to create particles that are not observable under regular terrestrial conditions and to measure their properties via their decay products. Huge detectors, such as the Compact Muon Solenoid (CMS), build around a collision point are capable of tracing the decay products back to the original process of interest.

The precise reconstruction of the underlying processes requires a precise measurement of the particles in the detector. But, various influences such as collateral processes, detector effects, or peculiarities of the algorithms used for reconstruction disturb these measurements.

A process that is less affected by some of the issues is the production of a Z boson that is additionally produced at an unprecedented rate at the LHC. Due to previous precision experiments, its properties are well-understood. Supplementarily, the possible decay products of the Z boson, two leptons, leave a clear signature in the detector that allows to accurately reconstruct it. As a consequence, it suits perfectly for calibration purposes.

One of the most important objects for most CMS analyses and the understanding of the fundamentals of the nature of the smallest building blocks are jets. These are collimated streams of particles that originate from a single gluon or quark. These compound objects are the most challenging in reconstruction since there are plenty of effects that bias their measurement. The CMS collaboration uses a sequence of sophisticated techniques to calibrate the jet energy, each of them dealing with a certain systematic distortion.

One of the most important calibration steps is the data-driven calibration with balancing techniques. Thereby, the analysis exploits the conservation of the transverse momentum and balances a jet to a well-measured reference object to correct the jet momentum.

This thesis concentrates on using Z + jets events to compensate for residual differences between MC simulated and recorded jets. By focusing on the decay of the Z boson into two muons, it exploits the strength of the CMS detector to precisely identify and reconstruct muons. Furthermore, this thesis proposes a new procedure for determining the jet momentum resolution using Z + jets events. It is expected that the new approach will allow cross-checks with established measurement approaches that use other production channels to detect and reduce channel-specific distortions. Besides, by combining the newly derived results with the established measurement procedures, the uncertainties on the jet momentum resolution can be decreased.

Furthermore, this thesis also focuses on the advantages of caching for data-intensive HEP end-user analyses. The calibration of the jet energy scale and the momentum resolution requires large input data sets to be processed iteratively. Shorten turnaround cycles is crucial since the results of multiple analysis groups need to be optimized and combined.

For HEP analysis workflows, the caching of input data, a common approach for boosting data-intensive workflows, suits well. However, conventional caches placed within the usually distributed HEP computing infrastructure may not be used efficiently. The research presented in this thesis suggests to coordinate data in the caches and schedule analysis workflows to the computing resource that already provides cached input data. This distributed coordinated caching concept is expected to avoid wasting limited cache space and to improve data processing as workflows processing data are served by caches close-by.

This thesis realizes this distributed coordinated caching concept in building dedicated high-throughput computing clusters, which allows to improve the performance for data-intensive HEP analysis workflows. Additionally, it studies the applicability of

the concept for shared resources such as High-Performance Computing (HPC) clusters and cloud resources to make them available for data-intensive processing.

Improving the processing of data-intensive workflows aims at shortening turnaround cycles and thus deriving physics results, e.g. the jet calibration results, faster. Simultaneously, the processing of large amounts of data on a shorter time scale is of general interest as it allows the HEP community to face the amount of data that current and future HEP experiments are recording.

Chapter 2

Particle physics at the terascale

Mankind has always searched for the basic building blocks of matter. Particularly in the last century, knowledge of the subatomic structure of matter has greatly developed, leading to a zoo of elementary particles and particle clusters that have been discovered and studied. This led to the development of a complex theory, commonly known as the *Standard Model of particle physics*, which today describes the known elementary particles and their interactions. Measurements in particle physics either confirmed theoretical predictions derived from the Standard Model or led to extensions that could be successfully included. The last extension for example introduced a mechanism for assigning masses to particles, the so-called Higgs mechanism. The discovery of its inevitably associated particle, the Higgs boson, in 2012 confirmed the extension of the model proposed in 1964. Such interplay between theoretical predictions and experimental measurements is a central component of particle physics.

The next section gives a brief overview of the Standard Model of particle physics as it stands at the time of writing this paper.

2.1 The Standard Model of Particle Physics

The basic principles of the Standard Model of particle physics are based on quantum field theory and gauge theory. By combining quantum mechanics with special relativity and assuming certain gauge symmetries, the Standard Model describes the fundamental particles of matter and their interactions within one Lagrangian

formula. At the time of writing this article, it already describes a broad spectrum of phenomena observed in the microcosm using basic principles.

The concept of the Standard Model is nowadays based on a small set of elementary particles, which are illustrated in Figure 2.1. There are two classes of particles: The




Figure 2.1: Overview of the fermions, the gauge bosons, and the Higgs boson within the Standard Model of particle physics [1]. The gauge bosons reflect three of the four fundamental forces. Gravity is still missing in the model. Fermions are classified into quarks and leptons according to the forces to which they are responsive.

basic building blocks of matter, so-called *fermions*, and the mediators of forces, the so-called *bosons*. While the elementary bosons are classified by the force they carry, the fermions are grouped by the forces that act on them.

The forces the mankind is currently aware of can be reduced to four fundamental ones: *Gravitational*, *electromagnetic*, *weak nuclear*, and *strong nuclear* force. By now, the Standard Model of particle physics includes the latter three. Gravitation, which is sufficiently describing particle collider physics, is currently not included in the Standard Model. As it is extremely weak on small scales, its influence on the particle production within collider experiments can be neglected.

Introduced by gauge symmetries each of the remaining three forces manifests in gauge bosons that couple to a certain charge. All fermions experience the weak nuclear force, which is mediated by the positively and negatively charged W bosons and the neutral Z boson. The weak interaction coupling to the weak isospin allows for changing the type of the fermion, the so-called particle flavor. A *photon*, the mediator of the electromagnetic force couples to charged particles. The mediator of the strong force, the *gluon*, couples to color charge. Fermions that carry color charge and thus are subject to the strong force are defined as *quarks*, whereas the remaining fermions are named *leptons*.

Most matter in the universe is build of the lightest and therefore stable fermions, which includes up (u) and down (d) quarks as well as the electrons and neutrinos as the corresponding leptons. The heavier muons, however, should also be mentioned here, as they live long enough to be detected and are of special interest for the presented analysis.

The theory of electroweak interactions and Quantum Chromodynamics (QCD), the theoretical constructs behind the gauge bosons assume massless particles. This inconsistency of the Standard Model with the observations of the massive gauge particles, the Z and W bosons, was eliminated by the Higgs mechanism. The electroweak symmetry breaking within the Higgs-mechanism assigns masses to the gauge bosons. Fermions acquire mass by coupling to the so-called Higgs-field. The self-excitation of this field is observed as the so-called Higgs-boson.

2.1.1 Quantum Chromodynamics and jets

The Quantum Chromodynamics (QCD), which describes the strong force, initiates special phenomenological characteristics that make interactions of quarks and gluons behave differently compared to the electromagnetic and weak nuclear force. It is based on an $SU(3)$ symmetry that introduces the color charge as a quantum number with the values red, green, and blue. The Group theory predicts two multiplets, a singlet and an octet state, of gluons that act as mediators between color-charged partons, i.e. quarks. Since $SU(3)$ is non-abelian, the gluons of the octet state carry color-charge themselves while interacting between partons. This allows the gluons of the octet state to couple to each other. The self-coupling of gluons restricts the range of the strong force to a small scale of few femtometers compared to a infinite range of an photon. Since a gluon of the singlet state would be color-neutral, it would lead to an infinite range of the strong force that would also affect color-neutral fermions. The singlet state is postulated as not existent as such long-range interactions of the strong force are not observed in nature.

From the self-coupling of the gluons, two important properties of the strong force emerge:

- **Confinement:** The confinement states that isolated color charged partons, quarks and gluons, can not exist freely in nature. Within the theory, additional terms for the gluon self-coupling make the potential of the strong force increase with the distance between the quarks. This enforces quarks and gluons to form color-neutral compounds, so-called hadrons. So far no theory has been developed that describes the process of hadronization exactly, but there are some successful approaches to modeling the phenomenon [2].
- **Asymptotic freedom:** With increasing energy scale, the coupling strength of the strong interaction becomes weaker. In the limit of high energies, quarks and gluons, thus, exist freely in nature. At low energy scales, the above-mentioned confinement takes effect. This makes it possible to describe collisions by perturbation theory.

Both effects are of special interest when studying particles produced within deep inelastic scattering, e.g. at hadron colliders. This thesis focuses on proton-proton collisions in which partons carrying a certain fraction of the proton momentum according to the PDF are scattered. The high center-of-mass energy during the scattering process allows the partons, quarks and gluons respectively, to behave asymptotically freely. This enables a *hard interaction* of isolated initial-state quarks and gluons within any particle physics process as illustrated in Figure 2.2. The final-state products of the hard interaction usually carry enough energy to also behave asymptotically free. High energetic final-state gluons and quarks lose their energy while radiating gluons or generating quark-antiquark pairs and generate so-called *parton-showers*. This showering process keeps on going until the energy is low enough that the confinement enforces the bunch of color-charged partons to recombine to color-neutral hadrons within the so-called *hadronization* process. The resulting hadrons are not necessarily stable so that unstable ones decay into more stable particles. As these decays include all kinds of interactions, they result in stable electrons, muons, hadrons, and photons. Each final-state parton, a quark or a gluon, thus, induces a plethora of stable particles, collimated in the direction of its initial parton momentum, called a *jet*. The same holds true for the color-charged proton remnants that also hadronize and form collimated streams of outgoing particles oriented along the collision axis.

Besides, isolated quarks and gluons in the initial and final state tend to radiate soft gluons or quarks, the so-called *Initial State Radiation (ISR)* or *Final State Radiation (FSR)* illustrated in Figure 2.3. This leads to additional hadronic activity besides the hard interaction that may cause additional jets.




Figure 2.2: An illustration of a proton-proton scattering with a hard interaction. Two partons, each of them carries a momentum fraction of the proton according to the Parton Distribution Function (PDF), interact within a hard scattering. The resulting gluons and quarks cause collinear or soft radiation that lead to parton showers, hadronization to derive color-neutral compounds, and the decay of unstable hadrons. Besides, color-charged proton remnants are created that also hadronize. The graph is based on [3].

2.2 Monte-Carlo simulation

The simulation of the hard process and the interactions of the resulting particles with the detector material that allow measuring them is crucial for comparing predictions with measurements in order to validate existing knowledge or derive new insights. Here the theoretical constructs of the Standard Model of particle physics are used in so-called Monte-Carlo (MC) simulations. These simulations are based on repeated, randomly generated events of the interaction processes to obtain a numerical description of the expected measurement results. Usually, particle physics analyses study a dedicated physics process with a set of initial-state and final-state partons, fermions, or gauge bosons, which decreases the number of hard interactions that need to be simulated drastically.




Figure 2.3: An illustration of a ISR and FSR, where a gluon is radiated by a quark in either the initial or the final state.

Most processes can not be calculated analytically using QCD, but approximations are required. Perturbative QCD uses a factorization ansatz to derive quantitative predictions based on the above-mentioned asymptotic freedom. It Taylor-expands each observable, e.g. the cross-section of the hard interaction process, as a function of the strong coupling constant α_s for a certain scale Q . The contribution of the lowest-order is called the Leading-Order (LO), the subsequent one the Next-To-Leading-Order (NLO), and so on. A more detailed insight into the simulation of the hard interaction with perturbative QCD can be found at [4].

Analogously to Figure 2.2, the resulting final-state partons of the simulation are modeled to fragment into a parton shower and hadronize afterward. The parton shower simulation generates the collinear or soft radiation around a final-state parton. For the hadronization process that recombines the color-charged partons to color-neutral hadrons, only semi-empirical models exist. The MC simulations used in this thesis are based on the Lund string fragmentation model [5–7].

To compare MC simulation with the events measured by complex detectors, a detector-simulation is applied on top of the previous steps. It models the interaction of particles with material using a 3-dimensional model of all components of the detector. This allows imitating the impact of the detector on the particles and the response of the detector in the MC simulation.

Particle physics analyses derive different insights from the processes at the various steps of the MC simulation. The final-state products of the hard process are called the *parton-level* in the further course. While the *particle-level* includes the hadronized parton-showers, the *detector-level* additionally applies the detector simulation. This makes it possible to track a single quark or gluon, which first develops into a particle jet, then enters the detector and is finally reconstructed from the signature of the

stream of particles in the detector. Thereby, the influences of the different effects can be studied separately, which is essential for the calibration of jets.

2.3 Z bosons at hadron colliders

This thesis concentrates on studying so-called Z + jets events, where a Z boson is produced within a deep inelastic scattering in association with at least one jet.

The Z boson as the neutral gauge boson of the weak force was discovered in 1983 [8, 9] at the Super Proton Synchrotron (SPS) collider at Conseil européen pour la recherche nucléaire (CERN). Precision measurements done at the subsequent collider Large Electron-Positron Collider (LEP) brought detailed knowledge of the properties of the Z boson. The world averages of its properties at the time of writing this thesis are given by [10].

$$\begin{aligned} \text{Mass } M_Z &= 91.1876 \pm 0.0021 \text{ GeV} \\ \text{Full decay width } \Gamma_Z &= 2.4952 \pm 0.0023 \text{ GeV} \end{aligned} \quad (2.1)$$

The high precision of the measured Z boson properties is ideally suited for detector calibration.




Figure 2.4: Feynman-graph of the production of a Z boson from a quark-antiquark annihilation decaying into a charged lepton pair.

At hadron colliders such as the LHC, the dominant contribution to the production of a Z boson is the Drell-Yan process shown in Figure 2.4. Here, a quark and an anti-quark annihilate and form a Z boson or a virtual photon. Within proton-proton collisions e.g. at the LHC, the anti-quark is a sea-quark of one of the protons. The quark is of the same type as the anti-quark and can either be a valence quark or also a sea-quark from the other proton. The Z boson or the virtual photon can decay into two oppositely-charged leptons. For Drell-Yan processes that have an invariant mass around the Z boson mass, the mass resonance increases the probability of producing Z bosons compared to virtual photons. However, the superposition of the two intermediate gauge bosons prevents a classification of the collision events into those that exchange Z bosons and those that exchange virtual photons.

This thesis focuses on using Z boson for calibrating jets. Hence, it requires additional hadronic activity besides the pure Drell-Yan process. Figure 2.5 (a), (b) and (c) show exemplary processes that contribute at the LO level to the production of a Z boson with one additional jet. They are relying either on two quarks or on one quark and one gluon in the initial state. Besides, Figure 2.5 (d) shows an exemplary NLO process that is based on two initial state gluons. Larger numbers of additional jets and higher levels of QCD calculations, e.g. the NLO level, drastically increase the number of Feynman graphs contributing to the hard process of interest.

The inclusive cross-section [11] for 13 TeV center-of-mass energy proton-proton collision for a Z boson decaying into two muons with additional jets was calculated at the next-to-next-to-leading-order with the tool kit NNPDF [12].

$$\sigma^{\text{theoretical}}(pp \rightarrow ZX) \times B(Z \rightarrow \mu^+ \mu^-) = 1870 \pm 50 \text{ pb} \quad (2.2)$$

Measurements [11] of the CMS detector show a perfect agreement of predicted and measured cross-section within the uncertainties.

$$\sigma^{\text{measured}}(pp \rightarrow ZX) \times B(Z \rightarrow \mu^+ \mu^-) = 1870 \pm 2(\text{stat}) \pm 35(\text{syst}) \pm 51(\text{lumi}) \text{ pb} \quad (2.3)$$

It is one of the largest cross-sections for LHC proton-proton collisions at 13 TeV, which results in a huge number of observed events. The Z + jets process, therefore, suits perfectly for calibrating jets with excellent statistical precision.




Figure 2.5: Examples of Feynman-graphs of $Z + 1$ jet process via different production channels. The quark-gluon production channel either via an intermediate quark (a) or exchange of a quark (b), the quark-antiquark production channel (c), and the gluon production channel (d). Except for the latter, all the channels shown can have a complementary process in which an anti-quark instead of a quark interacts. While (d) represents an NLO process, the others are LO production channels.

The experimental environment


3.1 Accelerating particles at the Large Hadron Collider

Founded in 1954, CERN [13], based in Geneva, Switzerland, is one of the largest international organizations for fundamental scientific research. Originally established for research in the field of nuclear physics, it soon became a pioneer in particle physics research.

Nowadays, it hosts the largest hadron collider, the Large Hadron Collider (LHC), which accelerates a beam of protons or lead ions. The LHC [15] was built in the former LEP tunnel with 26.7 km circumference in a mean depth of 100 m below ground. It is part of a complex for accelerating particles to supply several experiments as illustrated in Figure 3.1. The proton beams this analysis concentrates on, are prepared for acceleration by ionizing hydrogen and primarily accelerated within Linear accelerator 2 (LINAC). Afterward, the Proton Synchrotron Booster (BOOSTER), the Proton Synchrotron (PS), and the SPS step-wise increase the kinetic energy of the protons, before they are injected into the LHC accelerator ring.

The proton beam is not continuous but consists of multiple proton bunches, each one carries a huge number of protons. Thousands of proton bunches are traveling in opposite directions while being forced to cross at four collision points. The LHC is equipped with superconducting dipole magnets that bend the protons to their circular trajectory around the ring. Besides, focusing quadrupole magnets and multipole corrector magnets adjust the beam parameters and focus the proton bunches to prepare them for collisions. Alongside, radiofrequency cavities accelerate the proton

CERN's accelerator complex



European Organization for Nuclear Research | Organisation européenne pour la recherche nucléaire

© CERN 2008

Figure 3.1: Schematic view of the LHC accelerator complex [14]. A sequence of different accelerators prepare and pre-accelerate the hadron bunches before filling them into the actual LHC accelerator ring. It finally reaches up to 6.5 GeV kinetic energy per proton. Alongside, several adjacent experiments are supplied with hadron bundles of different energies.

bunches from initially 450 GeV to 6.5 GeV at the time of performing the presented analysis. Details about the individual components can be found in [15].

There are four major experiments located at the collision point along the ring. Two multi-purpose detectors, the Compact Muon Solenoid [16] and the A Toroidal LHC ApparatuS (ATLAS) [17] detector check for consistency of the Standard Model of particle physics and search for physics beyond the Standard Model. Both successfully verified independently the Higgs mechanism in 2012 [18, 19]. While both detectors almost cover all spatial directions by their cylindrical design, the Large Hadron Collider beauty (LHCb) [20] focuses on the heavy forward region to study b quark physics. A Large Ion Collider Experiment (ALICE) [21], however, is dedicated to heavy-ion collisions that allow to investigate plasma physics.

During the so-called Run 2 period from the beginning of 2015 to the end of 2018, the

LHC performed collisions every 25 ns and achieved center-of-mass energy of 13 TeV. This thesis concentrates on the data taking period by the CMS detector in 2018. In this period, the LHC accumulated about 68 fb^{-1} [22, 23] of integrated luminosity \mathcal{L} . It is measure for the overall performance of the accelerator, which is proportional to the number of events as follows.

$$N_{\text{events}} = \int L dt \cdot \sigma_{\text{process}} = \mathcal{L} \cdot \sigma_{\text{process}} \quad (3.1)$$

This allows estimating the number of observed events N_{events} of a certain hard interaction that is studied using its cross-section σ_{process} .

For the future, there are plans for a High-Luminosity LHC (HL-LHC) [24], where a highly increased integrated luminosity enables, among other measurements, high precision measurements of Standard Model parameters. Here, slight deviations may hint at particle physics processes beyond the Standard Model.

3.2 Detecting particles with the Compact Muon Solenoid

The Compact Muon Solenoid (CMS) is designed as a multi-purpose detector to search for the Higgs boson, to check for consistency of the Standard Model as well as to measure its parameters, look for physics beyond the Standard Model, e.g. supersymmetric particles, and other measurements.

Compared to the competitive ATLAS detector, the CMS detector has a compact design, which is reflected in a smaller size, about 21 m long with a diameter of about 15 m, at a higher weight, 14 000 t [26]. It was emphasized to measure close to the actual point of interaction to reduce external influences. Key features of the CMS detector are the muon chambers that take a major part of the detector and are designed for accurate measurement of muons, and the tracking system that is completely based on silicon. Besides, great importance was attached to installing most subdetectors including the tracking, ECAL, and HCAL within the magnetic coil. The only subdetector outside the magnetic coil is the space-consuming muon system.

Figure 3.2 gives a schematic overview of the CMS detector with its cylindrical-shaped layout consisting of the different subdetector layers. This section gives a brief overview of the subdetector properties that are important for the calibration presented in this thesis. Detailed descriptions of the individual components of the CMS detector can be found at [16, 22, 27]. Together, they enable the measurement of the charge, the




Figure 3.2: Schematic overview, adapted from [25, 26], of the CMS detector shows the different subdetectors: The inner most tracking system, the electromagnetic calorimeter (ECAL), the hadronic calorimeter (HCAL), the magnetic coil, and the muon chambers within the magnetic return yoke. While the layout is almost uniform in radial plane, the detector is split into the central region, the outer end-caps, and the forward components at both ends.

momentum and the energy of the incoming particles, photons, electrons, muons, and hadrons.

Before coming to the subdetectors, the coordinate system as well as the so-called pile-up and underlying event are introduced.

3.2.1 Coordinate system of the detector

The CMS detector is described using a Cartesian coordinate system, whose origin is set to the center of the detector, the nominal interaction point, respectively. While the x-axis horizontally points to the center of the LHC ring, the y-axis vertically faces




Figure 3.3: Visualization of the conventions for the CMS coordinate system taken from [28]. The collision point in the center of the detector defines the origin, whereas the orientation of x-, y-, and z-axis is towards the surface, the center of the LHC ring, and the Jura mountains. Besides the Cartesian coordinate system, the rotational symmetries of the cylindrical detector allow to define a spherical coordinate system based on the angles θ and ϕ and the radius.

the surface, and the z-axis is horizontally aligned along the beam pipe in westward direction.

Exploiting the cylindrical symmetry of the detector, a spherical coordinate system is introduced which is more suitable for analysis purposes. Therefore, the azimuthal angle ϕ is defined in the x-y-plane starting from the x-axis towards the y-axis. In the y-z plane, the polar angle theta describes the angle from the y-axis to the z-axis. Besides, the transverse radius that denotes the distance from the z-axis in the x-y-plane completes the coordinate system.

Both coordinate systems are illustrated by Figure 3.3. Based on these, there is a special coordinate system that is based on the physical properties of the processes to be observed and commonly used in CMS analyses. It uses the transverse radius, the azimuthal angle ϕ , and the rapidity y or the pseudorapidity η , respectively.

The rapidity y of a particle is given by the following equation.

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right) \quad (3.2)$$

The difference between the rapidities of two particles is invariant under a Lorentz boost along the z -axis. At hadron collider experiments, the rapidity is an important observable, where the colliding partons carry different longitudinal momentum fractions of e.g. the proton. Hence, the products of individual interactions have different longitudinal Lorentz-boosts in the lab frame. The above-mentioned feature of the rapidity compensates for the longitudinal boost when comparing the products of the interaction.

Since the energy of the detected particles can not always be measured accurately, the pseudorapidity η is used, which excludes the energy from its calculation.

$$\eta = \frac{1}{2} \ln \left(\frac{|\vec{p}| + p_z}{|\vec{p}| - p_z} \right) = - \ln \left(\tan \left(\frac{\Theta}{2} \right) \right) \quad (3.3)$$

The pseudorapidity is particularly suitable for muons, of which detectors can only accurately measure their momentum. For massless particles with high momentum that hold $m \ll p$ and, thus, $E \approx p$, both, rapidity and pseudorapidity approach each other.

$$\lim_{p \rightarrow E} \eta = y \quad (3.4)$$

Both y and η give the advantage that the particle flux is more evenly distributed as a function of them compared to Θ . This allows a more uniform treatment of different detector regions during detector calibration.

Using the pseudorapidity and the azimuthal angle, the spatial distance ΔR between two objects is calculated.

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} \quad (3.5)$$

This allows searching for certain event topologies within the analysis.

The interacting partons within the proton-proton collisions at the LHC meet almost in a frontal encounter, their transverse momentum \vec{p}_T is negligible. Thus, analyses can exploit conservation of the transverse momentum for the products of the hard process. As already mentioned above, the longitudinal momentum fraction of the partons is not known initially at measurement. Thus, the kinetic observables that are commonly used are p_T , η , ϕ , and the invariant mass M .

3.2.2 Pile-up interactions

At hadron colliders such as the LHC, a major bias on measurement are particles coming from so-called pile-up vertices and underlying event. An overview of the




Figure 3.4: Beside the hard process at the primary vertex, the beam remnants can interact, the so-called underlying event. Additional proton-proton collisions are called in-time pile-up. Effects in the detector material from previous collisions are summarized as out-of-time pile-up. In-time pile-up, underlying event, out-of-time pile-up, and the beam remnants themselves superimpose the hard process distort a proper detection.

contributing processes is shown in Figure 3.4.

Besides the process of interest, the hard process at the *primary vertex*, multiple soft proton-proton interactions occur with the remaining protons in the bunches. These additional collisions are called *pile-up*. According to the timing of the pile-up events relative to the hard process, different kinds of pile-up events are differentiated.

In-time pile-up events are soft collisions that occur within the same bunch crossing as the hard process.

Out-of-time pile-up effects are caused by previous (*early out-of-time pile-up*) or later (*late out-of-time pile-up*) bunch crossings. Due to the cooldown time of detector components such as the calorimeters, energy deposits of previous collisions can remain and are added to the current event. The same effect requires the detector readout to gather information from the detector components during their cooldown time. If the cooldown time is longer than the time between the bunch crossings, parts of the later event can be clustered into the event.

Since only two partons of the composite protons interact within the hard process, the remaining partons can interfere with the particle detection. Each hard process leaves behind two *beam-remnants* that consist of residual partons recombining to

color-neutral hadrons. These beam remnants move mainly in the forward direction along the beam pipe and interfere with detector components located there.

Furthermore, additional interactions between the remaining partons can occur in addition to the hard process. Particles that originate from these so-called *underlying event* at secondary vertices contribute to the measured result.

In 2016, the CMS detector recorded 23 interactions per bunch crossing on average [23]. This number increased in 2017 and 2018 to 32 interactions per bunch crossing. Since most of these interactions are pile-up events, the mitigation of the pile-up influence on the jet energy is a major task of the jet calibration. The pile-up mitigation is addressed in detail within Section 4.3.

3.2.3 The tracking system

The inner-most subdetector of the CMS detector is responsible for measuring the trajectories of charged particles and their momenta via their bending in the magnetic field. Thereby, a high spatial resolution of the tracks is especially important since the interaction vertices need to be reconstructed from the tracks. This is crucial for distinguishing between particles originating from the hard process and particles from pile-up and underlying events.

The CMS detector uses silicon-based sensors [22, 27, 30] that detect charged particles passing the material via the ionization they cause. So-called *particle hits* are detected by creating a charged current within an electric potential and amplifying the signal. Hits within multiple layers allow reconstructing the particle track using a *track-finding algorithm*. Therefore, the CMS experiment uses an adaption of the Kalman filtering technique [31].

Due to the high particle flux that crosses the CMS tracking system, it needs to be radiation hard. Furthermore, it must be lightweight to not disturb the measurement of the following subdetectors. Therefore, the so-called *Pixel Tracker* consisting of multiple layers were placed very close to the beam pipe. They are designed to provide an excellent spatial resolution. The Pixel Tracker is surrounded by a *Strip Tracker*, which is built more cost-effectively because the particle flux is lower as it is further away from the interaction point. Moreover, their spatial resolution does not have to match the granularity of the Pixel Tracker.

Both tracking components are designed to cover the full azimuthal ϕ range and the pseudorapidity range up to $|\eta| \leq 2.4$. Up to $|\eta| \leq 1.3$, the silicon sensor layers are




Figure 3.5: Side view of a cross section through one quadrant of the CMS detector [29]. It visualizes the different layers of subdetectors in size and spacial coverage in η : From inside to outside, there are the tracking system (light blue), the ECAL (green), the HCAL (yellow), and the muon chambers (red). The superconducting solenoid (dark blue) encloses all subdetectors except the muon chambers, which are located in the return yoke (gray).

arranged parallel to the beam pipe, beyond they are placed vertically. There is a small region between both, where the layers overlap.

3.2.4 The electromagnetic calorimeter

The electromagnetic calorimeter (ECAL) is designed to measure the energy of photons and electrons. Therefore, it uses absorber material that enforces incoming particles to create an electromagnetic shower. Within the Coulomb-potential of the detector material, the incoming particle interacts via bremsstrahlung and electron-pair production and create a bunch of multiple electrons and photons. The particles of this electromagnetic shower deposit their energy in the detector volume via the Compton scattering or photoelectric effect. Scintillating detector material forwards the energy deposits to arrays of photodiodes that convert the electromagnetic shower to digital

signals. Thereby, the number of registered photons is proportional to the energy of the incoming particle. Photons and electrons are distinguished since the latter ones are charged and leave hits in the tracking system. When assigning energy deposits to tracks, the energy losses of electrons due to bremsstrahlung are taken into account by the reconstruction algorithm.

The CMS uses lead tungstate (PbWO_4) crystals as both, absorber and scintillator material [32–34]. As this material provides a low scintillation time and a short radiation length, it suits perfectly for designing a compact ECAL. This is important because the design of CMS intends to place the ECAL and also the HCAL inside the solenoid.

The ECAL crystals are arranged in two components as illustrated in Figure 3.5. The central region, the so-called *barrel* region, covers $|\eta| \leq 1.48$ and the *end-cap* region $1.48 < |\eta| \leq 3.0$. In the end-cap region, a *pre-shower* detector is installed in front of the ECAL to increase the spatial resolution of and distinguish between photons of the hard process and photons originating from pion decays. Since the electromagnetic shower may hit several ECAL crystals, a calibration [34] is applied within the ECAL reconstruction algorithm to compensate for effects e.g. caused by alignment, electron drifting, or aging of the crystals.

3.2.5 The hadronic calorimeter

The hadronic calorimeter (HCAL) is constructed for measuring the energy of hadrons [35–37]. Electrons and photons are absorbed within the ECAL and, thus, do not reach the HCAL.

An absorber material, in case of the CMS brass plates, enforce incoming hadrons to interact with the material and create a so-called hadronic shower. During the showering process, any kind of hadronic interaction is possible, which leads to a bunch of hadrons, electrons, and photons. They enter the scintillation material, in case of the CMS detector a plastic scintillator, and ionize it. The resulting scintillation photons are registered by silicon photomultipliers and converted to digital signals. For optimal and compact measurement of the hadron energy, the absorber and scintillator material alternates a few times.

The CMS HCAL consists of a central part that covers $|\eta| \leq 1.3$, end-caps that cover $1.3 < |\eta| \leq 3.0$, and outer parts that extend the HCAL with $3.0 < |\eta| \leq 5.2$. However, the outer HCAL part is located outside the magnetic recoil and is constructed

slightly different with steel as absorber material and quartz-fibers as the scintillator. Thereby, the central region provides the highest granularity. Similar to the ECAL, the reconstruction algorithm of the HCAL includes a calibration.

3.2.6 The superconducting solenoid and the muon system

The tracking system, the ECAL, and the HCAL is surrounded by the superconducting magnetic coil that is made of niob-titanium (NbTi) and provides a magnetic field with 3.8T. It bends the trajectories of charged particles and, thus, allows to determine their charge and measure their momenta. The high energy of the particles entering the detector makes a strong magnetic field necessary to precisely measure the momentum. While the magnetic field inside is almost homogeneous, an *iron return yoke* outside the solenoid returns the magnetic field and ensures an almost homogeneous outer magnetic field.

The *muon chambers* are incorporated between the layers of the return yoke. Since muons produced within the collision are minimum-ionizing, they barely interact with the detector materials and cross the ECAL and HCAL unimpeded. Their lifetime is sufficiently long that the majority leaves the detector without decaying. Thus, their track, respectively their momentum and charge, is measured by the tracking system in the inner and the muon chambers in the outer detector region. The energy of the muons needs to be reconstructed from the momentum measurement assuming a certain mass. While reconstructing the muon tracks, effects such as the energy loss due to bremsstrahlung, multiple scattering within the detector material and inhomogeneities in the magnetic field are considered.

The muon system [38, 39] consists of a central part and a part in the end-caps covering in total $|\eta| \leq 2.4$. There are three types of muon detectors used within the CMS detector. *Drift Tubes* are placed in the barrel region, while *Cathode Strip Chambers* handle the higher particle flux and the non-uniform magnetic field in the end-cap region. Additionally, *Resistive Plate Chambers* are placed in both regions up to $|\eta| \leq 1.6$. However, the complete muon system covers $|\eta| \leq 2.4$.

3.2.7 The Trigger and Data Acquisition System

During LHC Run 2 period, the proton beams cross every 25 ns, each time producing about 60 simultaneous interactions on average between the partons [40]. In order to

reduce the amount of data for reconstruction and to store only events of interest, a two-level triggering system is used at the CMS [41, 42].

The so-called Level-1 Trigger (L1 Trigger) must select interesting events from the 40 MHz Bunch-Crossing-Rate and reduce the rate for further processing to 100 kHz. It triggers further processing when registering e.g. high-energetic particles or unusual particle combinations. The reduced output rate allows the Data Acquisition System (DAQ) system of CMS to read out the detector electronics, synchronize the information of different detector modules, and assemble the events to the High-Level Trigger (HLT). Whereas the Level-1 Trigger (L1 Trigger) is a pure hardware trigger, the High-Level Trigger (HLT) is based on software running on a dedicated computing cluster connected by a high-speed, low-latency network with 200 Gbit s^{-1} . Consisting of compute nodes, it runs the full CMS reconstruction software stack. Consuming about 1 MB per event, the HLT reconstructs the particles and jets and classifies the events by specific signatures, e.g. two muons with certain energy thresholds. The selected events are written to permanent storage with a rate in the order of 100 Hz, a few hundred Megabytes per second, respectively.

Further distribution of the recorded data for backup and analysis purposes is taken over by the Worldwide LHC Computing Grid (WLCG) described in Section 8.1.

Jet measurement and calibration at CMS

At hadron-colliders such as the LHC, a collision of high-energetic hadrons initiates interactions between the underlying initial state partons, the so-called hard process. Interactions of such color-charged partons consisting of quarks and gluons, the so-called strong force is described by QCD as described in Section 2.1.1. The theory points out that the potential of the strong force between two quarks increases with the distance between both. This also applies to the final state partons that are produced in the hard process.

While these quarks fly apart towards the detector, the strong force can spontaneously create color-neutral quark anti-quark pairs out of the vacuum. Thereby, the potential of the strong force has to reach required energy threshold for the production of these two new quarks. Due to confinement, color-charged particles can not exist freely and therefore recombine to color-neutral structures or hadrons, respectively. This process of creating and recombining quarks is repeatedly performed and creates a bunch of hadrons. Since most of the newly generated hadrons are short-lived, they decay into more stable hadrons in flight. This forms a shower of stable hadrons, leptons, and photons that enters the detector.

A transversal boost from the primary parton interaction causes all resulting particles to pass through the detector layers in a locally restricted, cone-like shape. The axis of this so-called jet indicates the direction of the four-momentum of the original parton. Hence, the four-momenta of all particles that contribute to the jet need to be summed up to derive the properties of the original parton. Dedicated algorithms are used to reconstruct jets by clustering the tracks and energy deposits recorded by the detector.

Thereby, various effects bias the measurement of jets, such as:

- detector noise and miscalibration
- non-linear calorimeter response
- particles from additional interaction vertices, pile-up or underlying event, being clustered into the jet
- particles originating from the hard process and not considered by the clustering algorithm, i.e. neutrinos or unregistered hadrons, electrons, muons, or photons

A precise measurement of the properties of the particle level jets requires the reconstructed jets to be corrected for these effects. To project the properties of the particle-level jet to the corresponding parton of the hard process, additional theoretical corrections are required.

This thesis focuses on the first, correcting the reconstructed jets to the particle-level jet. Data-based or simulation-based methods can be used to obtain correction factors that correct both recorded and simulated data. In the MC simulation, especially their simulation of the detector response, the influence of the correction can be validated by comparing the reconstructed jet and the particle-level jet. However, these MC and detector simulations use dedicated models to simulate the hadronization process, the detector response, and the contribution from pile-up or underlying events as described in Section 2.2. Hence, the simulations do not necessarily represent reality in detail.

To compensate for residual differences between simulation and reality caused by the modelling, the properties of the simulated jets need to be corrected to the ones of the measured jets. Data-based approaches, which exploit known processes and symmetries, can be used to calibrate jets on the average.

At CMS, multiple well-known physical processes are used to correct for different influences mentioned-above. The results of the individual analyses, each based on one well-known physical process, are combined in a factorized calibration approach that extracts correction factors. Thereby the jet energy scale and the jet momentum resolution are calibrated separately.

Typical well-known processes suitable for calibration purposes are the dijet and γ +jets channels due to their high number of events produced at the LHC. Additionally, the Z+jets channel, where the very high accuracy of the event reconstruction compensates for a lower production rate, is perfectly suited for the calibration of the jets. It profits from the precisely defined properties of the Z boson measured by previous HEP experiments. Furthermore, the reconstructed Z bosons are perfectly suited for the

calibration of jets since the CMS detector is especially designed to accurately measure the decay products of the Z boson, electrons and muons.

This $Z + \text{jets}$ analysis focuses on extracting absolute jet energy corrections to correct for residual differences between MC simulation and data. Furthermore, it presents a new procedure to extract the jet momentum resolution using $Z + \text{jets}$ events.

4.1 Reconstruction of physical objects by the Particle-Flow algorithm

Multiple layers of the CMS detector that measure different quantities. At CMS, there are three different kinds of reconstructed jets [43], which differ in the input. The input can be the energy deposits that are derived from the hadronic and electromagnetic calorimeter modules of the detector, or the tracking information from the CMS tracking system.

Calorimeter jets (CALO jets) are reconstructed using energy deposits in the calorimeters. Geometrically combining ECAL and HCAL cells allows for clustering energy deposits and reconstructing jets.

Jet-Plus-Track jets (JPT jets) improve the CALO jet precision by including information from the CMS tracking system. All tracks attributed to a CALO jet are utilized to correct the jet energy and momentum and derive the corresponding JPT jet.

Particle-Flow jets (PF jets) are clustered using objects that are previously reconstructed by the Particle-Flow (PF) algorithm [44].

This algorithm exploits the performance of all subsystems of the CMS detector to classify each object and reconstruct its four-momentum. The PF algorithm applies the following recipe to differentiate between particles.

- As **muons** are minimal-ionizing when crossing the detector layers, their tracks are reconstructed combining information of the inner tracking system and the outer muon chambers. Afterward, the assigned tracks and energy deposits are removed from the list of recorded objects that must be processed to avoid double counting.
- **Electrons** are reconstructed by combining tracker information and the corresponding energy deposits in the ECAL. Bremsstrahlung photons radiated by electrons being bent by the magnetic field are also taken into account.

- Isolated **photons** that show entries in the ECAL without an appropriate track and no corresponding energy deposits in the HCAL are identified in the same step.
- **Charged hadrons** are characterized by tracks pointing to energy deposits in both calorimeters.
- **Neutral hadrons** remain when coinciding energy deposits in the calorimeters can not be matched to any track.

Neutrinos can not be measured directly. Due to conservation of the momentum in the transverse plane, the transverse momentum of all contributing neutrinos can be deduced from the missing transverse momentum ($p_{\text{T}}^{\text{miss}}$). $p_{\text{T}}^{\text{miss}}$ is determined by the negative vectorial sum of the four-momenta of all particles of the event.

The PF approach is more precise than a pure calorimetric approach, as it combines the response and resolution of all detector components. This makes PF jets, which are then clustered from all PF candidates by a clustering algorithm, the standard for most CMS analyses. As almost all CMS analyses concentrate on using PF jets, this analysis focuses on the calibration of such jets.

4.2 Jet clustering

Jet algorithms can be mainly divided into two classes. On the one hand, *cone algorithms* associate all particles within a fixed distance in the $\eta - \phi$ plane, respectively associate them to a cone around the jet center. On the other hand, *sequential recombination algorithms* avoid a fixed geometric form but calculate for each pair of candidate objects a variable distance measure that is weighted with the transverse momentum of the particle.

Jet clustering algorithms that are not influenced by splitting up an input particle into two collinear particles are called *collinear safe*. Equally, algorithms are called *infrared safe*, which means, that the result of the clustering algorithm does not change, when adding a soft input particle. Collinear and infrared unsafe examples are illustrated in Figure 4.1. When using collinear or infrared unsafe jet clustering algorithms for experiments, it is difficult to calculate the required predictions within the simulations since perturbative QCD is not applicable. Thus, modern HEP experiments focus on infrared and collinear safe sequential recombination algorithms. Cone algorithms,




Figure 4.1: Example of collinear and infrared unsafe jets. Collinear safety requires jet clustering to be independent of splitting input constituent. In the collinear unsafe case, the algorithm does not reconstruct a jet if the momenta of the split constituents are below the minimum threshold. Infrared safety prevents jet clustering result to change when adding soft object. Illustration inspired by [29].

even if they also can be infrared and collinear safe, are not favored as they are usually slower than concurring sequential recombination algorithms.

At CMS, the anti- k_t algorithm [45] is used as the default jet clustering algorithm, since it is collinear and infrared safe. This sequential recombination algorithm calculates the distance parameter $d_{i,j}$ between two input objects i, j according to Equation 4.1.

$$d_{i,j} = \min(p_{T,i}^{2n}, p_{T,j}^{2n}) \frac{\Delta_{i,j}^2}{R^2} \quad (4.1)$$

The anti- k_t algorithm is defined by $n = -1$. $\Delta_{i,j}$ is the distance between the objects in the y - ϕ plane.

$$\Delta_{i,j} = (y_i - y_j)^2 + (\phi_i - \phi_j)^2 \quad (4.2)$$

Furthermore, the cone size R is chosen adequately for the collision environment and the purpose of the analysis. The default setting for CMS run 2 is $R = 0.4$, which is commonly used by most analyses. Besides, $R = 0.8$ is used for dedicated analyses and cross-check.

The algorithm proceeds with the object pair with the minimal value of $d_{i,j}$. This distance parameter is compared to $d_{i,B}$ that proves the distance between the object entity and the beam

$$d_{i,B} = p_{T,i}^{2n} \quad (4.3)$$

If the distance parameter to the object j , $d_{i,j}$, is smaller than the distance parameter to the beam $d_{i,B}$, the objects i and j are merged. However, if $d_{i,B}$ is lower than $d_{i,j}$, the object i is moved from the list of objects to the list of resulting jets. This

procedure is repeated until the list of objects is empty and all objects are clustered to jets.

Before applying the algorithm, it must be ensured that the list of input objects is cleaned of isolated particles such as electrons, muons, or photons that most likely are created in the hard process. These isolated particles need to be treated separately within the selection process of the analysis.

At CMS, however, the algorithm is applied to the full list of input objects. Thus, the list of resulting jets needs to be cleaned of jets that contain single objects like electrons, muons, or photons later.

4.3 Pile-up mitigation

Pile-up events, described in Section 3.2.2, have a major influence on the jet energy since jet algorithms cluster those pile-up particles into the jets. To avoid distorting the jet properties, analyses need to subtract the pile-up contribution before clustering jets.

At CMS, the Charged Hadron Subtraction (CHS) method [46, 47] is well-established to reduce the influence of in-time pile-up on the reconstruction of particles and jets. The algorithm removes all charged hadrons whose tracks point to pile-up vertices. This reduces the number of pile-up particles clustered into the jets. Residual contribution of pile-up and underlying events to the jet properties need to be corrected by the jet calibration.

Alternatively, the Pileup Per Particle Identification (PUPPI) method [48] aims at removing charged pile-up particles from the event and weighting neutral pile-up particles down. Charged pile-up particles are identified by searching for tracks not originating from the primary vertex. Neutral particles are weighted according to the calculated probability that they originate from pile-up events. This probability is calculated considering the global pile-up density as well as locally adjacent charged particles of the hard process.

Both methods require the reconstruction of the tracks of charged particles to significantly reduce the influence of pile-up particles. Since the CMS tracking system covers $|\eta| \leq 2.5$, the pile-up contribution in the outer spacial region $2.5 < |\eta| \leq 5.191$ can not be mitigated using the CHS or puppi method.

The CMS jet calibration focuses on covering the majority of CMS analyses when providing unified corrections that simplify jet reconstruction for analysts. Therefore, it primarily concentrates on the CMS default settings for analyses, which include the CHS method. The PUPPI method is studied in a second working step.

Hence, this analysis also focuses on the CHS method. However, the procedure for deriving absolute residual corrections described in this thesis is applicable to PUPPI corrected jets, too. This holds also true for the proposed method for extracting jet momentum resolution using $Z + \text{jets}$ events. The proposed procedure puts a handle on the influence of pile-up remaining in CHS corrected jets. Hence, it is also able to handle PUPPI corrected jets that usually come with a reduced pile-up contribution.

4.4 Balancing methods used for jet calibration

Balancing jets against more precisely reconstructable reference objects in the transverse plane plays a major role in estimating the energy scale and the momentum resolution of a jet. Figure 4.2 illustrates the required ideal event topology in the




Figure 4.2: An ideal event topology balances a precisely measured reference object with a jet in the transverse plane. Assuming that the transverse momentum is conserved in the event allows for calibrating the latter using a balancing method.

transverse view.

The balancing approach assumes that the transverse momentum is conserved during the collision. Thus, the transverse momentum of the final state parton that forms the jet is assumed to be equal to the transverse momentum of the reference object.

This allows balancing methods to calibrate the transverse jet energy with reference to the precisely measured reference object. Two methods that are commonly used by various calibration steps: The p_{T} -balance and the MPF method.

4.4.1 p_{T} -balance method

The p_{T} -balance method aims at directly comparing the jet with a well-defined and precisely measurable reference object, such as a photon or the Z boson decaying into two muons or electrons. It is required, that the energy and momentum of the reference object are measured with higher resolution and lower systematical uncertainties than the jet to be calibrated. As equation 4.4 shows, the p_{T} -balance method directly calculates the ratio between the transverse momentum of the reference object and the reconstructed jet.

$$R_{p_{\text{T}}\text{-balance}} = \frac{p_{\text{T}}^{\text{jet}}}{p_{\text{T}}^{\text{ref}}} \quad (4.4)$$

Whereas the p_{T} -balance method just requires the transverse momentum of the reference object to be precisely measured, it is sensitive to ISR and FSR. Radiated objects that are clustered into an additional jet affect the transverse momentum of the jet and, thus, distort the p_{T} -balance.

4.4.2 Missing E_{T} Projection Fraction method

The Missing E_{T} Projection Fraction (MPF) method takes into account the environment of the event to reduce the influence of ISR and FSR. It assumes an ideal event topology, where one reference object is balanced against one jet.

$$\vec{p}_{\text{T}}^{\text{jet1, gen}} + \vec{p}_{\text{T}}^{\text{ref, gen}} = \vec{0} \quad (4.5)$$

Each object is measured with a certain detector response, which transforms the properties of the object at the particle-level into the properties of the reconstructed object.

$$\vec{p}_{\text{T}}^{\text{ref, gen}} \rightarrow \vec{p}_{\text{T}}^{\text{ref}} = R_Z \cdot \vec{p}_{\text{T}}^{\text{ref, gen}} \quad (4.6)$$

$$\vec{p}_{\text{T}}^{\text{jet1, gen}} \rightarrow \vec{p}_{\text{T}}^{\text{jet1}} = R_{\text{jet}} \cdot \vec{p}_{\text{T}}^{\text{jet1, gen}} \quad (4.7)$$

In this case, any mismeasurement of the transverse momentum and the energy of the jet is directly represented by the missing transverse momentum $\vec{p}_{\text{T}}^{\text{miss}}$.

$$R_{\text{jet}} \cdot \vec{p}_{\text{T}}^{\text{jet1, gen}} + R_Z \cdot \vec{p}_{\text{T}}^{\text{ref, gen}} = -\vec{p}_{\text{T}}^{\text{miss}} \quad (4.8)$$

Thereby, the reference object is assumed to be accurately measured so that its response R_Z is set to 1 and $\vec{p}_T^{\text{ref, gen}} = \vec{p}_T^{\text{ref}}$. By solving this system of equations for the jet response, the MPF response R_{MPF} is obtained.

$$R_{\text{MPF}} = R_{\text{jet}} = 1 + \frac{\vec{p}_T^{\text{miss}} \cdot \vec{p}_T^{\text{ref}}}{(p_T^{\text{ref}})^2} \quad (4.9)$$

In comparison to the p_T -balance method, this method is based on \vec{p}_T^{miss} instead of \vec{p}_T^{jet1} . This reduces the sensitivity on ISR and FSR but requires an accurate measurement of \vec{p}_T^{miss} . Therefore, the detector needs to be sufficiently calibrated before deriving residual corrections.

Comparing the results of both methods allows to register detector effects and cross-check the influence of ISR, FSR, and pile-up on the jet measurement.

4.5 Factorized calibration approach for jet momentum scale at CMS

At CMS jets are calibrated in a factorized approach [43, 47, 49–51] as shown in Figure 4.3. Consequently, several correction levels are applied to the raw reconstructed




Figure 4.3: The factorized approach to calibrate the jet energy scale extracts correction factors in multiple stages. Each stage is either applied to MC simulated or recorded data. Illustration inspired from [47].

jet. Each correction level treats a dedicated systematic bias on the transverse momentum of the jet by applying a correction factor. While the first two correction levels are mainly based on MC simulation, the latter ones compare data to MC simulation to correct for residual differences. The multiple stages of jet energy correction levels are derived and applied in the following order.

1. **Pile-up offset corrections:** Comparing simulated response to detector level with the inputs at the generator level allows for subtracting the excess energy

that is not associated with the hard process. This excess energy is mainly coming from pile-up and electronic noise of the detector. Additionally, a random cone method is used to estimate the noise contribution using simulated and recorded data.

2. **Simulated detector response correction:** The detailed MC simulation of the detector allows for adjusting the energy of the reconstructed jets on average to the energy of the generated initial-state parton. This corrects for detector effects and non-linear calorimeter response.
3. **Relative residual correction:** Due to the tracking system, jets within $|\eta^{\text{jet}}| < 1.3$ can be efficiently corrected for pile-up and detector effects. Correlating jets in the forward region to the central ones allows for widening the area covered by the calibration and derive η -dependent correction factors.
4. **Absolute residual correction:** Balancing jets against a precisely measured reference object allows for correcting the absolute jet energy scale as a function of the transverse momentum $p_{\text{T}}^{\text{jet}}$.
5. **Flavor correction (optional):** Differences between jets originating from gluons, light quarks, charm quarks, or bottom quarks are corrected using channels that preferably produce a single so-called jet flavor.

While the stages 1 to 4 are mandatory for all CMS analyses, the flavor corrections can be applied optionally depending on the analysis purpose.

The main focus of the CMS jet energy calibration lies on the calibration of PF jets that are clustered with the anti- k_t algorithm at $R = 0.4$ and CHS applied. Support for other settings, such as using PUPPI instead of CHS or $R = 0.8$ instead of $R = 0.4$, is given after fixing the corrections for the default settings. This is done in an iterative way by successively applying the correction levels and deriving correction factors for the next level.

Within the next sections, the individual stages are shortly presented.

4.5.1 Determination of pile-up offset corrections

The first stage of the jet energy scale calibration corrects for contributions to the jet energy that come from pile-up, underlying events, and detector noise. There are two complementary methods used to estimate the distorting energy offset.

- **Jet area method:** This method determines correction factors on a jet-by-jet basis. It assumes that low p_T objects from pile-up and detector noise are isotropically distributed in the y - ϕ plane. In this case, their contribution to the jet energy can be estimated by multiplying the jet area with the average p_T density of the event. Contributions from underlying events need to be subtracted from the p_T density since they form an intrinsic part of the proton scattering process.
- **Average offset method:** This method extracts correction factors that are on an event basis. It estimates pile-up and detector noise by using a collection of randomly triggered events while vetoing processes of interest, i.e. high activity from a single vertex. Placing random cones in the η - ϕ plane allows for summing up the noise contribution. After averaging the results in ϕ , correction factors that depend on the number of primary vertices n_{PV} are derived. Each event is corrected according to n_{PV} .

Both methods are combined within the **Hybrid jet area method**. A dedicated η^{jet} dependent factor, which is estimated from a modulation of the average offset method, is introduced into the jet area method. This compensates for non-uniformity of the energy response. The derived correction factors depend on the p_T density and n_{PV} of the event and η^{jet} .

4.5.2 Determination of simulated detector response corrections

The MC simulation of the full detector response enables a direct comparison between a reconstructed jet at the detector-level and its corresponding particle jet on the generator level. Therefore, each reconstructed jet in simulated QCD events is spatially matched to an MC particle in the η - ϕ plane. The ratio between the average detector p_T^{jet} with the p_T of the matching generator particle is used as a correction factor. Due to detector layout and response, this method is applied in bins of η^{jet} and p_T^{jet} .

4.5.3 Determination of relative residual corrections

The first two stages of jet energy scale calibration assume that MC simulation describes the detector response precisely enough to derive correction factors. However, this correction stage compensates for residual differences between data and MC simulation.

The central region ($|\eta| \leq 1.3$) of the CMS detector provides the highest uniformity of the detector and the highest accuracy and range to measure the jet energy. Hence, it is chosen as a reference region to calibrate jets in the forward region. Therefore, the p_T balancing method introduced in Section 4.4 is used for dijet events. One jet lying in the central region of the detector $|\eta| \leq 1.3$ is balanced with a jet at arbitrary $|\eta|$. Comparing the p_T -balance extracted from MC simulation and data allows to determine a η dependent correction factor.

4.5.4 Determination of absolute residual corrections

In contrast to the previous calibration stage, the absolute residual corrections [28] derive p_T dependent correction factors. Balancing a jet against a precisely measured reference object allows an increase in the accuracy within the central region. The previously applied relative residual corrections propagate the accurate jet measurement from $|\eta| \leq 1.3$ to the forward region.

At CMS, three types of reference objects are used to extract absolute residual correction factors.

- **Multijet** events extend the measurement range to very high p_T jets
- γ + **jets** events provide reference photons with a high p_T range
- $Z(\rightarrow \mu + \bar{\mu})$ + **jets** and $Z(\rightarrow e + \bar{e})$ + **jets** events use the combined accuracy of the tracker, the muon system and the ECAL and extend the p_T range down to very low values.

Both balancing methods described in Section 4.4 are used: The p_T -balance method, which is also used for deriving relative residual corrections, and the MPF method. Typically, both are sensitive to additional jets, which disturbs the balancing of the jet to the reference object. The large center-of-mass energy of the collision supports ISR and FSR. Furthermore, multi-parton interactions contribute in terms of pile-up and underlying events. Thus, there is typically more than one jet associated with the event. As the balancing methods assume an ideal event topology with only one jet balanced to the reference object, the analysis procedure has to account for those additional jets.

Therefore, the second jet activity α defined in Equation 4.10 is used to estimate the influence of additional jets.

$$\alpha = \frac{p_T^{\text{jet2}}}{p_T^{\text{ref}}} \quad (4.10)$$

Extrapolating to zero second jet activity ($\alpha \rightarrow 0$) allows determining corrections for the ideal event topology that consists of only one jet and the reference object.

The extrapolated results of both methods, p_T -balance and MPF, for all channels are combined within a global fit.

4.5.5 Determination of flavor corrections

This optional correction stage uses a pure MC based calibration approach. MC simulated Z+jets and γ +jets events within $|\eta| \leq 1.3$ allow for separating different jet flavors. u and d quarks are corrected in combination but s, c and b quarks are treated separately. Dedicated correction factors are extracted by comparing the particle-level p_T of the jets with the reconstructed p_T .

There are at the time of writing the thesis efforts to extract flavor dependent corrections within the procedure of deriving absolute residual corrections.

4.6 Estimation of the jet momentum resolution at CMS

Detectors such as CMS measure the properties of particles at a certain finite resolution only. The effects of each component of the subdetectors involved in the measurement contribute to the overall resolution of the measurement.

The measurement of the transverse momentum of leptons or photons is based on a few hits within the CMS tracking system. This subdetector is designed for high precision measurement of particle tracks and, thus, the resolution of the transverse momentum of these particle types is also accurate. Since the jets consist of large numbers of particles that are measured separately in the CMS detector, the transverse momentum is reconstructed from all particles involved. Effects of the tracking system, the ECAL, and the HCAL, which all contribute to the jet measurement, as well as influences of the jet clustering algorithm, affect the measurement of the momentum and impair the resolution of the jet momentum. Therefore, the resolution of the jets is rather poor compared to non-compound objects like single leptons or photons.

These effects also complicate the correct modeling of jet resolution in MC simulations, which is necessary for precise jet analysis. The modeling of the jet momentum resolution suffers from the dimensionality of the effects to be considered and is

therefore simplified in the simulation. Hence, it does not necessarily represent the actual resolution of the CMS detector.

To correct for that, scale factors have to be derived to smear the jet momentum resolution in MC simulation relative to data as defined in Equation 4.11.

$$SF(Data/MC) = \frac{JER_{Data}}{JER_{MC}} \quad (4.11)$$

The jet momentum resolution JER describes the detector resolution for jets. It is defined as the variance of the ratio of the reconstructed jet momentum and the momentum of the corresponding particle-level jet.

$$JER = \sigma \left(\frac{p_T^{\text{jet}}}{p_T^{\text{gen-jet}}} \right) \quad (4.12)$$

At the CMS experiment, two different determination procedures are applied to dijet and γ + jets events. Both procedures balance jets in the transverse plane similar to the methods used for correcting the jet energy scale for estimating the jet momentum resolution. Whereas the analysis of dijet events focuses on a dedicated asymmetry method balancing two jets in the transverse plane, the γ + jets channel exploits the p_T -balance method described above. Detailed descriptions of the established determination of the jet momentum resolution at CMS can be found at [43, 47, 49, 51, 52].

While both procedures are shortly presented in the following, this thesis proposes a new procedure to determine the jet momentum resolution using Z + jets events in Chapter 6. Adding the new Z + jets results allows to extend the p_T range of the scaling factors for the jet momentum resolution to low values. The medium p_T range covered by the Z + jets channel enables a cross-check of the results of the different channels.

4.6.1 Extracting the jet momentum resolution using dijet events

In the case of dijet events, the detector resolution is incorporated in both jets. Assuming that the momentum resolution of both balanced jets is equal, the asymmetry A between both jets is used to estimate the average jet momentum resolution.

$$A = \frac{p_T^{\text{jet1}} - p_T^{\text{jet2}}}{p_T^{\text{jet1}} + p_T^{\text{jet2}}} \quad (4.13)$$

The variance of the jet asymmetry is directly proportional to the jet momentum resolution.

$$JER = \sigma \left(\frac{\sigma(p_T^{\text{jet}})}{p_T^{\text{jet}}} \right) = \sqrt{2} \cdot \sigma_A \quad (4.14)$$

The variance of the asymmetry distribution σ_A is extracted using the Root Mean Square (RMS) truncated on two sides to 98.5% of the integral to compensate for statistical outliers in non-Gaussian tails. These may artificially increase the RMS.

To fulfill the above-mentioned assumptions, dijet events with additional jets are extrapolated to the ideal event topology with only two balanced jets. Therefore, the third jet activity α is defined analogously to the second jet activity described in Section 4.5.4.

$$\alpha = \frac{p_T^{\text{jet3}}}{p_T^{\text{jet2}}} \quad (4.15)$$

When looking at the particle level of the dijet events in MC simulation, a small asymmetry is remaining, which is not originating from the detector resolution. This so-called Particle Level Imbalance (PLI) is caused by underlying events, initial and final state radiation, and the jet clustering algorithm. The pure detector response of the jet momentum is derived when removing PLI from the jet asymmetry for reconstructed jets. This is done by quadratically subtracting their variances as shown in Equation 4.16.

$$JER = \sqrt{2} \cdot \sqrt{\sigma_A^2 - \sigma_{PLI}^2 \quad (-\sigma_{detector}^2)} \quad (4.16)$$

The PLI is thereby defined as the asymmetry A between the particle-level jets instead of the reconstructed jets.

When balancing a jet from the forward detector region with a jet in the central region, an additional asymmetry $\sigma_{detector}$ that compensates for the asymmetric detector design needs to be subtracted. This allows extending the measurement of the jet momentum resolution to the forward region.

4.6.2 Extracting the jet momentum resolution using $\gamma + \text{jets}$ events

In contrast to the dijets channel, the analysis of $\gamma + \text{jets}$ event to determine the jet momentum resolution concentrates on balancing one jet with a reference object, in this case, a photon. This requires a different procedure to extract the jet momentum

resolution. Similar to the absolute residual jet energy corrections, the p_T -balance method is used to compare the jet with the photon. While the jet energy scale is determined by the mean of the p_T -balance distribution, the resolution can be extracted from the spread of the p_T -balance distribution. By splitting up $R_{p_T\text{-balance}}$ into different terms, the detector response of the jet is directly visible in Equation 4.17.

$$\begin{aligned}
 R_{p_T\text{-balance}} &= \frac{p_T^{\text{jet}}}{p_T^{\text{ref}}} \\
 &= \frac{p_T^{\text{jet}}}{p_T^{\text{gen-jet}}} \cdot \frac{p_T^{\text{gen-jet}}}{p_T^{\text{gen-ref}}} \cdot \frac{p_T^{\text{gen-ref}}}{p_T^{\text{ref}}} \\
 &= R_{jet} \cdot R_{PLI} \cdot R_{ref}
 \end{aligned} \tag{4.17}$$

Similar to the dijet channel, there is a PLI contribution originating from underlying events, initial and final state radiation, and the jet clustering algorithm. The influence of this effect is determined from the asymmetry between the reference object and the jet at particle-level. Additionally, the reference object contributes to the overall response of the process.

Each component of the detector contributes to the overall detector response. Since the responses of the detector components that contribute to the measurement of a particle or jet are compatible and act individually, the generalized central limit theorem can be assumed. It states that the sum of random variables that are independent, identically distributed, and have finite variances tend toward a Gaussian distribution if the number of variables grows. Thus, the resulting detector response is Gaussian shaped even if the individual responses of the detector components are not. Since all terms contributing to Equation 4.17 represent a detector response in a certain manner, they are typically assumed to be Gaussian shaped. Furthermore, it is assumed that log-Normal and Gaussian distributions are approximately the same in the area the analysis focuses on. Therefore, Equation 4.18 must hold for all response terms within Equation 4.17.

$$\frac{\sigma(N(x))}{\mu(N(x))} \ll 1 \tag{4.18}$$

Here, each response term is distributed according to $N(x)$ with its variance σ and its mean μ . This allows to reorder Equation 4.17 and calculate the jet momentum resolution according to Equation 4.19. Assuming log-Normal distributions, the variance of these three distributions can be quadratically summed up. Here, it utilizes that the product of log-Normal distributions results in a log-Normal distribution and the resulting variance is given by the quadratic sum of all variances. By reordering the resulting formula, the jet momentum resolution is derived by quadratically subtracting

the contributing variances.

$$\begin{aligned}
 JER &= \sigma(R_{jet}) \\
 &= \sigma\left(\frac{p_T^{jet}}{p_T^{gen-jet}}\right) \\
 &= \sqrt{\sigma\left(\frac{p_T^{jet}}{p_T^{ref}}\right)^2 - \sigma\left(\frac{p_T^{gen-jet}}{p_T^{gen-ref}}\right)^2 - \sigma\left(\frac{p_T^{gen-ref}}{p_T^{ref}}\right)^2} \\
 &= \sqrt{\sigma\left(R_{R_{p_T\text{-balance}}}\right)^2 - \sigma(R_{PLI})^2 - \sigma(R_{ref})^2}
 \end{aligned} \tag{4.19}$$

This method assumes that the resolution of the reference object and PLI is small compared to that of the jet. Typically, the momentum resolution of the reference object is negligible small due to the precise ECAL of the CMS detector. The contribution of PLI extracted from MC simulation is required to be smaller than the jet momentum resolution.

Similar to the dijets extraction procedure, the γ +jets analysis uses the RMS truncated on two sides to 98.5% of the integral to remove the influence of non-Gaussian tails. To account for additional jets that distort the balancing of the jet and the reference object, the resulting jet momentum resolution scale factors are extrapolated to zero second jet activity α defined in Equation 4.10.

Preparations for calibrating jets using $Z + \text{jets}$ events

Modeling the detector response and resolution for a compound object such as jets is complicated by the high dimensionality of effects that need to be considered. The simulation of jets needs to take into account multiple correlated effects, such as the detector response and resolution of all involved detector components, or distortions caused by the clustering of jets, or contributions from underlying events and pile-up. As precise predictions of jets are essential for almost all CMS analyses, a dedicated calibration procedure compensates for such effects and adjusts the energy scale and the momentum resolution of jets in MC simulation to the real conditions.

This $Z + \text{jet}$ analysis focuses on absolute residual jet energy corrections that compensate for residual differences between MC and data within the last step of the CMS calibration procedure. Therefore, scale factors are extracted from $Z + \text{jets}$ events by balancing a jet with a precisely measured Z boson as reference objects. This allows using the balancing methods, p_T -balance and MPF, described in section 4.4 for calibrating a jet. Exploiting the strength of the CMS detector in precisely identifying and reconstructing muons, this analysis concentrates on Z boson reconstructed from two muons. The Z boson decay channel into two electrons can be used to cross-check and to lower the uncertainties by combining the results. For applying the balancing methods, an ideal event topology shown in Figure 5.1 is required. There, a well-measured Z boson is facing the jet that needs to be calibrated. The conservation of the transverse momenta allows for comparing the kinematics of both physical objects and, thus, correcting the jet energy scale.

Besides the jet energy scale calibration, a new method is presented that uses $Z + \text{jets}$ events to adjust the jet momentum resolution in MC simulation to data. The method




Figure 5.1: Balancing a precisely measurable reference object, in this case, a Z boson, with a jet allows calibrating the latter using the p_T -balance or MPF method. To exploit the strength of the CMS detector, this analysis concentrates on Z bosons reconstructed from two muons.

presented in Chapter 7 extends the above-mentioned approach of balancing a jet with a Z boson. This allows to extract the jet momentum resolution in data and MC simulation. Comparing both allows for extracting scale factors, which are used to smear the jet momentum resolution in MC simulation and, thus, adjust it to data.

The scale factors for correcting both, the jet energy resolution and jet momentum resolution, are derived using almost the same data set and selection. Hence, common analysis settings are described in the following sections. Afterward, the individual methods as well as the corresponding results are presented.

5.1 Data samples

This analysis concentrates on extracting the jet energy scale and momentum resolution using $Z + \text{jets}$ events recorded by the CMS detector during the LHC Run 2 period in 2018 at center-of-mass energy $\sqrt{s} = 13 \text{ TeV}$. Both measurements balance a jet with a well-measured reference object, in this case, a Z boson. Z bosons are precisely reconstructed from two oppositely charged muons. The measurement of muons profits from the combination of the inner tracking system and the outer muon system, which

reduces the fake rate and increases reconstruction precision. The balancing of jets with the reconstructed Z boson requires a dedicated event selection, which is described in detail in Section 5.3.

All data sets used within the analysis are listed in the appendix A.1.1.

Both measurements use a recorded data set that passed a dedicated trigger preselecting events with at least two muons. While the event reconstruction was updated for the second half of the data taking period, data recorded in the first half was reprocessed offline to ensure analogous treatment of detector effects by the reconstruction algorithm. A list of the used data samples and their global tag that summarizes the detector conditions can be found in Table A.1.

Corrections for the jet energy scale are derived by comparing data to a LO MC simulation generated with MadGraph [53], while PYTHIA 8 [54] was used as the parton shower generator. To light the phase space more evenly, the LO MadGraph simulation is divided according to the number of jets generated in the hard process and each part is appropriately filled with simulated events. The separately generated MC samples are stitched together considering the different cross-sections and number of events listed in Table A.2.

The extraction of scale factors for the jet momentum resolution is based on an extended NLO MC simulation generated with aMC@NLO [55], an extension of MadGraph for NLO calculations, in combination with PYTHIA 8. In addition to the binning in the number of jets generated within the hard process, the MC simulation listed in Table A.3 is also split into different bins of the transverse momentum of the Z boson. This allows illuminating the phase space more uniformly and accurately examine jet momentum resolution scale factors in a wide p_T and $|\eta|$ range despite the complex analysis procedure.

Using these very slightly preselected events, this analysis applies advanced selection criteria, which are presented in Section 5.3.

5.2 Data processing workflow

The CMS trigger system [56] reconstructs preselected events. Events of interest are identified by both, the online hardware at the detector, the so-called Level-1 trigger [57], and the offline computing cluster, the so-called HLT [58]. The results are

Workflow step	Processing time	Processing frequency	Data reduction
Preselection	2 – 3 days	every 4 months	17 TB → 6.4 TB
Selection	3 – 5 hours	every 3 weeks	6.4 TB → 65 GB
Conversion	45 minutes	every week	65 GB → 12 MB
Plotting	10 – 45 minutes	multiple times/day	12 MB → 2.8 GB

Table 5.1: Overview of the data reduction, the average processing time, and the frequency of reprocessing the analysis steps. All values are estimated for a single run of the jet momentum resolution analysis for the muon channel. This includes the recorded double-muon data set and the aMC@NLO MC sample mentioned in section 5.1. The frequency is roughly estimated and depends highly on the activity within the CMS collaboration.

stored in a CMS specific format within Data analysis tool (ROOT) files, which are processed by the analysis workflow as input.

To shorten turnaround cycles when applying optimizations or changing methods, the analysis workflow is divided into different steps as shown in Figure 5.2. Each step




Figure 5.2: The workflow for calibrating the jet energy scale and momentum resolution is processed in different steps. While the amount of data is reduced during each step, the number of times each step is processed increases.

reduces the amount of data to be processed but at the same time the number of iterations increases due to the adaptation of the analysis. While the first step has to handle a large amount of input data and is rarely processed, the last step is processed very frequently, consuming a relatively small input file. An overview of the amount of data reduction and the processing time per step is listed in Table 5.1.

The division of the analysis into independently processed steps makes it possible to implement analysis sequences, which frequently undergo changes, as late as possible in order to generate results faster in case of code changes. Each step, therefore, deals with a specific part of the analysis workflow as described below.

- **Preselection:** This so-called "skimming" step converts the CMS specific data format into an analysis specific data format while applying a preselection. In

this process, the KARLSRUHE PACKAGE FOR PHYSICS ANALYSES (KAPPA) [59, 60] produces the basic observables for the analysis and reduces the amount of data by roughly a factor of 4.

- **Selection:** Using the output of the previous step, the EXCALIBUR framework [61] applies the full analysis selection, reconstructs analysis specific observables and allows processing different analysis scenarios simultaneously. EXCALIBUR extends the functionality of the ARTUS framework [62, 63] with Z + jets specific analysis tools. Within this step, the amount of data is reduced by a factor of 100.
- **Conversion:** A conversion of event-based ROOT tuples to histograms with LUMBERJACK provided by the KARMA framework [64] allows to transparently use modern optimization techniques such as multi-threading. Reducing the amount of information to a minimum in the same step helps to shorten the processing time of the plotting steps.
- **Plotting:** Within this last step, final results are derived and cross-checks are performed using the tool PALISADE included in the KARMA framework. The small size of the input files allows applying changes in the analysis procedure on-the-fly and get resulting ROOT files and plots within a short time scale.

The pre-selection and selection steps are trivially parallelized by submitting jobs in the order of thousands to an HTCONDOR batch system. At the time of writing this thesis, the later steps are executed locally on individual computing resources. However, there are efforts to also parallelize them for facing an even larger amount of data expected in future LHC runs. While large output files are stored on Grid storage resources, small ones are stored and accessed locally.

The iterative optimization of the jet energy scale factors within the CMS experiment requires the extraction of results as fast as possible. Therefore, the presented workflow is highly optimized to enable fast turnaround cycles while dealing with a huge amount of data.

While selections and corrections are technically split into the different steps, the following sections present a summarized overview including all steps.

5.3 Event selection

The event selection used for the determination of the jet energy scale and the jet momentum resolution is equal in most parts. Both calibration procedures are based on

the above-mentioned balancing approach, which requires a specific event topology. A well-measured reference object, in this case, a Z boson, needs to face a jet. Therefore, kinematic selections as well as a few corrections are applied to recorded and simulated data samples.

5.3.1 Data certification

Alongside the triggering system at CMS that selects events with at least two muons, the quality of data is monitored during the operation of the accelerator and the detector. This allows selecting periods, where the detector and accelerator run properly. Based on a list of validated run periods, all recorded events are selected for further processing.

This analysis uses a list of certified run periods provided by the data certification group of the CMS collaboration. The recorded and certified data recorded in 2018 corresponds to an integrated luminosity of 58.83 fb^{-1} [65].

5.3.2 Pile-up reweighting

At CMS, MC simulations are usually produced simultaneously to data taking. During the production of MC simulations, a certain pile-up distribution is assumed that roughly fits the expected pile-up distribution in recorded data. Afterward, the pile-up distribution in MC samples is re-weighted to match the recorded pile-up distribution.

The average number of pile-up interactions $\langle n_{\text{PU}} \rangle$ is a common measure for re-weighting the MC pile-up distribution. While in MC simulation $\langle n_{\text{PU}} \rangle$ is directly set to a specific value during generation, it has to be measured individually per luminosity section when recording data. Therefore, $\langle n_{\text{PU}} \rangle$ is calculated from the time-dependent luminosity multiplied by the average minimum bias cross-section. The minimum bias cross-section accounts for any interaction during the bunch crossing that was registered by any trigger. Thus, it reflects the rate of hard proton-proton interactions without accounting for underlying-events. For 2018 data taking, the minimum bias cross section was determined to be 69.2 mb .

Summing up the individual $\langle n_{\text{PU}} \rangle$ values of all luminosity sections for the year 2018, a pile-up distribution is derived from data, which is compared to the $\langle n_{\text{PU}} \rangle$ distribution in MC simulation. Thus, dedicated pile-up weights are derived, which are applied to adjust the simulation to data.

5.3.3 Selection criteria for leptons and reconstructed Z boson

The properties of the Z boson are reconstructed from its decay products, in this case, two muons. To guarantee the precise reconstruction of a Z boson, some kinematic and detector dependent selections are applied to the muons and the reconstructed Z boson.

Each of the two muons has to pass the following selection criteria.

- **Number of muons:** For reconstructing muons, at least 2 muons are required. To exclude muons originating from different hard processes and hadron decays, not more than 3 muons are allowed per event.

$$2 \leq N_\mu \leq 3 \quad (5.1)$$

- **Muon transverse momentum:** Each of the two muons has to have a transverse momentum p_T^μ greater than 20 GeV to suppress noise contributions such as electronic noise, pile-up contribution, or underlying events, which mainly occur at low energies.

$$p_T^\mu > 20 \text{ GeV} \quad (5.2)$$

- **Muon pseudorapidity:** Muons are reconstructed by the PF algorithm combining information from the muon system and the tracker. This shrinks the region acceptable for muon detection $|\eta|_\mu$ to the barrel and forward region of CMS.

$$|\eta|_\mu < 2.3 \quad (5.3)$$

The region does not include the full region covered by the corresponding sub-detectors to reduce boundary effects.

- **Muon identification:** The CMS collaboration has established a standardized catalog of criteria that identify a muon, the so-called Muon ID. A detailed description of the Muon ID for CMS run 2 is found at [66–68]. The Muon ID ensures that the muon originates from the primary vertex to exclude muons from pile-up vertices or underlying events. Furthermore, the muon must be detected by several detector components to reduce the fake rate. The specific criteria used within the muon ID are listed at A.1.3. This analysis uses a tight working point to ensure that the measured muons are suitable for the precise reconstruction of the reference object in the balancing approach.

- **Muon track isolation:** Heavy-flavour quark decays that produce a vast amount of secondary particles are a major source for muons. Such decays can produce muons within the hadron shower, which are therefore surrounded by other particle shower products. To reject muons originating from hadron decays and other background processes, so-called isolated muons are selected. Similarly to the Muon ID, the CMS collaboration has established dedicated selection criteria to identify isolated muons [66, 67]. Therefore, all PF particles are taken into account within a cone around the muon track. Using a tight working point ensures that muons originating from heavy-flavor quark decays are excluded.

As events with two or three muons are selected, there are additional restrictions to the muon pair that is taken for reconstructing the Z boson.

- **Charge of the muon pair:** The muons need to have the opposite charge to derive a charge-neutral muon pair suitable for reconstructing the Z boson.
- **Invariant mass of the muon pair:** If more than one muon pair exists within the event, the pair that provides an invariant mass closer to the Z boson mass $M_{Z, \text{PDG}} = 91.1876 \text{ GeV}$ derived from [10] is taken.

Additional selection criteria are applied to the reconstructed Z boson.

- **Z boson mass:** The reconstructed mass of the Z boson M_Z has to be within a mass window of 20 GeV around the Z boson mass $M_{Z, \text{PDG}} = 91.1876 \text{ GeV}$ derived from [10].

$$|M_Z - M_{Z, \text{PDG}}| < 20 \text{ GeV} \quad (5.4)$$

- **Z boson balancing:** To enable a balancing of the reconstructed Z boson to the leading jet, a back-to-back topology is required. Therefore, the analysis selects events that satisfy a certain angular distribution in the ϕ plane of the CMS detector.

$$|\Delta\Phi(Z, \text{leading jet}) - \pi| < 0.34 \quad (5.5)$$

- **Z boson momentum**(only applied for jet energy scale corrections): To reduce contributions of low energetic jets originating from soft parton interactions, a minimum transverse momentum is required for the Z boson.

$$p_{\text{T}}^Z > 30 \text{ GeV} \quad (5.6)$$

This selection is not applied to the analysis of the jet momentum resolution in order to not reduce the number of events too much and to check which minimum threshold can be reached.

5.3.4 Selection criteria for jets

Before applying selections on jets, jets need to be corrected using jet energy scale factors. While the determination of the jet momentum resolution uses fully corrected jets, not all stages of jet energy corrections must be applied for the calibration of the jet energy scale. In this case, the absolute residual corrections that will be extracted within the analysis are not applied.

Afterward, the following selection criteria are applied to the data sets.

- **Number of jets:** To balance and thus calibrate jets, at least one jet per event is required.

$$0 < N_{\text{jets}} \quad (5.7)$$

- **Jet transverse momentum:** Especially in the low p_{T} region, the contribution of jets originating from pile-up increases dramatically. To reject events, where the leading jet corresponds to a pile-up jet, the leading jet must have at least 12 GeV.

$$p_{\text{T}}^{\text{jet1}} > 12 \text{ GeV} \quad (5.8)$$

- **Jet identification:** Similar to muons, the CMS collaboration provides a set of criteria for identifying jets. A detailed description of the jet identification at CMS can be found at [69, 70]. This analysis uses a tight working point with lepton veto for the jet identification. The specific criteria are listed in Appendix A.1.4.

- **Second jet activity:** For balancing the leading jet with the Z boson to calibration purposes, the contribution of additional jets to the event needs to be sufficiently low. Therefore, the analysis is limited to events where the transverse momentum of the second jet is less than 30% of the transverse momentum of the Z boson. The ratio of the transverse momenta of the Z boson and the second jet is called second jet activity α .

$$\alpha = \frac{p_{\text{T}}^{\text{jet2}}}{p_{\text{T}}^{\text{Z}}} < 0.3 \quad (5.9)$$

- **Jet pseudorapidity** (only applied for jet energy scale corrections): Within the calibration procedure of the jet energy scale, the correction factors are

propagated to the forward region by the relative residual corrections. Hence, jets are restricted to the central detector region.

$$|\eta^{\text{jet1}}| < 1.3 \quad (5.10)$$

That allows rejecting pile-up and precisely measuring the jet properties. However, the determination of the jet momentum resolution aims at covering the complete pseudorapidity range provided by the CMS detector. It is therefore not limited to a certain pseudorapidity range.

5.3.5 Missing transverse momentum corrections

Within CMS analyses, the missing transverse momentum ($p_{\text{T}}^{\text{miss}}$) is calculated by the negative sum of all four-momenta of the reconstructed objects. Thereby, objects that have at least 15 GeV to minimizing the contributions coming from pile-up are taken into account.

Analogous to the reconstruction of jets, $p_{\text{T}}^{\text{miss}}$ is highly affected by any distortion coming from pile-up, underlying event, the detector response, or the event reconstruction. Therefore, counterparts of the jet energy scale calibration procedure are applied to the reconstructed $p_{\text{T}}^{\text{miss}}$.

The so-called Type-0 correction reduces the influence of pile-up by applying the CHS method to the reconstructed objects before calculating $p_{\text{T}}^{\text{miss}}$. As described in Section 4.3, this removes charged hadrons originating from pile-up vertices. Thus, the sensitivity of $p_{\text{T}}^{\text{miss}}$ to pile-up is reduced and its resolution is improved.

During the initial calculation of the $p_{\text{T}}^{\text{miss}}$, the jet energy scale corrections that were not considered. Any change of the p_{T} of an object within the event like the jet energy correction requires a recalculation of the $p_{\text{T}}^{\text{miss}}$. Hence, the so-called Type-I modification propagates the jet energy scale corrections that were applied to the jets in the event to $p_{\text{T}}^{\text{miss}}$. This is done by subtracting the changes in p_{T} of the jets from the original $p_{\text{T}}^{\text{miss}}$.

$$\vec{p}_{\text{T}}^{\text{miss,Type-I}} = \vec{p}_{\text{T}}^{\text{miss}} - \sum_{\text{jets}} (\vec{p}_{\text{T}}^{\text{jet,corrected}} - \vec{p}_{\text{T}}^{\text{jet}}) \quad (5.11)$$

By propagating the jet energy scale corrections via the Type-I modification to the $p_{\text{T}}^{\text{miss}}$, it also compensates for non-linear detector response and noise, remaining pile-up influences, and residual differences between data and MC simulation. This improves the resolution and precision of $p_{\text{T}}^{\text{miss}}$ further, which is important when using the observable for calibration purposes.

5.4 Kinematic studies of Z+jet events

Before using the selected and corrected events for the calibration, a look at the kinematic distributions allows a brief validation of the input data. Therefore, the kinematic distributions of the muons, the reconstructed Z boson, the jets, and E_T^{miss} are studied.


For all kinematic distributions shown in this section, the selection criteria that are dedicated to the jet energy scale calibration are applied as described in detail in Section 5.3. Hence, there are also no absolute residual corrections applied to the jets.

For the calibration of jets, the primary focus lies in the reconstruction performance. Thus, this thesis concentrates on comparing the shape of a particular kinematic distribution instead of looking at the absolute values. Trigger or selection efficiencies can be neglected for calibration purposes. To estimate the reconstruction performance more efficiently, all simulated distributions are normalized to data in the following.

The distributions presented in the following use the aMC@NLO MC simulation, on which the determination of the jet momentum resolution described in Chapter 7 is based on. However, the analogous distributions containing the MadGraph MC simulation, which are used for the determining the jet energy scale corrections, can be found in the Appendix A.1.5. They mainly show the same behavior as the aMC@NLO distributions.

5.4.1 Characteristics of muons and Z bosons

The p_T and η distributions of the muons allow checking their suitability for reconstructing a Z boson. In Figure 5.3, the p_T and η distributions for the positively charged muon are exemplarily shown for the data samples used for the jet energy scale correction. A good agreement between data and MC simulation within their uncertainties is observed over a wide p_T range. As Figure 5.3 (a) shows, the steeply falling p_T spectrum caused by the limited phase space is almost perfectly described by the simulation. The geometry of the detector itself has a direct influence on the detection and reconstruction efficiency of muons. While the CMS detector is fairly symmetric in the transverse angular ϕ , it shows a clear dependency on the longitudinal angular that is proportional to η . Figure 5.3 (b) shows the distribution for different bins of the muon pseudorapidity η^μ . Especially the central region of the detector $|\eta^\mu| < 1.3$ is well described by the MC simulation, while the agreement



(a) p_T distribution of positively charged muon (b) η distribution of positively charged muon

Figure 5.3: A shape comparison of the aMC@NLO MC simulation and the 2018 data set for p_T and η of the positively charged muon. MC simulation and data distributions are normalized to each other for shape comparisons.

between data and MC simulation in the forward region is slightly fluctuating. As the presented analysis mainly focuses on the central region of the detector, this proves to be a good basis for a precise reconstruction of the Z boson required for calibrating jets.

Looking directly at the distribution of the reconstructed Z boson mass in Figure 5.4, the mass peak at about 91 GeV is visible. Comparing the mass peak in data and




Figure 5.4: A shape comparison of the aMC@NLO MC simulation and the 2018 data set for the reconstructed mass of the Z boson. MC simulation and data distributions are normalized to each other for shape comparison.

MC simulation, a shift of the mean value is observed. This effect is caused by a

slightly distorted muon energy scale that needs to be corrected. While the method for extracting the jet momentum resolution is independent for such an effect, there is a minimal effect on the determination of jet energy scale corrections. Thus, the shift is corrected within the global fit that combines the Z + jets results with other channels.

In total, the reconstructed Z bosons are suitable for calibration purposes as remaining differences between MC simulation and data either not affect the measurements or will be included in systematic uncertainties later on.

5.4.2 Jet characteristics

The properties of the leading jet balanced to the Z boson is of primary focus for the jet calibration. Within the calibration methods, the second leading jet is just required for estimating the additional jet activity within the event that distorts the balancing.




Figure 5.5: A shape comparison of the aMC@NLO MC simulation and the 2018 data set for p_T and η of the leading jet. The distributions of MC simulation and data are normalized to each other for comparing the shapes.

As Figure 5.5 (a) presents, the p_T distribution of the leading jet shows an excess in data compared to the MC simulation at low p_T values. This may indicate that the contribution of pile-up jets, which concentrate at low p_T values, is slightly underestimated in MC simulation compared to data. Thus, it should not cause any concern at this point. The η^{jet1} distribution shown in Figure 5.5 (b), however, offers a perfect agreement between data and MC simulation within the central detector region. This points out that detector effects within the central region are well understood

5 Preparations for calibrating jets using $Z + \text{jets}$ events

within MC simulation. While the absolute residual jet energy scale corrections focus on the precise central detector region, the jet momentum resolution is extracted for the complete $\eta^{\text{jet}1}$ region. Due to the lack of pile-up mitigation in the very forward detector regions, it is expected that the closure between MC simulation and data there will degrade. This may result in larger systematic uncertainties within this region.




Figure 5.6: A shape comparison of the aMC@NLO MC simulation and the 2018 data set for p_T and η of the second-leading jet. The distributions of MC simulation and data are normalized to each other for comparing the shapes.

The kinematic distributions of the second-leading jet shown in Figure 5.6 are affected by jets originating from pile-up. This pile-up contribution causes an excess at low $p_T^{\text{jet}2}$ values. Beyond the central detector region $\eta > 2.3$, pile-up mitigation gets impossible due to the missing tracking system as described in Section 4.3. Because pile-up contributions can not be well modeled in MC simulation, the differences between data and MC simulation become greater the larger the pile-up contribution is. This also affects the shape comparison done in Figure 5.6.

The contributions of second jets that distort the balancing of the Z boson and the leading jet are treated specially within the calibration methods for energy scale and momentum resolution as described in Section 5.5.2. Thus, deviations between simulated and recorded second jets are compensated within the calibration methods.

5.4.3 E_T^{miss} characteristics

While the p_T -balance method directly balances the Z boson to the leading jet, the MPF method takes into account the whole environment of the event using \vec{p}_T^{miss} . This observable requires the objects of the hard process to be reconstructed accurately and is sensitive to additional objects from pile-up vertices or underlying events. As a short cross-check, the scalar observable E_T^{miss} , which is shown in Figure 5.7 is used. Here, one registers a shift of the E_T^{miss} distribution in data to higher values compared




Figure 5.7: A shape comparison of the aMC@NLO MC simulation and the 2018 data set for E_T^{miss} . The distributions of MC simulation and data are normalized to each other for comparing the shapes.

to MC simulation. A reason for that may be a slightly underestimated contribution of pile-up particles in the MC simulation.

5.5 Study of the jet response using $Z + \text{jet}$ events

Balancing the leading jet and the Z boson is done via two observables, p_T -balance and MPF, as described in Section 4.4. Both observables assume an ideal event topology, where the absence of additional jets allows one to directly compare the momenta of the two objects. Within this section, both observables are compared, and their sensitivity on additional jets spoiling the back-to-back topology is studied.

5.5.1 Comparison of p_T -balance and MPF method

While the p_T -balance method directly compares the momenta of the Z boson and the jet, the MPF method takes into account the complete particle collection of the event. Thus, both methods are sensitive to different distortion effects.

Since both observables represent a balancing topology that allows conclusions to be drawn about the detector response, they also represent a response in a certain way. Typically, such detector responses are assumed to have Gaussian shapes. Therefore, the same behavior is expected from the balancing observables. For an ideal event topology, where the Z boson is balanced against one perfectly calibrated jet without any additional jet, it is presumed that the Gaussian shapes of MC simulation and data are equal.

Assuming that the Z boson is precisely reconstructed, a shift in the Gaussian mean indicates that the properties of the jet do not correctly represent those of the corresponding particle-level jet. This makes the calibration of the jet energy scale necessary. The process is described in the following sections.

Similarly, a difference in the Gaussian width of the simulated and recorded distributions indicates that the jet momentum resolution is not correctly modeled within the MC simulation. Since the momentum resolution highly depends on the time-dependent detector performance, it is difficult to predict in simulation and, thus, need to be adjusted to the corresponding data sample.

Figure 5.8 shows the distributions of both observables, which both confirm approximately Gaussian shapes. The mean is observed to be slightly shifted to lower values in data compared to MC simulation within both observables. This gives a clear hint that the jet energy scale needs to be corrected in data. Furthermore, the variance of the Gaussian shape slightly deviates between MC simulation and data. Especially the ratio between both samples emphasizes the differences within the variances. Hence, the jet momentum resolution in simulation has to be adjusted to represent the detector performance.

Additionally, the p_T -balance distribution shows an excess on the left side of the Gaussian peak and a cut-off at 0.25. The excess is more clearly visible in the data distribution than in MC simulation. A closer look reveals that the excess is caused by mainly two effects. When the leading jet is coming from a pile-up vertex, the assumptions of the balancing approach are not valid anymore. Even if the amount of pile-up jets is dramatically reduced by the CHS method and the event selection,




Figure 5.8: A shape comparison of the aMC@NLO MC simulation and the 2018 data set for p_T -balance and MPF. The distributions of MC simulation and data are normalized to each other for comparing the shapes.

there is a small contribution coming from remaining pile-up jets. Since pile-up jets concentrate at low p_T values, their contribution in p_T -balance is shifted to lower values, which contributes to the above-mentioned excess at low values. The influence was studied by comparing a data sample with a low pile-up contribution to one with a high pile-up contribution. Additionally, the jet clustering algorithm can lead to a jet being split into multiple jets by high-energetic FSR. This reduces the fraction of energy assigned to the separated jets. If a fraction of the leading beam is split off, the p_T -balance value is shifted towards low values, which also contributed to the above-mentioned excess at low values. Enforcing additional jets to have a certain minimum distance in the η - ϕ -plane to the leading jet excludes events with such split leading jets. While this selection criteria reduces the excess dramatically and thus confirms the assertion, it introduces a dramatic reduction of the number of events. Hence, this selection criterion is not used for further analysis.

Furthermore, it is evident that the p_T -balance distribution and thus the excess is truncated at 0.25. This is caused by the selection criteria, in particular the restrictions $p_T^Z > 30 \text{ GeV}$ and $\alpha < 0.3$.

The MPF method does not directly balance the leading jet to the Z boson but assumes that p_T^{miss} reveals the momentum component missing in the jet. An accurately measured MPF observable, thus, requires all detector subcomponents to already be well-calibrated. While the pile-up contributes directly to p_T^{miss} and distorts the MPF method, the random distribution of the pile-up cancels out its influence on average.

The MPF distribution, however, is not that sensitive to contribution coming from jet splitting. When a jet splitting occurs, the missing energy fraction enters the balancing calculation through p_T^{miss} . In total, the MPF method is less sensitive to jet splitting and pile-up contribution compared to the p_T -balance method.

Since both observables are sensitive to different effects, using both within the calibration procedure allows them to cross-check and combine the results.

5.5.2 Study of second jet activity

Additional jets within the event topology distort the required back-to-back topology and, thus, the balancing of the Z boson and the jet. Such jets can arise from ISR, FSR, and pile-up vertices. The p_T -balance method is expected to be sensitive to additional jets due to the direct balancing of the jet and Z boson momentum. In contrast, the MPF method is expected to be less sensitive, since the whole event environment is taken into account.




Figure 5.9: A shape comparison of the aMC@NLO MC simulation and the 2018 data set for the number of jets per event. The distributions of MC simulation and data are normalized to each other for comparing the shapes.

However, the ideal event topology with only one jet and a Z boson rarely occurs due to a large number of expected pile-up vertices during LHC run 2 and the affinity of QCD interactions to cause additional hadronic activity within ISR and FSR. In addition, the probability of ISR and FSR is highly increased due to the large center-of-mass energy of 13 TeV. This leads to a large number of jets registered per event as shown in Figure 5.9. The average number of about 75 additional jets per event is sufficiently described within MC simulation, even if it is difficult to model their exact

distribution. While most of these jets carry low p_T , the influence of additional jets with relatively high p_T compared to the leading jet can not be neglected. Thus, the analysis procedure has to handle the distortions caused by these additional jets.

Both analysis procedures, the extraction corrections for the jet energy scale and the jet momentum resolution, therefore, apply a selection on the second-jet activity α as mentioned in Equation 5.9. This reduces the influence of additional jets on the balancing but does not completely extinguish it. A common approach to completely minimize their influence is to split the analysis results in bins of α and extrapolate to zero second-jet activity.


While this approach eliminates the dependency of the analysis results on second-jets, it may introduce a bias within the pile-up distribution. The applied selection criteria, especially the restriction on second-jet activity α can lead to an asymmetric cut on the pile-up distribution. However, eliminating the second jets within the jet to Z boson balancing topology simultaneously reduces the pile-up contribution within the jets. The derived jet energy corrections and jet momentum scale factors are therefore applicable particularly for analyses whose data sets have a relatively low pile-up contribution. This covers the major amount of all analyses done within the CMS experiment.

5.5.3 Time dependency of the jet response

The environment and, thus, the performance of both, the accelerator LHC and the CMS detector, slightly change during operation. Such changes may occur on a short-term or long-term basis during the fixed data taking period of several months.

Short-term changes such as a drop in beam intensity or fluctuations in the magnetic field are partly taken into account in the reconstruction of the events. Since the jet corrections derived within this analysis are based on statistical methods, they require a large number of observed events and, thus, can not face short-term fluctuations.

Long-term changes such as detector elements damaged by radiation or a slightly different detector alignment caused by the high magnetic fields, can for some degree be covered within the jet calibration. Therefore, the recorded data set is split up into different parts, e.g. by run period A, B, C, and D. Each one provides enough number of recorded events to apply the calibration procedure. A good measure for estimating the suitable granularity of the splitting are the uncertainties. The derived statistical uncertainties should be in the same order as the systematic uncertainties.




(a) Mean of p_T -balance versus p_T of the Z boson (b) Mean of MPF versus p_T of the Z boson

Figure 5.10: Mean of the balancing method versus p_T^Z shown for the MadGraph MC simulation and the 2018 data taking periods.

The studies presented in this section concentrate on the MadGraph MC simulation since they are of interest for determining the jet energy scale. In the case of correction of the jet energy scale, observed time dependence is used to split the scale factors into periods with similar behavior. Since the determination of the jet momentum resolution presented in this thesis is a feasibility study, it uses the complete data set recorded by the CMS detector in 2018.

Comparing the mean of the balancing method, p_T -balance or MPF, versus the p_T of the Z boson for different run periods in 2018 allows detecting long-term changes. Each run period outlines a time period, where the measurement conditions are approximately the same. Figure 5.10 shows a good closure between all run periods within the statistical uncertainties. The differences between the MC simulation and the run periods will be covered by the jet energy scale corrections determined within this thesis. Most differences between the run periods are covered by statistical uncertainties. This agreement will be improved by adding systematic uncertainties, that are expected to be in the same order as the statistical ones. A closer look at figure 5.11, which shows the mean value of the balancing method compared to $|\eta^{\text{jet}1}|$, indicates a small difference between the run period D and the remaining run periods at $2.7 < |\eta^{\text{jet}1}| < 3.0$.

Thus, the jet calibration results officially used within the CMS collaboration are split



(a) Mean of p_T -balance versus $|\eta|$ of the leading jet (b) Mean of MPF versus $|\eta|$ of the leading jet

Figure 5.11: Mean of the balancing method versus $|\eta|^{\text{jet1}}$ shown for the MadGraph MC simulation and the 2018 data taking periods.

up into two time periods, runs A to C and run D. Because the absolute residual jet energy corrections are derived for $|\eta|^{\text{jet1}} < 1.3$, the results shown in this thesis use the full data set.

The jet momentum resolution, however, is extracted using the variance of the p_T -balance method. This analysis applies the extraction procedure to the complete pseudorapidity range provided by CMS. Due to a complex extraction method, the analysis lacks of enough recorded events. Thus, it neglects the minor differences between the run periods by taking the full data set. This may change in future iterations of the proposed method for extracting jet momentum resolution using $Z + \text{jets}$ events.

Determination of absolute residual corrections using $Z + \text{jet}$ events

While previous jet energy corrections concentrate on calibrating recorded and simulated data separately, the absolute residual corrections focus on eliminating remaining differences between MC simulation and data. As described in Section 4.5, the results of various channels are combined within a global fit. The $Z + \text{jets}$ analysis, in which the Z boson decays into muons or electrons, plays an important role because it provides two within the four channels that are combined. While the extraction procedure of the $Z + \text{jets}$ channel is almost the same as in $\gamma + \text{jets}$ channel, multijet events are analyzed in a different manner.

The procedure of extracting absolute residual jet energy corrections using $Z + \text{jets}$ events is described in detail in the following sections. It compares the mean of both balancing methods, the p_T -balance and MPF method in recorded and simulated data. Thereby, the influence of additional jets that distort the balancing approach has to be mitigated. The derived results are validated and cross-checked before combining them with other channels.

6.1 Extrapolation to zero second jet activity

The influence of additional jets to the transverse momentum of the event is estimated as already mentioned in Section 5.5.2 using the second-jet activity α . This ratio between the p_T of the second jet and the one of the Z boson allows extrapolating to zero second jet activity α and, thus, avoiding any interfering additional jet. As




Figure 6.1: Linear extrapolation of the mean values of both balancing methods, p_T -balance and MPF, comparing data and MC simulation. The ratio at $\alpha = 0$ corresponds to the absolute residual correction factor, which corrects for remaining differences between the MC and data distributions. While this plot exemplarily shows inclusive results, the official correction factors are binned in $|\eta^{\text{jet}1}|$ and p_T^Z .

Figure 6.1 shows, the mean of the balancing methods depends linearly on the second-jet activity below $\alpha < 0.3$.

While the MPF method is nearly independent of α , the p_T -balance method is highly sensitive to second jets. Comparing the trends in MC and data distribution, we observe slightly different slopes. If we extrapolate the activity of the second jet to $\alpha = 0$, we find that the p_T -balance and MPF for the data match within their statistical uncertainties. We observe the same for both methods within the MC simulation. This confirms that all interfering effects were covered by the event selection and the α -extrapolation. The remaining difference between data and MC simulation is used to calibrate the jet energy scale. Thus, the data-MC ratio at $\alpha = 0$ directly corresponds to the absolute residual correction factor. While Figure 6.1 exemplarily shows the derivation of an inclusive correction factor, the data sets are split into $|\eta^{\text{jet}1}|$ and p_T^Z bins for the official correction factors. This allows a differentiated mapping of detector performances within the different detector subsystems in the MC simulation.

6.2 Jet energy scale corrections extracted from Z + jet events

While the CMS detector is approximately uniform in ϕ , its cylindrical shape leads to an η dependent detector response. Within the very central detector region below $|\eta^{\text{jet1}}| < 1.3$, the structure is uniform and provides the most accurate measurement results. Thus, we shrink the valid region for the absolute residual corrections to that region and propagate the corrections to the forward detector region via the relative residual corrections. Usually, different detector subsystems perform slightly differently depending on the p_{T} of the registered particle. Thus, the scale factors are derived in bins of the transverse momentum of the jet. The balancing allows using p_{T}^{Z} instead, which can be more accurately measured than the momentum of the jet.

The α extrapolation as described in the previous section is applied to each p_{T}^{Z} bin. Taking the ratio data versus MC simulation extrapolated to $\alpha = 0$, we derive one correction factor per p_{T}^{Z} bin.

6.3 CMS combination of absolute jet energy corrections by global fit

While the above-mentioned procedure is done at the Karlsruhe Institute of Technology (KIT), the combination with other channels is done by the CMS collaboration. As described in Section 4.5.4, this procedure is also applied to Z + jets channel decaying into electrons and the γ + jets channel. The multijet channel was skipped when deriving jet energy corrections for 2018 data samples. All correction factors of the individual channels are combined within a global fit as illustrated in Figure 6.2. This includes the results of both balancing methods, p_{T} -balance and MPF for each channel. The results of both Z + jets channels decaying into muons and electrons are combined separately before including them into the global fit. Within this global fit, the systematic uncertainties are introduced per channel and taken into account in addition to the statistical uncertainties. This includes systematic uncertainties coming from jet momentum resolution, FSR and ISR, pile-up contribution, jet flavor, and luminosity measurement.




Figure 6.2: Combination of p_T^Z dependent absolute residual correction factors for the jet energy scale done via a global fit. The results of the $Z + \text{jets}$ channel decaying into muons and electrons were pre-combined into $Z + \text{jet}$. While the solid black line indicates the extracted correction factors, the yellow band shows the absolute jet energy scale uncertainty. The graph is taken from [51].

6.4 Validation of the absolute residual jet energy corrections

For validation purposes, the correction factors derived from the global fit are applied to the data sets used. Figure 6.3 shows the results for both balancing methods. The closure between data and MC simulation at $\alpha = 0$ gives a clear hint that the derived corrections fulfill their purpose.

A more detailed view is provided by Figure 6.4, which shows the mean of the p_T -balance method versus p_T^Z and $|\eta^{\text{jet}1}|$ after applying the derived correction factors. Comparing the results with the corresponding uncorrected ones within Figure 5.10 and Figure 5.11 confirms that the closure between data and MC simulation was drastically improved. Within most p_T^Z and $|\eta^{\text{jet}1}|$ regions, the corrected MC simulation matches to data within the uncertainties. Differences between the run periods are also reduced by splitting the data sets into two parts, run periods A to C and D, and derive separate correction factors. The overall agreement between recorded and simulated data is dramatically increased. Remaining slight differences may be caused by differences between the channels used as input for the global fit. Adding systematic uncertainties, which are likely to be in the order of the statistical uncertainties, may cover the




Figure 6.3: Linear extrapolation of the mean values of both balancing methods, p_T -balance and MPF, comparing data and MC simulation. The absolute residual corrections derived within the global fit are applied and the closure between MC simulation and data at $\alpha = 0$ validates the results.

remaining differences. There is room for improvement while the presented method for deriving absolute residual corrections is iterated to gradually eliminate distorting influences.

6 Determination of absolute residual corrections using $Z + \text{jet}$ events




Figure 6.4: Mean of the p_T -balancing method versus p_T^Z and $|\eta^{\text{jet}1}|$ shown for the Mad-Graph MC simulation and the 2018 data taking periods. The absolute residual corrections derived within the global fit are applied and MC simulation and data closes dramatically better compared to Figures 5.10 and 5.11. Iterating the analysis procedure and gradually eliminating distorting influences allows the closure to be further improved. Systematic uncertainties, which are likely to be in the order of the statistical uncertainties, may cover the remaining differences.

Extraction of jet momentum resolution from $Z + \text{jet}$ events

While jet energy scale corrections at CMS are derived and cross-checked with independent channels using different data sets, the extraction of jet momentum resolution scale factors in 2018 is based on a single dijet data sample. The missing possibility of a comparison with other channels complicates troubleshooting and reduces the validation possibilities. To overcome this limitation, the following section presents a new approach to estimating the jet momentum resolution using a $Z + \text{jets}$ data set.

The approach is based on former studies done with $\gamma + \text{jets}$ events [52, 71]. While the similar event topology of $\gamma + \text{jets}$ and $Z + \text{jets}$ channels allow sharing the same analysis method, their applicability for CMS run 2 needs to be approved. As described in the following section, we observed effects that hinder applying the extraction procedure without adjusting it. Therefore, a new procedure was developed that handles the distorting effects.

7.1 Proposed extraction procedure for $Z + \text{jets}$ events

The theoretical background of the extraction procedure is already described in Section 4.6.2.

As shown in equation 4.17, the p_T -balance response is divided into three terms respectively responses: The momentum response of the jet, the particle-level imbalance,

and the momentum response of the Z boson. As the variance of the jet momentum response corresponds to the jet momentum resolution, we assume log-Normal distributions and derive Equation 4.19. This means, that we can subtract the variances of the PLI and the Z boson momentum response from the variance of the p_T -balance response to derive the jet momentum resolution.

While the p_T -balance response can be directly measured in data, the three split terms contain particle-level information. They are therefore derived from the MC simulation, assuming that it describes the particle level sufficiently accurately.

Similar to the approach of deriving jet energy corrections using $Z + \text{jets}$ events, we check for any dependency on second jets that distort the balancing approach. Analogously, an extrapolation to zero second jet activity is performed to derive scale factors that correct the momentum resolution in MC simulation to data.

The results of each analysis step are described in detail in the following section. They are based on an extended aMC@NLO MC sample to enrich the number of events in a wide p_T^Z and $|\eta^{\text{jet1}}|$ range. The full jet energy corrections stack is applied to data and MC simulation to ensure already well-calibrated jets. To study the maximal observations range, the selection requirements $p_T^Z < 30 \text{ GeV}$ and $|\eta^{\text{jet1}}| < 1.3$ are not applied.

7.2 Extraction of response variance

As a basic step for calculating the jet momentum resolution, the variance of the response distributions needs to be extracted. We assume log-Normal response distributions to be able to reorder Equation 4.17 and calculate the jet momentum resolution as mentioned above. The variance of the log-Normal distribution can be approximated by the Gaussian variance if Equation 4.18 holds for the distribution $N(x)$ with its variance σ and its mean μ . This means, that the variance of the distribution has to be much smaller than its mean value. We expect the mean to be close to 1.0 and all resolutions to not exceed a third of the mean value. A look at the response distributions listed in Figure 7.1 shows that this requirement is fulfilled in all cases.

When having a closer look, we observe various effects that hinder using the RMS value as a simple estimate of the Gaussian width.




Figure 7.1: Response distributions for the aMC@NLO MC simulation and the data of 2018 used for extracting their variance and calculate the jet momentum resolution. As example, the distributions are shown inclusively in p_T^Z for the bin $0.25 < \alpha < 0.3$, $|\eta^{\text{jet}1}| < 0.522$. We observe non-Gaussian tails that originate from FSR in all shown distributions. Additionally, the p_T -balance and PLI distributions show an excess at the left side of the peak caused by the jet splitting, pile-up contribution, and a cut-off at low values due to the applied selection criteria. Their influence on the variance extraction is reduced by choosing $[\text{Max} - \text{RMS}_{\text{corr}}, \text{Max} + 2 \cdot \text{RMS}_{\text{corr}}]$ as the range for fitting a Gaussian function.

- **Excess on the left side of the response peak:** The p_T -balance method assumes that the balanced jet is correctly reconstructed by the clustering algorithm. Any FSR or pile-up contribution can cause the jet to be split into at least two jets at reconstruction. Since in this case parts of the four-momentum




Figure 7.2: Generated jet momentum response distribution from the aMC@NLO MC simulation used for comparison with the extracted jet momentum resolution. As example, the distributions are shown inclusively in p_T^Z for the bin $0.25 < \alpha < 0.3$, $|\eta^{\text{jet}1}| < 0.522$. We observe non-Gaussian tails that originate from FSR. Their influence on the variance extraction is reduced by choosing $[\text{Max} - \text{RMS}_{\text{corr}}, \text{Max} + 2 \cdot \text{RMS}_{\text{corr}}]$ as the range for fitting a Gaussian function.

of the jet are missing, the p_T -balance shown in Figure 7.1 (a) and (b) and PLI response shown in Figure 7.1 (c) is shifted towards lower values. While PLI as a pure particle-level variable is only affected by jet splitting due to FSR, the p_T -balance distribution is also affected by pile-up contribution. Here, a pile-up jet can also mimic the balancing jet and forces the leading jet to be replaced. This so-called jet switching increases the excess observed in the p_T -balance distribution, especially the one derived from data, compared to the PLI distribution.

- **Cut-off at low response values:** Especially the p_T -balance and the PLI distributions show a cut-off at 0.3. The selection requirements listed in Section 5.3 causes that cut-off by rejecting events with second-jets that carry more than 30% of the p_T^Z . This directly enforces the leading jet to have at least 0.3 times p_T^Z .
- **non-Gaussian tails:** If radiation effects such as ISR or FSR happen and are clustered into the jets during reconstruction, the response is shifted randomly to lower or higher values. The effect is visible as non-Gaussian tails in both directions. This symmetric effect can be seen clearly in the Z boson momentum resolution shown in Figure 7.1 (d), where all other effects are negligible.

These effects distort the measurement of the jet momentum resolution and must therefore be excluded. Since these effects add an asymmetric behavior to the distributions,

a symmetrical truncation to 98% of the peak area as done for the $\gamma + \text{jets}$ channel is not suitable. Thus, we propose a new procedure fitting a Gaussian fit within an asymmetric fit range. In the first step, the fit range is determined by extracting the RMS value for the upper half of the distribution starting at the maximum. This removes the excess and the cut-off at low values. Truncating a distribution affects the RMS value, which thus needs to be corrected.

We, therefore, use the formula for a two-sided truncation of a Normal distribution [72, 73] shown in Equation 7.1.

$$\sigma_{\text{truncated}}^2 = \sigma^2 \left[1 + \frac{A \cdot \phi(A) - B \cdot \phi(B)}{\Phi(B) - \Phi(A)} - \left(\frac{\phi(A) - \phi(B)}{\Phi(B) - \Phi(A)} \right)^2 \right] \quad (7.1)$$

The truncation is illustrated in Figure 7.3. Here, A describes the difference between




Figure 7.3: Illustration of a two-sided truncation of a Gaussian distribution with the mean μ and the variance σ . a marks the lower truncation threshold, b the upper one, respectively.

the lower truncation value and the mean value in terms of the variance σ and B the difference for the upper value, respectively.

$$\begin{aligned} A &= \frac{a - \mu}{\sigma} \\ B &= \frac{b - \mu}{\sigma} \end{aligned} \quad (7.2)$$

Besides, the Gaussian probability density function $\phi(x)$ and the Gaussian cumulative distribution function $\Phi(x)$ is required.

For the case of a half truncated distribution with $a = \mu$ and $b \rightarrow \infty$ the formula is simplified to Equation 7.3.

$$\sigma_{\text{truncated}}^2 = \sigma^2 \left[1 - (2 \cdot \phi(0))^2 \right] \approx 0.6 \quad (7.3)$$

This value is used to correct the RMS value for the truncation effect.

In the next step, a Gaussian function is fitted to the range $[\text{Max} - \text{RMS}_{\text{corr}}, \text{Max} + 2 \cdot \text{RMS}_{\text{corr}}]$. This allows avoiding any influence of the cut-off effect, the excess at low values, and the non-Gaussian-tails. With this procedure, we derive a robust estimation of the variance for each response distribution while we exclude the distorting effects.

Since this procedure is done for a two-dimensional binning separated by α and p_{T}^Z or $|\eta^{\text{jet1}}|$, some bins lack statistics. Here, the procedure for extracting the variance may get unstable. These bins are either rejected from the further analysis or their fit range is manually tuned to get valid results. An overview of the rejected and tuned bins is shown in the appendix A.1.6. The same procedure is applied, when binning 3-dimensional in all three variables α , p_{T}^Z and $|\eta^{\text{jet1}}|$. Here, the finer binning increased the fraction of bins that either need to be removed or require manual tuning.

The derived fit ranges are shown as a red area in Figure 7.1, the fitted Gaussian function in blue. At the Z boson momentum response, which is not sensitive to switched or split jets and the jet selection, we observe an exclusion of the non-Gaussian tails. This holds also true for the generated jet response. Especially within the p_{T} -balance and PLI distributions, the Gaussian fit excludes the excess and the cut-off at low values. Thus, the proposed method is suitable for extracting the variance of the response distributions without taking into account the mentioned distortions.

7.3 Resolution estimation using quadratic subtraction


As described in Section 7.1, the jet momentum resolution is calculated according to Equation 4.19. Therefore, we need to quadratically subtract the Z boson momentum resolution and the variance of the PLI from the variance of the p_{T} -balance observable.

We derive separated jet momentum resolution values for data and MC simulation by using the corresponding variance of p_{T} -balance distribution, while taking PLI variance and the Z boson momentum resolution purely from the simulation. To ensure that the assumptions of this method hold true, the results are validated within an MC closure test described in 7.5.

7.4 Extrapolation to zero second jet activity

Within the next step, the dependency on second jets is studied for the individual variance values taken into the quadratic subtraction as well as the derived jet momentum resolution.

We expect the p_T -balance and PLI resolution to be sensitive to second jets since both compare the leading jet to the Z boson either on reconstruction or particle level. Any second jet that distorts the balancing approach directly affects the measured quantity. In the case of p_T -balance, this may be caused by contributions to the leading jet coming from ISR, FSR, or pile-up. The PLI is, however, just affected by ISR and FSR contribution, since pile-up jets are not part of the particle-level simulation. Because the Z boson momentum resolution focuses on comparing the particle-level and reconstructed Z boson, it is not sensitive to additional jets at all.



(a) Dependency of the resolution values extracted from MC simulation (b) Dependency of the resolution values extracted from data

Figure 7.4: As an example, the dependency of the extracted variances that go into the input response distributions versus the second jet activity is shown inclusively in p_T^Z for the bin $|\eta^{\text{jet1}}| < 0.522$.

An exemplary bin that represents the dependency on second jets of most p_T^Z and $|\eta^{\text{jet1}}|$ bins is shown in Figure 7.4 (a) for MC simulation and Figure 7.4 (b) for data. Both, the variance of p_T -balance and the PLI confirm a strong sensitivity on second jets. Furthermore, the Z boson momentum resolution is independent of α as expected. The extracted jet momentum resolution, however, is slightly depending on the second-jet activity. While the quadratic subtraction of PLI from the p_T -balance variance may cancel out the second-jet dependency originating from contributions of FSR and ISR, the dependency on pile-up will remain.

The remaining dependency on second jets makes extrapolation to zero second jets necessary. We focus on deriving scale factors for correcting the simulation to data and not the individual jet momentum resolution of data and MC simulation themselves. Thus the ratio of jet momentum resolution extracted from data and MC simulation is directly extrapolated to zero second-jet activity. Here, we require at least four of the five α bins to be correctly extracted and available for the linear fit used for the extrapolation. This ensures that the degrees of freedom are sufficiently covered in the linear fit and the influence of fluctuations is reduced.

7.5 Validation with MC closure test

Within MC simulation, the jet momentum resolution extracted by the quadratic subtraction, called extracted jet momentum resolution in the following, can be compared to the true value. Therefore, we determine the generated jet momentum resolution using the variance extraction procedure described in Section 7.2.




Figure 7.5: As an example, the dependency of the generated and extracted jet momentum resolution on second-jet activity is shown inclusively in p_T^Z for the bin $|\eta^{\text{jet}1}| < 0.522$.

As Figure 7.5 shows, the generated and extracted resolution of the jet momentum indicate a slightly different dependence on second jets. This may be influenced by pile-up contributions within the leading jet as described in the previous section. For the extraction method to work, we expect both to agree on $\alpha = 0$ within their statistical uncertainties.

By directly extrapolating the ratio of generated and extracted jet momentum resolution to zero second-jet activity, we derive an MC scale factor. For the method




Figure 7.6: The closure of the MC scale factor with 1.0 validates the method. We, therefore, study its dependency on $|\eta^{\text{jet}1}|$.

to be valid, this factor must agree with 1.0. Figure 7.6 shows the MC scale factor versus $|\eta^{\text{jet}1}|$. Except for two conspicuous $|\eta^{\text{jet}1}|$ bins, the MC scale factor agrees with 1.0 within the statistical uncertainties propagated through the fits. The two bins at $1.93 < |\eta^{\text{jet}1}| \leq 2.043$ and $2.964 < |\eta^{\text{jet}1}| \leq 3.139$ show a non-closure of about 3σ . This may be caused by defective regions in the hadronic calorimeter caused by radiation damage. Such effects are difficult to calibrate and may distort the measurement of the jet momentum resolution at these two conspicuous bins. In general, we observe a good MC closure for the majority of $|\eta^{\text{jet}1}|$ bins. This measurement is, thus, able to extract the valid jet momentum scale factors.

Some detector regions, especially in the transition regions between the components may be sensitive to p_{T}^Z . Therefore, we analogously study the MC scale factor in a 2-dimensional map simultaneously binned in $|\eta^{\text{jet}1}|$ and p_{T}^Z . Figure 7.7 shows the pull of the MC scale factor from the expected value of 1.0. As defined in Equation 7.4, it measures whether the deviations to the expected value are covered by the uncertainties.

$$\text{Pull}(x) = \frac{|x - 1.0|}{\sqrt{\sigma_x^2}} \quad (7.4)$$

The individual values for the MC scale factors and their uncertainties are given in the appendix within Figure A.8 and Figure A.9.

The 2-dimensional map shows, that especially the high and low p_{T}^Z regions suffer from too few events per bin to extract valid results. Besides, the forward region beyond $|\eta^{\text{jet}1}| > 2.8$ is almost not populated with valid results. In both cases, either the Gaussian fit or the linear α extrapolation that requires four of five bins being available is not able to converge. The central region of the CMS detector is, however,




Figure 7.7: Full 2-dimensional MC closure test simultaneously binned in $|\eta^{\text{jet}1}|$ and p_T^Z allows searching for detector regions that provide a p_T dependent jet momentum resolution. Therefore, the pull of the MC scale factor from 1.0 is calculated, which shows whether the deviations are within the statistical uncertainties.

well-populated with events and provides a p_T^Z range from 60 GeV up to 500 GeV. We observe a majority of the 2-dimensional bins to agree within 1σ uncertainties and, therefore, validate the extraction method. Even the fact that a few bins exceed the 2σ uncertainties, the overall agreement of $\chi^2/ndf \approx 1.3$ is a clear hint for the validity of the procedure.

7.6 MC correction by jet momentum resolution scale factors

For the $|\eta^{\text{jet}1}|$ and p_T^Z region validated in the previous section, the final scale factors for the jet momentum resolution comparing data and MC simulation are derived. Similar to the MC scale factor, the ratio of the jet momentum resolution extracted from data and simulation is directly extrapolated to zero second-jet activity. Thereby, the trends can be slightly different in MC simulation and data as Figure 7.8 shows for an exemplary bin. Within the MC simulation, a certain pile-up distribution that does not necessarily perfectly match the one in data is assumed. Simulating




Figure 7.8: As an example, the dependency of the jet momentum resolution extracted from data and MC simulation on second-jet activity is shown inclusively in p_T^Z for the bin $|\eta^{\text{jet1}}| < 0.522$.




Figure 7.9: The dependency of the jet momentum resolution scale factor on $|\eta^{\text{jet1}}|$.

pile-up is quite tricky since the number of pile-up events and the processes that occur within the events are difficult to predict. The small pile-up contribution within p_T -balance, which is slightly different in data and MC simulation, changes the slope of the second-jet dependency. While the extrapolation compensates for such effects, we observe a small difference between data and MC simulation at $\alpha = 0$. The ratio between the extrapolated values in data and simulation corresponds to the scale factor used for correcting the jet momentum resolution in MC simulation to fit the one in data.

We study the dependency of these scale factors versus $|\eta^{\text{jet1}}|$. Figure 7.9 (a) shows a clear dependency on the pseudorapidity. The central region below $|\eta^{\text{jet1}}| < 1.75$, which includes the barrel region and the first half of the forward region, must be

corrected uniformly by about 15%. However, the forward region at $1.75 \leq |\eta^{\text{jet}1}| < 2.8$ indicates that the correction increases by about 50%. The very forward region does not provide such uniform corrections and shows highly fluctuating results with large uncertainties.




Figure 7.10: Illustration of 2-dimensional jet momentum resolution scale factors simultaneously binned in $|\eta^{\text{jet}1}|$ and p_{T}^{Z} . This allows searching for detector regions that provide a p_{T} dependent jet momentum resolution.

As already mentioned in the previous section, we check for detector regions that require p_{T} dependent scale factors. Therefore, a 2-dimensional map of the required corrections factors is provided in Figure 7.10. The corresponding maps showing the uncertainties and the pull of the 2-dimensional scale factors can be found in the appendix within Figure A.10 and Figure A.11.

We observe that the scale factors are uniform within the covered p_{T}^{Z} and $|\eta^{\text{jet}1}|$ range. The amount of events recorded by the CMS detector further shrinks the region validated within the MC closure test. Especially the large p_{T} region suffers from too few events being filled into the bins to derive results from the proposed extraction method. Here, p_{T} dependent results up to $p_{\text{T}}^{\text{Z}} < 500$ GeV are only extracted in the very central region of the detector. In the very forward region, the number of events prevents studying the p_{T} dependency.

Looking at the central region, we observe no dependency of the jet momentum resolution scale factor on p_{T}^{Z} . Thus, the full 2-dimensional binning is not required and




Figure 7.11: Comparison of the jet momentum resolution scale factors derived from $Z + \text{jets}$ events as presented in this thesis and the official results from the dijets channel. Both are binned in $|\eta^{\text{jet}1}|$. The jet momentum resolution scale factors derived from $Z + \text{jets}$ events are based on the full 2018 data set. For comparison, the official jet momentum resolution scale factors that are based on the first half of the CMS data set of 2018 (Run A,B, and C) are used. A comparison with the second half of the data set (Run D) can be found in the appendix in Figure A.12. However, the official scale factors of both parts of the data set show similar trends. Differences are covered by the uncertainties.

the purely $|\eta^{\text{jet}1}|$ dependent scale factors are suitable for correcting the jet momentum resolution in MC simulation to data.

Within Figure 7.11, these $|\eta^{\text{jet}1}|$ dependent jet momentum resolution scale factors derived from $Z + \text{jets}$ events are compared to the dijet results published by the CMS jet energy scale and resolution group. While the jet momentum resolution scale factors derived from $Z + \text{jets}$ events use the full CMS data set recorded in 2018, the official jet momentum resolution scale factors are based on the first half of the CMS data set of 2018. A comparison with the second half of the data set can be found in the appendix in Figure A.12. However, no significant difference in the trends of both parts of the data recorded by the CMS detector in 2018 are observed. Within the central detector region up to $|\eta^{\text{jet}1}| < 1.75$, both official dijet scale factors and the $Z + \text{jets}$ results of the proposed extraction procedure close perfectly within the uncertainties. The same holds true for the very forward detector regions. There, the resulting scale factors of both channels agree within relatively large uncertainties. The region $1.75 \leq |\eta^{\text{jet}1}| < 2.5$, however, shows a constant offset between both channels that roughly reaches 2σ deviation. This may be caused by so-called hot or cold zones, where some parts of the detector did not work optimally and introduce excesses or lacks in the reconstructed jet energy. Such effects are difficult to predict in MC simulation and hard to calibrate within the jet energy corrections. The different

topology of both channels and the resulting different treatment of such effects within the analyses may explain the observed differences. While the shown $Z + \text{jets}$ results contain statistic uncertainties only, the dijet results also take into account systematic uncertainties originating mainly from the FSR estimation in simulation and the jet energy scale.

This thesis is a feasibility study, which clarifies that $Z + \text{jets}$ events can be used for deriving jet momentum scale factors. Therefore, it focuses on extracting validated results and reducing their statistical uncertainties within a complex extraction process that requires a large number of events as input. Future studies may concentrate on newly released data sets for the 2018 data taking period, where the event reconstruction was adjusted to handle damaged detector elements more efficiently. Within these studies, systematic uncertainties for the proposed extraction method may be estimated to combine both channels and derive more accurate jet momentum scale factors.

High-throughput computing in high energy physics

Particle physics research is challenging in several ways. On the one hand, engineers and physicists are tasked with the required complex acceleration and detection methods. On the other hand, theoretical and experimental physicists are faced with elaborate measurements and their complicated interpretations. Another major challenge is computing, which has to provide the infrastructure for handling the enormous amounts of data and the capacities for processing the necessary calculations.

Faced with unprecedented demands for computing capacities and storage in such a scope, the LHC collaborations jointly developed a huge distributed computing infrastructure dedicated to HEP. This so-called WLCG [74] connects worldwide collaboration partners to allow the LHC experiments to distribute data and computing tasks.

The ever-increasing amount of data generated by HEP experiments requires constant efforts to improve and expand the computing infrastructure. This concerns both the storage capacities for storing the vast amounts of data as well as the computing resources for the reconstruction of the recorded events, simulation tasks, and the analysis of the data.

While the LHC is heading for high-luminosity runs to enable precision studies of the standard model of particle physics, computing demands increase drastically. An estimate of the required storage and computing capacity during current and future LHC runs is illustrated in Figure 8.1 [75, 77]. In preparation for the dramatic increase of the resources needed for the HL-LHC from 2026 onwards, the HEP community




Figure 8.1: Estimated demand for storage and computing resources during LHC run 2 and the HL-LHC run [75] required by the CMS experiment. The required CPU hours are given in terms of HS06, a HEP specific performance measure [76]. The amount of required resources will increase dramatically in 2026 while preparing for HL-LHC. Although computing and storage technologies are constantly improving over time, this development does not cover such a strong increase in demand.

at the LHC develops several solutions [75, 78]. While one aspect is working on enhancing code performance to reduce the number of required CPU resources, others focus on making additional resources available for HEP usage. This can either be public or private clouds or HPC systems. Most of these systems do not provide managed storage that allows permanent placement of data for processing. In this case, caching techniques can enable processing of data-intensive workflows by automatic and temporal storing required data on small-sized cache volumes. In addition to its advantages in this case, caching technologies also enable automatic distribution of data. This allows reducing the volume of data sets and their managed replication on permanent data stores provided by the WLCG.

The work presented in this thesis concentrates on the advantages of caching with regard to data-intensive HEP end-user analyses such as the jet energy calibration. The calibration of the energy scale and the momentum resolution of jets presented in the previous chapters requires large data sets to be processed iteratively. Reducing the processing time by caching input data close to the processing CPU, enables short turnaround cycles. This is crucial since the results of multiple analysis groups need to be optimized and combined. Processing the individual analysis tasks faster also means converging to calibration results faster and making them available for the physical analyses that depend on them. Thus, for efficiently processing the presented physics analysis, this thesis also focuses on boosting data-intensive end-user analysis

workflows with distributed coordinated caching. This includes that the applicability and the benefits of placing caches in a distributed computing environment are studied.

The following sections give an overview of the computing environment the analysis is performed in, as well as the data locality concept used for optimizing the data throughput.

8.1 The Worldwide LHC Computing Grid

The WLCG is designed to provide the computing infrastructure to run large HEP experiments connected to the LHC [74, 79, 80]. Therefore, it enables handling vast amounts of recorded and simulated data as well as the corresponding computing capacities to reconstruct, simulate, and analyze the underlying physical processes. It is designed as a globally distributed computing infrastructure organized in a tiered hierarchy, as shown in Figure 8.2. CERN, which houses the LHC and its detectors,




Figure 8.2: The WLCG is organized in a tiered structure. While the Tier 0 is responsible for collecting and storing raw data, the Tier 1 computing centers focus on providing computing capacities for event reconstruction and Grid storage for distribution of data to the end-users. The Tier 2 computing centers provide storage and CPU capacities for MC simulation and analyses of end-users. Outside the strict WLCG structure, the main part for end-user analysis processing is provided by institutes, summarized in the so-called Tier 3 layer.

operates the Tier 0 computer center to collect and store raw data. To ensure uninterrupted operation, parts of the data center are located in Budapest, Hungary. In addition to storing raw data, the Tier 0 computing center provides the infrastructure for the first reconstruction of the recorded events and distribution of the data sets to the next layer, the Tier 1 computing centers. These are distributed all over the world provided and funded by the member states of the LHC collaborations. 13 Tier 1 computing centers contribute to the WLCG. They all hold a shared custodial copy of

the raw data, while the original raw data is stored at the Tier 0. The Tier 1 data center is also designed to reconstruct events from raw data and to process organized analyses of the experiment collaborations, in particular large-scale reprocessing of data sets. They are also responsible for the distribution of the data sets to Tier 2 computing centers. These Tier 2 computing centers provide the computing resources for MC simulation and for end-user analyses. Because both the Tier 1 and the Tier 2 layers are essential for preprocessing and providing data sets, only a fraction of the compute resources is available for end-user analysis workflows. End-user analyses are, thus, mainly developed and processed on computing resources operated and provided by the institutes themselves or national analysis facilities, e.g. the National Analysis Facility (NAF) at Deutsches Elektronen-Synchrotron (DESY) [81]. Such processing and storing capacities that are not directly part of the WLCG infrastructure are summarized as Tier 3 resources. Since they are financed locally, access to such resources is usually limited to the members of the respective institute. However, the combined contribution of the Tier 3 layer to the total computing capacity is not to be underestimated [82, 83]. While the WLCG shares a concept and infrastructure optimized for HEP usage, the Tier 3 resources are individually designed and may not have a common structure.

8.2 Dynamic Tier 3 landscape

Although the production of huge MC simulation campaigns and the distribution of enormous amounts of data is handled by the LHC collaborations via the WLCG, a large part of the data processing is still performed at the Tier 3 level [79, 82–84]. Here, end-users, e.g. the scientists working on physics analyses, process a huge amount of input data to derive their results. Thereby, each contributing institute usually funds, designs, and operates its own infrastructure. This results in a very heterogeneous computing environment using different hardware and software.

Facing the increasing amount of data that will have to be processed in future HEP experiments such as the HL-LHC, there are approaches to unify parts of the Tier 3 computing landscape at least within regional collaborations [85, 86]. This includes collaborative work on analysis tools, provisioning of non-HEP resources for end-user analysis, and optimization of data-intensive workflows. A set of supported tools allows to share expertise, simplify maintenance and application and facilitate setting up new computing resources.

A large fraction of institutes operating HEP computing clusters use the batch system HTCONDOR [87–89] to manage their computing capacities. Large-scale grid storage

systems at the Tier 1 and Tier 2 WLCG data centers focus on providing data over Wide Area Networks (WANs) using the XROOTD [90, 91] and Storage Resource Manager (SRM) [92] data transfer protocols. The Tier 3 storage systems can either offer the same data access options or be limited to local availability within the computing clusters.

Because Tier 3 resources provided by the institutes themselves are limited due to funding and limited manpower, they can be extended with additional computing resources not dedicated to HEP, e.g. provided by the universities themselves for shared use. Such so-called opportunistic resources can be shared computing infrastructures like HPC clusters and commercial or private cloud resources. This is evolving into an increasingly distributed infrastructure that combines heterogeneous resources into one pool of resources managed by a batch system. Because cloud or HPC resources usually do not provide permanent, managed storage space, the processing of data-intensive workflows relies on the access of remote Grid storage resources. However, such resources can temporarily provide local storage volumes. As cloud and HPC infrastructures are usually shared among different users, their fluctuating performance of shared storage, computing, and network resources presents a challenge for optimization.

At KIT, several opportunistic resources were successfully made available for processing HEP analysis workflows. Examples are the HPC center `bwForCluster` for Elementary Particle Physics, Neuroscience, and Microsystems Engineering (NEMO) [83, 93] and the OpenTelekom Cloud (OTC) [94] in the scope of the Helix Nebula Science Cloud (HNSciCloud) project [95, 96].

WLCG computing resources are optimized for processing data-intensive workflows that access data stored within the same computing center. Tier 3 resources in institutes, however, may experience bottlenecks in accessing data at Grid storage elements. This gets especially significant when workflows are processed within a distributed infrastructure, e.g. including HPC clusters and cloud resources, and need to access data remotely since it is not available locally. In this case numerous jobs share a limited bandwidth to remote storage servers and congest network and storage resources. Bringing data processing and storage resources close together avoids bottlenecks and optimizes data processing, which gets increasingly important facing vast amounts of data to analyze. Close, in this case, means that the bandwidth between the storage and the processing unit is optimized for serving the data throughput required for the job. The closeness depends on the possibilities offered by the surrounding infrastructure.

In the scope of this thesis data-intensive calibration workflows are iteratively processed, which can easily be negatively influenced by inefficient data access. By caching the

input data on the computing resources, data locality and, hence, faster turn-around times are achieved.

8.3 Processing end-user analyses

In HEP, each end-user analysis is usually organized in a sequence of *workflows*, each performing a task specific to the analysis. The major part of HEP end-user workflows focuses on statistical data analyses of event-based information. Typically, all events are independent of each other, which allows trivial parallelization of the analysis by splitting up the input data set into subsets. Each subset of the complete data set is processed separately within a so-called *job*. A job holds an instance of the workflow to process a particular subset of data.

Jobs are distributed to *worker-nodes*, hosts that provide computing capacities for batch-processing. Each job is assigned to a *slot*, a fraction of the computing capacity provided by the worker node where it is processed, e.g. fraction of the provided number of CPU cores or part of the memory. Thereby, a batch system manages the computing resources and schedules jobs to suitable worker-nodes.

Although the processing of the jobs is distributed, all participating worker-nodes share a single namespace for accessing and storing data. This is introduced by the WLCG concept of data storage and is typically also applied to Tier 3 storage resources.

By examining factors that limit the processing speed of end-user analysis jobs, it is possible to classify the jobs into two main types. CPU-intensive jobs perform complex calculations on a relatively small input data set. In this case, the processing time is limited by the performance of the CPU cores. On the other hand, jobs that use little computational overhead on a relatively large input data set are limited by the data throughput. While CPU limitations can be overcome by adding additional computing resources, reducing data throughput limitations requires optimizing the complex interplay between network, storage, and computing resources. Data localization concepts, as described in the next section, must adapt the interplay between the components to optimize the data throughput.

8.4 Data throughput optimization via data locality

The efficient utilization of computing resources is challenging when processing data-intensive workflows. In this context, the efficiency of such jobs depends on the

data throughput. Data throughput itself highly depends on the performance of the connections to the servers that provide data storage [94]. The available resources such as storage and network are both limited and shared among multiple users. Limitations arise from network congestion due to bottlenecks in shared WAN bandwidth or massive parallel access to storage resources, especially when a large number of jobs access data over the same shared connections. This can cause the workflows not to receive the data rate required for efficient processing.

Bringing processing and storage resources as close as possible together can avoid such limitations. This so-called *data locality* can be applied on various scales to circumvent the above-mentioned bottlenecks, i.e. globally within a computing community, regionally within a computing cluster, or locally on a per-host basis. This allows to individually tune each computing infrastructure for data-intensive workflows by adjusting the level of data locality to the surrounding conditions.

The techniques to achieve data locality vary from submitting jobs directly to storage servers for processing to pre-placement of data on worker nodes either manually or automatically via caching approaches. The performance of these solutions highly depends on the surrounding computing infrastructure. Examples for existing solutions within HEP are presented in the following.

8.4.1 Data locality within HEP computing

The WLCG considers data locality on the scale of computing centers. Each computing center is responsible for optimizing their internal infrastructure by balancing the number of CPU cores, the Local Area Network (LAN) bandwidth, and the performance of the Grid Storage elements to each other. Since all Tier 1 computing centers hold a shared copy of raw and simulated data and Tier 2 store replicated parts, production jobs of the LHC experiments are scheduled to the computing center that provides the required input data. The replication and management of the data in the various data centers is done automatically supervised by a data management team of the collaboration. It considers the maximal volume of data that can be stored per computing center, the geographical distribution of the data centers, and policies for data save-keeping. Additional replicas required for processing of the jobs in a certain computing center usually require manual interventions. This concept is well suited for the WLCG because data distribution and production workflows are centrally managed by the experiments. In this case, the workflows are scheduled to the computing center that holds a replica of the required input data sets. Usually, the output data sets produced in the jobs are also passed on to the same data distribution infrastructure.

End-user analysis workflows typically start with processing recorded and simulated data sets that are provided by WLCG data centers. However, intermediate results or final output files are not managed centrally by LHC experiments, but locally by the end-users themselves. The overhead of centrally managing such a large amount of short-lived intermediate results or output files outweighs its benefits. Besides, since Tier 3 storage systems are non-uniform and not centrally managed, a centralized approach to data management for Tier 3 resources is not applicable. However, there are efforts to accelerate data-intensive end-user analysis workflows by optimizing data access within local computing clusters. In Chapter 9, a highly scalable concept for boosting end-user analysis workflows within the Tier 3 infrastructure is proposed.

8.4.2 Existing data locality approaches for end-user analyses

Most efforts for increasing the data throughput for data-intensive workflows are done in the WLCG context.

This includes that storage is brought close to the processing hosts via distributed storage systems such as HADOOP [97]. Such a system requires high-performance storage on processing worker nodes that is made available to the distributed file system and dedicated programming of the workflow. Because the HEP infrastructure and software tools have been optimized for large-scale data processing over a long period of time, it is hard to achieve significant improvements by switching to new systems. In the case of HADOOP, no significant improvement was achieved while data formats and software tools had to be adapted to benefit from the data localization approach [98–100]. However, it enables data locality on the scale of computing clusters and individual hosts and circumvents bottlenecks accessing remote storage systems.

A different approach is to process data-intensive workflows directly on the storage servers. Although the Batch on EOS Extra Resources (BEER) [101] project was not designed to achieve data locality, it allows jobs to use a fraction of the CPU cores of storage servers for data analysis processing. Since this project does not pay special attention to data locality, it mainly profits from additional computing resources for job processing. However, it might be beneficial to coordinate jobs to the storage node that provides input data. For scaling up such a system, the storage nodes have to provide sufficient processing power. Nevertheless, it may suffer from insufficient resources when facing the vast amounts of data produced by future HEP experiments. In addition, any influence of the job processing on the performance of the Grid storage elements needs to be avoided, which restricts the approach to a certain fraction of CPU cores. The scalability of the approach is currently limited by the number of Grid storage servers available within the individual WLCG computing centers.

A common approach for bringing storage and data close together suitable for HEP workflows is the *caching* of input data. On first access, a copy of the data is stored on high-performance cache volumes close to the processing CPU. Iterative processing of the same data profits from improved data access. At the same time, it reduces the load on the network when accessing the locally cached copy instead of data provided by remote storage servers. Reducing the load on the network might also increase the over-all performance of the whole system. Thus, the caching fits well for HEP workflows that repeatedly process large amounts of data as input.

The XROOTD framework contains the software for providing storage servers for the XROOTD data transfer protocol, one of the standard data transfer protocols in HEP. Since the XROOTD framework already provides basic caching functionality, current efforts of the LHC experiment communities are focused on its application and improvement. For this, XROOTD caching proxy servers [102] can be placed between the client and remote storage servers to cache accessed data on-the-fly or redirect accesses to already cached data. First applications of single large-scaled XROOTD caches within WLCG computing clusters show excellent performance and scalability of the caching proxy servers [103–105]. Load on network and storage resources was reduced and the data throughput improved.

Accelerating end-user analyses with coordinated distributed caching

Caching of input data is a common approach to realize data locality and thus improve the data throughput for data-intensive workflows. It fits especially well for optimizing HEP workflows that repeatedly process the same input data by reducing their processing time. Access to cached data is expected to improve latency and increase read speed compared to transferring data from remote storage servers. It can also reduce the network load when accessing remote storage servers so that other jobs can also benefit from the freed resources.

Conventional caches deployed in a distributed computing infrastructure such as Tier 3 resources, however, might not necessarily be used efficiently. If repeatedly processed jobs are accessing different caches, data is redundantly stored in multiple caches. Besides wasting limited cache space, duplication can lead to a decreased the cache hit rate, since data already cached on other cache volumes is ignored. Hence, efficient utilization of limited cache space requires a complex management of data duplication that avoids the unnecessary replication of data. However, the more the distributed infrastructure is horizontally scaled, the more individual caches are required and the more duplication occurs. Coordinating data distribution into the caches and scheduling jobs to already cached data is therefore of special interest for optimizing the overall data throughput.

This thesis proposes a concept for distributed coordinated caching, which is based on former research done at KIT [106–111]. These studies were dealing with the optimization of job and data coordination when using multiple caches. While the latter focused on local file caches, this thesis exploits the advantages of XROOTD

caching that enables caching in a distributed infrastructure. This approach suits perfectly for Tier 3 use cases, where the distributed infrastructure requires the management of multiple caches.

Former researches suggest that the granularity of data locality should be best based on individual worker-nodes. Depending on the hardware of the worker-nodes, the surrounding infrastructure, constraints by the resource provider, and other boundary conditions, this is not always feasible. This thesis assumes, that this fine granularity based on individual hosts is also not always necessary. It may be sufficient to achieve data locality on different scales [111, 112] such as on scale of a computing cluster or a region.

Challenges for efficiently utilizing caches within a distributed heterogeneous computing infrastructure are discussed in the following sections. Furthermore, the benefits of distributed coordinated caching within a Tier 3 environment are studied using a prototype setup that was developed in the scope of this thesis and deployed at the KIT.

9.1 Distributed coordinated caching

The distributed coordinated caching concept aims at optimizing the utilization of caches distributed across multiple computing resources within a computing infrastructure. Two main challenges need to be overcome as part of this concept. First, data that is most relevant for caching in terms of increasing the overall data throughput needs to be selected. Second, data-intensive jobs have to be scheduled to the most suitable computing resource that provides already cached input data. Although both challenges are nested, the presented concept treats them separately to simplify optimization and enable scalability.

The concept assumes that the caches are distributed within the computing infrastructure that is optimized. These caches are adapted to the computing resources, a set of worker-nodes, they serve. This allows us to treat a single cache with the computing resources it serves as an independent subsystem that can be optimized separated from others. In order to simplify the optimization of the data throughput, the concept neglects correlations that might occur between the individual caches in this work. Within each subsystem, the data and access management of the cache needs to be adjusted to the cache characteristics, the surrounding computing infrastructure, and the worker-nodes it serves.

Furthermore, the concept assumes that all worker-nodes within the computing infrastructure are managed by a batch system. It is responsible for matching jobs to suitable worker-nodes, preparing the software environment for the jobs, and monitoring their status.

The concept of distributed coordinated caching aims to introduce data-locality into the job scheduling process, taking into account whether the data required for the job is already cached. In this context, the scheduling of jobs is improved by including data locality meta-data into the matchmaking process. Including information e.g. about the distribution of data in the caches allows sending jobs to the computing resource, which is most suitable in terms of data locality.

The batch system, even if no input data of the jobs is cached, indirectly takes influence on the data placement within the caches by scheduling jobs to worker-nodes. Processing of jobs on a subset of the available resources triggers caching of their input data there depending on the used caching algorithm. Thus, the batch system is utilized for the coordinated placement of data in the distributed caches.

This allows using the batch system to coordinate both, data to caches and jobs to suitable resources. As both are optimized at the same time, the challenges described in the following sections need to be faced.

9.1.1 Data selection and placement

In HEP a batch system is usually confronted with a mixture of workflows to be processed. However, the distributed coordinated caching concept focuses on boosting only a fraction of these workflows. The computing infrastructure already satisfies the requirements of CPU-intensive workflows such as simulations, which require practically no bandwidth. In contrast, data-intensive workflows may suffer from congestion and insufficient performance of network and storage resources and may benefit from access to cached data. Studies showed that many workflows processing data are also not optimized enough to benefit from caching at all. Workflows that access few input data, that require only a low data transfer rate, or that are not sufficiently optimized do not benefit from caching. Caching input data of such workflows consumes limited cache storage and, thus, reduces the cache hit rate for other workflows. To increase caching efficiency, data from such workflows must be excluded from caching.

Focusing on the small number of workflows that will profit from caching reduces the storage volume required for caching. The cache hit rate is improved by reading cached

data more often as the cache keeps the required data for a longer period of time and deletes less cached to free up space for new files to be cached. Increasing the efficiency of caches allows even small-sized caches to achieve acceptable cache hit rates and thus be suited for this caching concept. This enables boosting data-intensive jobs also on commercial cloud resources and HPC systems, which provide relatively small or costly temporary storage resources to profit from caching.

While workflows selected for caching directly profit from increased data throughput provided by the caches, the residual jobs indirectly benefit from reduced congestion on network and storage resources. Hence, coordinated caching boosts the overall data throughput by caching a small fraction of input data.

One challenge within this caching concept is identifying suitable workflows and selecting their input data for caching. The coordination of data selection and placement can either be done by each cache making decisions on its own or by a central decision logic that coordinates all caches within the infrastructure. Because the coordination of jobs by the batch system already influences the placement of data in the caches, an additional central cache management is not explicitly necessary. On the contrary, it may interfere negatively with the job management. Hence, this thesis focuses on including cache meta-data into the matchmaking of the batch system and improving the local decision handling of the individual caches. This approach is based on asynchronously fetching of cache meta-data. It avoids synchronizing cache meta-data among all caches as it is required for a centralized cache management. This improves the scalability of the concept. Most of the data selection and placement is thus left to the already existing batch system or an extension of it. The prototype systems presented in this thesis, which were developed for a caching concept for Tier 3 resources, evaluate user data to identify suitable workflows or data suitable for caching, respectively.

Future efforts may concentrate on finding automatisms for the identification of such workflows. An important step that allows us to study the influence of different approaches and test different identification algorithms is to simulate the behavior of caches and the batch system. A tool called LAPIS, with which the behavior of cache and batch systems can be simulated, is being developed at KIT [113–115].

9.1.2 Coordination of jobs to distributed caches

Beyond selecting workflows or data suitable for caching, it is also necessary to coordinate workflows to computing resources that already provide cached data. There

are various metrics to rate the suitability of caches to the job. Numbers that can be considered are the number and size of files that are already cached, the fraction of input files that are already cached, or the boost of the processing time that is expected. Building an algorithm from such measurable or predictable input variables allows us to rank computing resources according to their suitability for the job in terms of data locality. The exact algorithm depends on the meta-data provided by the jobs, the surrounding computing infrastructure, the characteristics of the cache itself, or the political and financial aims within the HEP group.

Here, the simulation of the coordinated distributed caching as mentioned in the previous section may help to develop advanced coordination algorithms adjusted to the local circumstances in the future [113].

9.1.3 Optimization of overall data throughput

Beside coordinating jobs to the most suitable worker node in terms of data locality, the optimization potential of the caches needs to be considered. When placing a cache in a distributed computing environment, it needs to be adjusted to the surrounding infrastructure and worker-nodes it serves. As already mentioned in the beginning of this chapter, achieving data locality on the granularity of individual worker-nodes is not always feasible or necessary. The granularity of data locality needs to reflect the boundary conditions such as the surrounding infrastructure, the hardware of the worker-node, the resource provider, and the political or financial aims of the collaboration. Thus, the distributed coordinated caching approach aims at balancing the performance limits of caches and worker-nodes to optimize the overall throughput of the computing resource.

Therefore, the *level of data locality* that specifies how close data is brought to the processing CPU needs to be chosen. High-performance caches placed within single worker-nodes increase the I/O rate avoiding data access over the network. Proxy caches placed within computing clusters enable boosting data throughput for a set of worker-nodes by exploiting the high-performance local interconnection between the hosts. Such caching proxy servers can also be placed on larger scales, supplying e.g. computing resources of a whole region or a complete HEP experiment community.

The improvement of the overall throughput of the workflows has to be estimated for each of the mentioned cases. In this context, caches reduce concurrent data access on remote storage systems and, thus, reduce congestion on network infrastructure and storage systems. As caches serve as additional sources for accessing data, they do

not necessarily have to be faster than the direct data transfer from remote storage servers.

The improvement gain of caching can be estimated by comparing the limitations of the network infrastructure, the storage resources, the cache, the processing worker-node, and the workflow itself.

$$t_{\text{processing}} = \max \left[\frac{(1 - x_{\text{cached}}) \cdot V_{\text{data}}^{\text{total}}}{r_{\text{network}}^{\text{total}}}, \frac{x_{\text{cached}} \cdot V_{\text{data}}^{\text{total}}}{r_{\text{cache}} \cdot n_{\text{nodes}}}, \frac{V_{\text{data}}^{\text{total}}}{r_{\text{workflow}} \cdot n_{\text{slots}}^{\text{total}}} \right] \quad (9.1)$$

Equation 9.1 gives an estimate of the processing time for a complete workflow accessing a total volume of data $V_{\text{data}}^{\text{total}}$ on $n_{\text{slots}}^{\text{total}}$ slots in parallel. The influence of caching is displayed by the fraction of data that is served from the caches x_{cached} . Access to this part of the data is limited by the total shared bandwidth provided by the cache r_{cache} per worker-node. The total performance of caching relies on how many worker-nodes are supplied by caches n_{nodes} . The remaining fraction of data that is not cached is transferred from remote storage servers via a shared network connection with the bandwidth $r_{\text{network}}^{\text{total}}$. If the type or implementation of the analysis workflow itself limits data processing, it results in a threshold for the data processing bandwidth r_{workflow} . For benchmarks only targeting the cache performance, this limitation can be neglected.

However, achieving maximum data throughput performance is not always necessary or desired. Because the major part of HEP workflows use the scientific software toolkit ROOT [116] that is designed for big-data processing to access data, their data throughput is limited by the deserialization and extraction of ROOT files. Very few analyses would benefit from a data throughput beyond that limit.

Nevertheless, the estimated processing time in Equation 9.1 has its minimum at a certain fraction of cached data. At this point, the available bandwidths of caches and remote storage systems are both saturated to achieve maximum data throughput. Concluding, a certain fraction of data needs to be excluded from caching to achieve maximal performance. This task must be done by the data selection and job coordination discussed in the previous sections.

At shared computing resources such as HPC centers or private cloud resources, one may want to ease the pressure on network resources in order not to affect other users. Furthermore, network traffic as well as storage resources at commercial cloud resources may be cost-intensive especially. Here, the cheapest working point may be preferred.

9.2 Realization of distributed coordinated caching

In the context of this thesis, the elaborated distributed coordinated caching concept was realized within a prototype setup [112, 117–119]. Based on prior studies and experiences [106–111] it is designed to boost data-intensive analysis workflows within a distributed Tier 3 computing infrastructure. To guarantee general usability the prototype is constructed from software components that are commonly used in HEP: The batch system HTCONDOR [87–89] provides a highly scalable way to schedule jobs to compute resources. As described in section 8.2, the XROOTD [90, 91] framework establishes as the standard for handling storage and cache access in HEP. While the XROOTD data transfer protocol allows streaming block-based data efficiently, the XROOTD framework also provides basic caching functionality. The prototype setup combines both systems to coordinate data in distributed caches and match jobs to the most suitable computing resource in terms of the locality of input data. This allows for reducing replication of data within multiple caches and thus improve the efficiency of caching in a distributed computing infrastructure. Furthermore, it enables boosting data-intensive workflows that process the same data repeatedly by scheduling them to computing resource that already provide cached input data.

Although, the prototype setup is bound to the software components XROOTD and HTCONDOR commonly used in HEP, the derived experiences, and the concept can be transferred to different setups. Within the appendix A.2.1, a general coordination service design that is independent of the involved computing infrastructure and its software components is proposed.

9.2.1 Prerequisites of the implementation

A lot of HEP computing sites use HTCONDOR to schedule jobs to a set of computing resources. Those resources are usually arranged as computing clusters, which reach from a few to hundreds of worker-nodes. HTCONDOR is proved to scale up to a million jobs processed per day [120]. The scaling can in addition be improved by connecting multiple HTCONDOR pools and flocking jobs from one to the other HTCONDOR pool [89]. Hence, one of the main strengths of HTCONDOR is the efficient and highly scalable matching of jobs to computing resources. For coordinating jobs to already cached input data, the required data locality information needs to be included into the job scheduling of HTCONDOR. HTCONDOR provides so-called *hooks*, a feature to manipulate each submitted job and, thus, influence the decision of matching the job to a computing resource. The implementation of the data and job coordination

into HTCONDOR is based on three different kinds of such hooks. Those are called at various phases during the job scheduling process. A *translate hook* is called for each job directly after submission. This hook is used to integrate information that is required to rate computing resources according to their suitability in terms of data locality. While waiting for a matching resource to become available, the *update hook* is periodically called and updates the meta-data added to the job continuously. Each finished job invokes a *finalize hook*, which allows us to collect information about job and cache performance. For efficient handling of meta-data and evaluation of data locality, the prototype uses the HTCONDOR hooks as an interface to the batch system while the major part of the data locality ranking is done by a dedicated service. While the implementation of the prototype is bound to such hooks and thus HTCONDOR, the simultaneous coordination of data and jobs, in general, only requires a batch system that provides an interface for manipulating the job scheduling process.




Figure 9.1: XROOTD management server allows the connection of multiple data servers within a hierarchical structure. This enables the creation of a global namespace for all data stored on the data servers. Analogously, XROOTD allows designing a highly scalable caching infrastructure that combines multiple caching proxy servers. Each management server is capable of handling up to 64 data servers, caching proxy servers, or further management servers. Data accesses and metadata requests to any management server are redirected to its subordinate data servers or caching proxy servers.

The basic caching functionality of the XROOTD data transfer protocol and the associated software toolkit, which are optimized for large-scale data processing in HEP, enables the construction of a highly scalable caching infrastructure. Data management services such as *data servers*, *management servers*, and *proxy servers* are designed for creating a global namespace for data. While XROOTD data servers enable access to connected storage volumes, management servers connect multiple data servers within a hierarchical structure as illustrated in Figure 9.1. This allows for merging all storage space provided by data servers into one global namespace. Clients

requesting access are redirected to the data server that serves the corresponding file. The data transfer can also be conducted through an intermediate XROOTD *proxy server*. By connecting such a proxy server to a storage volume, it functions as a *caching proxy server* that caches accessed files on-the-fly. Depending on the individual settings, the caching proxy server either stores a complete copy of a file or caches the blocks of a file that are actually read. When repeatedly requesting the same file, the cached copy of the file or blocks is served. Since multiple caching proxy servers can also be connected to via management servers, a complex caching infrastructure can be created. Caching proxy servers can easily be placed at any kind of computing infrastructure as they only require some temporary storage space for caching files and a network connection. This allows flexible placement of caches within any computing infrastructure. By adjusting the number of caches and their caching behavior, the data throughput within any computing infrastructure can be optimized. Furthermore, flexible provisioning of XROOTD caches allows us to handle a dynamic infrastructure, where computing resources are added temporarily.

Distributed coordinated caching requires to influence the selection and placement of data within the caches as well as to adjust the scheduling of jobs to computing resources. As HTCONDOR allows to modify the scheduling of jobs to computing resources, it fits well for matching jobs to the most suitable computing resource in terms of data locality. The flexible provisioning of a caching infrastructure and adjustable cache decision logic makes XROOTD a perfect choice for realizing the concept. While the hook mechanism and the exchangeable scheduling logic is implemented in HTCONDOR for a long time, the XROOTD client version is required to be higher than 4.7. This allows to load plugins, e.g. for enabling caching functionality or for automatically redirecting XROOTD clients to a pre-configured proxy server transparently for the user. For older client versions this needs to be manually adjusted.

9.2.2 Conceptual design of the coordination service

Distributed coordinated caching requires jobs and data to be coordinated simultaneously and thus meta-data of both systems, the HTCONDOR batch system, and the XROOTD caching infrastructure, to be combined. Matching jobs to the most suitable computing resource in terms of data locality or cached data respectively requires information as input. This includes the content and status of caches, the properties of the jobs themselves, and can in addition include information on the surrounding computing infrastructure, e.g. network bandwidth. The distributed coordinated caching concept was realized by implementing a coordination service NAVIX [121] that intermediates between the caching and the batch system.




Figure 9.2: The NAVIX coordination service connects an XROOTD based caching infrastructure with the HTCONDOR batch system to introduce data locality into job scheduling. By this, it influences the placement of data in multiple caches avoiding replication of data and matching jobs to suitable worker-nodes that already provide cached input data. The graph was previously published in [117–119].

An overview of the structure of NAVIX is shown in Figure 9.2. Each job that is submitted to the batch system HTCONDOR triggers a translate hook for calculating a score that describes how each cache fits to the job. The hook integrates the data locality information derived from the coordination service into the job description to influence the job to computing resources matchmaking of HTCONDOR. While the jobs are waiting for a suitable computing resource to get available, update hooks keep the calculated scores up to date. The score itself is calculated by the NAVIX coordination service, which extracts information from the caches using the XROOTD infrastructure, such as the location of the cached files.

The hierarchical structure of XROOTD illustrated in Figure 9.1 redirects the meta-data request from the top-level manager down to the individual caching proxy servers. During this step, the meta-data extracted from the caching proxy servers is cached at all intermediate XROOTD management servers, which enables an efficient aggregation of meta-data. This allows the approach to scale well as meta-data needs to be requested for a huge amount of files. Studies revealed that aggregation and intermediate caching of meta-data may introduce a time delay in the job coordination and scheduling. There is an additional time delay from job scheduling itself since the batch system collects submitted jobs to efficiently schedule them in a single processing step. Additionally, it takes time until the job actually reads the data. The time until data is read by the job is usually much larger than the time delay caused by the coordination process or the job scheduling. As conventional HEP jobs usually run for several hours, time delays caused by the aggregation of meta-data and scheduling of jobs in the order of minutes can be neglected. Hence, the concept assumes that

metadata derived from the caches can be outdated to a certain degree.

NAVIX collects information about the job itself required for calculating the data locality score from the HTCONDOR hook. Each job needs to provide a catalog of input files that it will process. Since this list cannot be automatically extracted from the configuration and encoding of the job, NAVIX requires the user to add this information when creating the job. The remaining necessary information about the jobs is extracted from the job description.

After aggregating meta-data of the job and the caches, NAVIX calculates a ranking of suitable caches sorted by the data-locality score. This ranking allows us to modify the HTCONDOR job scheduling and force HTCONDOR to match the job to the most suitable computing resources in terms of data-locality. The calculation of the data-locality score and thus the ranking of suitable computing resources is easily exchangeable and can be adjusted to the infrastructure and the political or financial aims. To a certain extent, this introduces a pre-scheduling of jobs that takes place before the actual HTCONDOR job scheduling. All kinds of meta-data can be included in the calculation of the data-locality score to optimize the pre-scheduling decision. Within the prototype setup, NAVIX chooses the cache that provides the highest fraction of already cached files to be the best. Including additional observables, such as the load of the individual caches or the utilization of the computing resources may increase the overall data throughput while reducing congestion of the infrastructure.

On the one hand, this pre-scheduling done by NAVIX introduces the desired behavior of sending jobs to the most suitable computing resource. This reduces the load on the HTCondor batch system as the scheduling algorithm, which runs single-threaded, just has to handle one additional restriction introduced by NAVIX. The calculation overhead is efficiently done by NAVIX. Since HEP jobs do not depend on each other, the procedure of extracting meta-data and calculating the data-locality score is performed on a per-job basis and can be highly parallelized trivially. On the other hand, the pre-scheduling of NAVIX can interfere with the job scheduling of HTCONDOR. In some cases, this can lead to inconsistencies in the scheduling decision, which must be handled.

The pre-scheduling decision of each job is transferred via the HTCONDOR translate hook to the job scheduling. Afterwards, the NAVIX decision is repeatedly updated by the HTCONDOR update hook while the job is waiting for suitable computing resources to become available. This also allows NAVIX to change the pre-scheduling decision according to the new situation, e.g. when new files are available in the caches or old ones are removed. Furthermore, it allows us to detect and intercept the above-mentioned inconsistencies and soften the forcing of the pre-scheduling

decision in case a job remains in the scheduling phase for too long.. This allows the jobs to run on different resources and introduce caching of the input data there. Here, the duplication of data in the caches have to be balanced versus long waiting times of jobs that are waiting for the most suitable computing resource. Within this context, the optimization of the overall data throughput of the system competes with the benefits for users that want their jobs to finish as fast as possible.. Future research will address optimizing this field with the help of simulating the chronological sequence in the processing of jobs. After softening the forcing of the pre-scheduling decision step by step, the prototype setup completely releases jobs. This catches scheduling inconsistencies or enables the duplication of data in the caches to overcome the overload of computing resources. Especially if a huge number of jobs need to be processed, caching the files on multiple caches can increase the overall performance.

After the job has finished, the finalize hook is called to analyze the job and its performance. NAVIX as an intermediary between XROOTD and HTCONDOR suits perfectly for aggregating monitoring information about the job and cache performance. This allows concluding how the job profits from caching and how the scheduling of job to resources and data to caches can be fine-tuned.

9.2.3 Scalability of data and job coordination

While realizing distributed coordinated caching and developing the coordination service NAVIX, special attention was given to the scalability of the approach. The choice of HTCONDOR and XROOTD is based on both their wide distribution in the HEP community and their built-in scalability. HTCONDOR supports to send jobs from one pool of computing resources to another one. This makes it possible to merge different resource pools, each of which operating its HTCONDOR batch system, under a superimposed HTCONDOR batch system. The hierarchical scalability of XROOTD illustrated in Figure 9.1 allows accessing multiple caching proxy servers via management servers. Since multiple management servers can also be structured by a superimposed management server, the caches of multiples resource pools can be united.

Scalability is included in the design principles of NAVIX to enable the distributed coordinated caching approach to scale. Therefore, it exploits that XRootD and HTCondor already take care of hierarchical scaling. As Figure 9.3 shows in an exemplary view, multiple resource pools, each of them is managed by its own HTCONDOR and NAVIX instance, are combined in a superimposed resource pool. While the individual resource pools take care of data locality on the scale of individual worker-nodes or




Figure 9.3: XROOTD and HTCONDOR already support hierarchical scaling. This allows to combine multiple local pools of computing resources into regional superimposed systems. Several regional resource pools are merged into a national one. NAVIX considers only the data locality in the granularity of the direct subsystems, i.e. the computing resources one level down. This allows for scaling the complete approach easily.

computing clusters when scheduling jobs, the superimposed resource pool considers data locality at the scale of whole computing clusters. As multiple such superimposed resource pools can again be combined within an superior system that coordinates jobs and data on a higher level, the approach is highly scalable. As NAVIX only interacts between the HTCONDOR and XROOTD management on the same level and considers data location information of the corresponding subordinated systems, the scalability of the overall system is guaranteed. Here, various institute resources are integrated into a local resource pool, while multiple local resource pools are merged into a regional resource pool. Multiple regional resource pools are in this example merged into a national pool. Within each level, NAVIX just considers data locality at the granularity of the directly involved subordinated systems.

The approach of scalability of the setup was tested for a few computing resources building two pools with two independent NAVIX instances. This allowed to route jobs within the superimposed batch system to the resource pool that already provided cached files. The NAVIX instance of the subordinated system then decides which worker-node suits best in terms of data-locality. Larger scale testing is part of the future merging of German CMS Tier 3 resources [122].

9.3 Benchmarking of prototype systems

The advantages of the distributed coordinated caching concept was tested using a prototype system. This was realized at the Institute of Experimental Particle Physics

(ETP) using the above-mentioned coordination service NAVIX, an HTCONDOR batch system, and an XROOTD caching infrastructure.

The conclusions drawn from the prototype system were used to implement a dedicated high-throughput computing cluster. In such a cluster all hardware components are balanced and extra cache storage is installed to optimize the data throughput for data-intensive workflows. This approach was tested on a small-scale computing cluster, before building a production system at the Tier 1 WLCG computing center GridKa [123, 124].

Furthermore, the applicability of caching in dynamic environments such as cloud and HPC resources to optimize the limited data throughput of such resources was investigated. Here, well-placed caches can relieve the limited network bandwidth and make computing resources available for data-intensive workflows.

9.3.1 Small-scale testing setup




Figure 9.4: Small-scale testing setup with three worker-nodes connected to an HTCONDOR batch system. Each of them is equipped with a Solid State Drive (SSD)-cache via its own XROOTD caching proxy server. Jobs are matched to the worker-node that already provides cached files via the coordination service NAVIX.

This prototype system consists of three worker-nodes that are managed by an HTCONDOR management instance. Each of them provides its own cache space via an XROOTD caching proxy server as illustrated in Figure 9.4. All caches are kept within a hierarchical XROOTD infrastructure under a top-level management server. A NAVIX instance intermediates between the top-level management server and the HTCONDOR management. The benchmarks used a basic HTCONDOR and XROOTD setup, as well as a simple pre-scheduling logic implemented in NAVIX. Here, NAVIX

prefers the worker-node with the highest percentage of files already cached for each job. While it forces HTCONDOR for a while to assign the job to the most suitable worker node, it releases the constraint step by step afterward. All proxy servers use the default XROOTD caching logic. This means, that all accessed files are cached block-wise on-the-fly. If the cache volume reaches its limit, the least recently accessed files are deleted. Each cache volume consists of a software RAID 0 of SSDs that provide in total about 2 TB cache space per worker-node. The caches are accessed internally within the worker-node with about 9.6 Gbit s^{-1} . Access to the cached files is restricted to the respective caching proxy server of the worker-node. Hence, NAVIX treats all worker-nodes as separate systems, which allows testing the efficiency of coordinating data and jobs.

While all three worker-nodes share a 10 Gbit s^{-1} LAN interconnection, they can access data of remote Grid Storage elements with about 6 Gbit s^{-1} . The data sets used for benchmarking purposes are stored on such common WLCG storage resources. Benchmarks performed in preparation showed that up to 60 files stored on these storage systems can be accessed at an average of about 100 Mbit s^{-1} .

Two different types of benchmarks are performed. A performance benchmark measures the maximum achievable data throughput of the setup by transferring data from remote Grid storage elements without performing any calculations. In addition, the benefits of distributed coordinated caching for typical data-intensive HEP workflows is estimated by using the CMS jet energy calibration workflow as a benchmark. Out of the chain of tools used for determining both the jet momentum resolution and the jet energy scale, the data-intensive selection step described in Section 5.2 is used for benchmarking. Both benchmarks are repeatedly processed while the fraction of cached files is artificially increased within each step by 10%. As Equation 9.1 indicates, the best performance is expected at a certain fraction of cached files, at which both the bandwidth of the network and the cache are simultaneously saturated. This allows checking the functionality of coordinating jobs and data and comparing the predicted and measured performance.

In the following sections, the results of the performance benchmarks are discussed first and subsequently those of the HEP workflow benchmarks.

9.3.2 Benchmarking data throughput and coordination efficiency

The performance benchmark submits test jobs that only read data without further processing to the HTCONDOR batch system. Data is streamed from the remote WLCG




Figure 9.5: While a step-wise increasing fraction of cached files in multiple benchmark runs, the duplicity of files in all three caches is monitored. Each bar in the plot monitors a time window that corresponds to one step of the benchmark that covers a certain fraction of cached files. In each step, the number of files accessed and their duplicity in the caches is measured. Since there is no file cached more than once, the coordination of jobs to already cached data appears to succeed. Otherwise, processing jobs on different worker-nodes would increase file duplicity. The graph was first published in [119].

Tier 1 computing center GridKa via WAN connection to the computing resources of the institute located at the ETP. By repeatedly sending these performance benchmark jobs and increasing the fraction of already cached data, mismatches between jobs and caches within the coordination process can be detected. Any job that is not properly coordinated to the most suitable worker-node in terms of data locality would trigger additional caching of files on different caches. Over multiple submits of the same jobs, this would cause duplication of files in the caches. Figure 9.5 shows the file duplicity in the three caches for the repeated processing of the same jobs while increasing the fraction of cached files. The benchmark results show that the expected 10% increase in the fraction of cached data per submission step is visible. It is observed that files are only cached once and not duplication of files in multiple caches occurs. Hence, it can be concluded that the basic coordination of jobs to already cached data succeeds. The amount of data processed within the performance benchmark was adjusted to fit to the total cache volume of all worker nodes. As no file is missing in the cache at the last benchmark step, the XROOTD caching is working as expected.

Within the same benchmark, the average data throughput of the jobs depending on the fraction of cached files is measured, as shown in Figure 9.6 (a). The error bars shown in the plot indicate the variance of the set of jobs submitted for one benchmark run. Increasing the fraction of already cached files up to 80% increases the data throughput by roughly a factor of 5. At this point, the jobs saturate the




Figure 9.6: Within the performance benchmark, the data transfer rate and the processing time of the jobs are monitored while step-wise increasing the fraction of already cached files. Using Equation 9.1 allows us to predict the influence of the cached fraction of files on the processing time. Comparing prediction and measurement shows quite a good agreement. In both graphs shown, the error bars indicate the variance of the measured set of jobs. Graph (b) was first published in [117].

bandwidths of the network and caches simultaneously to derive the maximum possible data throughput. Below this point, the network bandwidth is fully utilized while the caches are not. Beyond 80%, the utilization of the bandwidth of the network decreases, and, thus, the data throughput is slightly reduced. The last data point at 100% has to be treated separately since no additional caching of files for the next step is performed. Therefore, the load it causes slightly differs from the one of other benchmarking steps, which has a minor influence on the results of the performance benchmark. For technical reasons, this implementation was chosen. However, conclusions drawn from the trend of the data points are not affected.

Furthermore, the average processing time of the benchmark jobs are measured versus the fraction of already cached files. As Figure 9.6 (b) shows, a large reduction of the average processing time is observed while increasing the fraction of cached files up to 80%. Beyond that point, the processing time slightly increases. An optimal working point is visible, at which the maximum data throughput is achieved by saturating both the network and the cache bandwidths. Equation 9.1 allows calculating a prediction of the processing time.

The following general conditions of the setup are used for this purpose.

- **Total amount of data:** 2.16 TB of data is processed by 60 jobs, each of them handles 10 files.

- **Batch resources:** The test setup provides 60 job slots on 3 worker-nodes via HTCONDOR.
- **Network data throughput:** Access to remote Grid storage elements at the WLCG Tier 1 GridKa is limited to 750 MB s^{-1}
- **Cache read rate per worker-node:** Each cache volume provides 1.2 GB s^{-1} read rate.
- **Data throughput limit of the workflow:** For purely transferring files via XROOTD, the workflow limitation can be neglected.

Comparing the data points of the prediction with the measured ones shows a good agreement within the variances extracted from a set of benchmarks. The predicted processing time confirms the observed optimal working point as well as the measured trends. This optimal working point is specific for the exact hardware setup and the workflow itself.

9.3.3 Benefits for data-intensive HEP analysis workflows

Using the jet energy calibration workflow presented in this thesis allows estimating the benefits for typical data-intensive HEP analysis workflows. The data-intensive selection step of the analysis described in Section 5.2 is submitted as benchmark to the throughput optimized testing setup. These jobs apply a basic selection on a large amount of input data and, thus, dramatically reduces the amount of data for further processing. Since the selection does not require advanced calculations, the workflow usually suffers in terms of data throughput. The workflow itself uses the ARTUS framework that is also used by other particle physics analysis groups at KIT and Rheinländisch-Westfälische Technische Hochschule Aachen (RWTH Aachen University). Hence, it suits perfectly for benchmarking the advantages of distributed coordinated caching for such data-intensive HEP workflows.

The conditions of the measurement are the same as for the performance benchmark in the previous section. However, access to the remote storage was reduced to about 375 MB s^{-1} due to congestion of the network during working hours. The jobs processed a total volume of data of about 1.2 TB of data.

The benefits are measured in terms of CPU efficiency and processing time of the jobs. It is assumed that data-intensive workflows suffer from inadequate data throughput,




Figure 9.7: Submitting the selection step of the jet energy calibration workflow to the testing system allows us to estimate the benefits of distributed coordinated caching for typical data-intensive HEP analyses. Therefore, the CPU efficiency and the processing time of the jobs is measured per given fraction of cached files. The prediction of the processing time is calculated according to Equation 9.1. In both graphs shown, the error bars indicate the variance of the measured set of jobs. Graphs were first published in [119].

which is reflected in the CPU efficiency of the jobs. While waiting for data to get transferred to the processing unit, the CPU is not fully utilized. Increasing the fraction of cached files, as shown in Figure 9.7 (a), increases the data throughput and, thus, the CPU efficiency. The jobs derive an increase of the CPU efficiency from 58% to 88% by caching at least 20% of the input files. A further increase of the CPU efficiency is limited by the algorithm efficiency of the workflow itself. To reach even higher CPU efficiencies, calculations done within the event selection need to be optimized. Beyond the optimal working point at 20% cached files, the data throughput provided by the caching setup exceeds the limitations of the workflow.

The same effect is observed in the processing time shown in Figure 9.7 (b). This shows that the workflow itself is limited to a maximum data throughput of about 9.5 MB s^{-1} due to its algorithm efficiency. Due to this limit, the improvement of the processing time reaches a plateau beyond 20% cached files.

Increasing the CPU efficiency absolutely by 30% dramatically reduces the processing time of the workflow by the same amount. Thus, the user gets the results earlier and the resources are freed faster to process other workflows. From a financial point of view, the costs of additional cache volumes (cost for 1 TB SSD \approx €200 per worker node) in combination with the data throughput optimization via software has to be

compared to around 30% more computing resources necessary to achieve the same speedup (cost per worker-node \approx €2500). In the case of the testing cluster, the caching solution would be about 76% cheaper compared to purchasing additional computing resources. Costs for improving the network and remote storage resources as well as for electricity are not considered in this calculation. However, adding worker-nodes instead of caching input files increases congestion of the network and Grid storage elements and decreases the achievable CPU efficiency further. Beside boosting data-intensive jobs, caching input data also reduces the load on network and remote storage resources. Thus, jobs that do not directly profit from caching do benefit from the freed network, computing, and storage resources.

This section showed that the general concept of distributed coordinated caching can be realized. Coordinating data and jobs allowed to reduce the duplication of files in the caches and, at the same time, increase the cache hit rate. This allows for improving the overall data throughput of the computing cluster.

The applicability of the concept for other resource types such as HPC computing centers or cloud resources is studied in the next section. Since the test cluster was highly optimized by the caching approach, a production system for a high-throughput cluster was designed and realized as described in Section 9.3.5.

9.3.4 Applicability for shared computing resources

Computing resources such as HPC clusters and commercial or private cloud resources are designed for shared usage by multiple users of different communities. Using such resources for processing HEP jobs enables to satisfy the dramatically increasing demand for computing resources. At cloud resources and HPC clusters, the network, computing, and storage resources are shared among different users. Depending on the utilization by other users, fluctuating performance of the network, computing, and storage resources poses challenges to the optimization of data throughput. Congestion and limits of resources especially affect the processing of data-intensive workflows, while fluctuating performance of resources complicates optimizing data throughput of the setup.

This section studies the suitability of the distributed coordinated caching concept for computing resources that are not designed for data-intensive HEP analyses. Two example setups are used to test the applicability and benefits of the concept: The NEMO HPC computing cluster and the OTC.

Studies revealed inefficient processing of data-intensive HEP workflows on such resources [83]. Caching may allow making this kind of computing resources available for processing data-intensive HEP analyses.




Figure 9.8: Caching structure that is applicable for shared resources such as HPC clusters or cloud resources, where the worker-nodes do not provide a suitable storage volume for caching purposes. Instead, a dedicated caching proxy server that is connected to a suitable cache volume serves as a global cache for all worker-nodes within the computing cluster.

For testing purposes, the mentioned resources are made available within the HT-CONDOR cluster of the institute [125–127]. This allows NAVIX to coordinate jobs to these resources according to data-locality. Since the worker-nodes of both resources do not provide extra storage volumes that can be used as cache volumes, one dedicated caching proxy server is placed at each of the resources. This proxy is running on an extra host that provides additional disk space that is not accessible from the worker-nodes directly. As outlined in Figure 9.8, this single cache serves all worker-nodes within the computing resource. The performance gain is analogously measured to the testing setup described in the previous section. Benchmark jobs for determining the maximum achievable performance are submitted repeatedly while increasing the fraction of already cached files to scan for an optimal working point.

As an example of an HPC computing center, the NEMO computing cluster [83] is utilized. This cluster is located at the University of Freiburg and was funded by the state of Baden-Württemberg for research in different institutes and subject areas. As a typical HPC cluster, it is designed for CPU-intensive jobs that may require hundreds of CPU cores to communicate with each other. Here, data-intensive jobs suffer from limited network bandwidth to remote storage servers, since all 900 worker nodes,

with 18000 CPU cores in total, share two 40 Gbit s^{-1} external network connections. The cache uses an HPC storage volume that offers a work-space of about 20 TB per user to cache job input data. This allows us to speed up data-intensive HEP analysis workflows that process data repeatedly, as iterated access exploits the internal bandwidth provided by the HPC infrastructure. Using the internal bandwidth of 1 Gbit s^{-1} between the worker-node and the caching proxy server allows reducing pressure on the limited external network interface. The exact bandwidths available between the CPU core and the remote storage resources or the internal cache highly depends on the network utilization caused by other jobs and users.

As an example of a commercial cloud resource, OTC resources were used for testing the benefits of caching. Such general-purpose resources are usually not designed for high data throughput as they should serve a wide range of computing activities. While network, storage, and computing resources at commercial or private clouds are usually limited, high network bandwidth or excessive network traffic may cause additional costs. To circumvent network traffic between external remote storage resources and the worker-nodes spawned at OTC, a dedicated host is installed that is connected to a 2 TB cache volume for testing purposes. The interconnection between the worker node and caching proxy server or remote storage system is unknown since no detailed information is provided by OTC for the storage, network, and computing capacities.

The results of the maximum performance benchmark for both the NEMO HPC cluster and the OTC computing resources are shown in Figure 9.9 (a) and (b), respectively. After repeating the measurement after one week, a different performance behavior was observed within both shared systems. As mentioned above, a varying performance over time according to the load on storage and network resources caused by other users within the shared infrastructure is expected. At the NEMO HPC cluster, the performance benchmarks indicated that the performance of the distributed file system used as cache volume highly depends on parallel utilization by other users. The effect is clearly visible by comparing the processing times of the jobs at 100% cached files. Furthermore, the difference of the points at zero cached files indicates that the network connection to remote resources slightly varies over time. This is most likely caused by the different amounts of other jobs transferring data from or to remote storage systems. Both effects cause the optimal working point to be time-dependent and vary between caching 40% and 80% of the input files. At the OTC worker-nodes, an even stronger deviation in the performance of both test runs is visible. While the first one suffers from a limited I/O-rate of the cache volume, congestion on the network connection to remote storage resources restricts the performance of the second test run. The optimal working point varies from 20% to 80% already cached files.




Figure 9.9: A maximum performance benchmark was executed on the HPC center NEMO and at OTC resources. The average processing time of the jobs was measured while step-wise increasing the fraction of cached files. The benchmarks are repeated after a week to estimate the time dependency of the results. All results confirm a fluctuating performance of the shared computing resources. In both graphs shown, the error bars indicate the variance of the measured set of jobs. Parts of the measurements were first published in [119].

Both performance measurements showed that the cache utilization needs to be balanced with network performance. This decision must be dynamically adapted over time to the evolving conditions of the systems to profit from caching. Whether improvements by caching input files can be achieved and how much benefit data-intensive HEP jobs get strongly depends on the resource provider and the utilization of computing resources. Caching makes resources that provide enough temporary storage space for caching purposes available for processing data-intensive HEP workflows. Distributed coordinated caching, however, efficiently uses multiple such resources at the same time, since caches are managed and jobs are scheduled to the most suitable cache or worker-node, respectively. At commercial cloud providers, the additional costs due to the increased data transfer must be weighed against the costs for additional cache storage space. Especially HPC computing resources, which usually provide a distributed file system as a work-space for users, can, therefore, be easily made available for data-intensive HEP jobs with improved performance. In both cases, caching a certain fraction of input data and, thus, reducing congestion on the network improves the processing of data-intensive HEP workflows. Optimizing such shared computing resources that are not dedicated to HEP increases the total amount of resources that can be utilized for analysis processing. This allows us to face the high demand for computing resources within current and future HEP experiments.

9.3.5 Dedicated high throughput computing resources

The experience gained in optimizing the data throughput of the test cluster led to the design of a production system for high-throughput data analysis that benefits from distributed coordinated caching. This so-called Throughput Optimized Analysis System (TOpAS) cluster was installed at the WLCG Tier 1 data center GridKa in cooperation with the local operations team. This cluster is dedicated to data-intensive HEP workflows to reduce their processing time and enable short turn-around cycles. Hence, it fits perfectly for processing parts of the jet energy calibration that is presented within this thesis and requires repeated processing on a short time-scale.




Figure 9.10: Design of the TOpAS cluster dedicated to high-throughput data analysis. Two cache layers allow us to boost data throughput and reduce the load on Grid storage and network resources. A distributed file system spread across all 11 worker-nodes serves as 1 PB cache. Additionally, one high-performance Non-Volatile Memory express (NVMe) SSD per worker-node provides a fast accessible 1 TB cache volume. Both caches can be accessed via XROOTD caching proxy servers internally within the worker-node.

The overall design of the TOpAS cluster is outlined in Figure 9.10. Here, the optimizations focus on two layers of caches that exploit a fast accessible distributed file system and one high-performance NVMe SSD per worker-node. For building the distributed file system, a Ceph file system (CephFS) that connects 176 Hard Disk Drives (HDDs) to a total cache volume of about 1 PB is used. Even though HDDs are not made for high-performance operations, both the large number of HDDs and a 100 Gbit s^{-1} interconnection between the worker-nodes lead to a huge storage volume that provides large data transfer rates. Each SSD cache, however, only supports the worker-node where it is installed. This requires the SSD caches to be coordinated to avoid wasting limited cache space due to the duplication of data in multiple caches.




Figure 9.11: Analysis-like benchmarks that apply a basic selection on a data sample were performed to determine the maximum performance when accessing the TOPAS cache volumes. The NVMe SSDs performed slightly better than the distributed file system CephFS. However, both show a relatively stable, high data transfer rate up to 560 parallel processes, which corresponds to the number of logical cores delivered by the setup. For the CephFS benchmarks, a classical network setup was used without the modern optimization RoCE that allows remote direct memory access via network. The average CPU usage of the performance benchmark represents its CPU efficiency since no advanced calculations are performed with data. The performance benchmarks applied to a set of 10 TOPAS worker-nodes were published in [128].

First performance benchmarks shown in Figure 9.11 apply a basic selection to a data set to mimic a data-intensive HEP analysis workflow. The performance benchmarks read data directly without an intermediate XROOTD caching proxy server. This allows us to measure the maximum performance jobs may gain when accessing data provided by the caches. The benchmarks showed that the performance of the SSD caches is better than the distributed file system. In both cases, however, a relatively stable, high data throughput rate was observed when performing multiple benchmark processes in parallel. Even for 560 processes reading data either from the SSD caches or the single distributed cache, a stable data read rate of more than 25 MB s^{-1} per process was measured as illustrated in Figure 9.11 (a). This roughly corresponds to the case when all CPU cores of the cluster are processing analysis workflows that read from the cache. If the caches are only accessed by a part of the jobs while other jobs are reading data via the network, the setup can process data with up to 55 MB s^{-1} . This bandwidth suits well for data-intensive analysis workflows since most HEP analyses are based on the scientific software toolkit ROOT that at the time of writing this thesis limits the data throughput to about 50 MB s^{-1} per CPU core at maximum. The CPU efficiency illustrated in Figure 9.11 (b) is almost ideal at 1.0, especially when accessing the SSD caches. In general, the caches at the

TOPAS cluster deliver a high data throughput rate that fits perfectly for processing data-intensive HEP workflows. Additionally, the cluster profits from a fast network connection to WLCG Grid storage resources that provides up to 200 Gbit s^{-1} transfer speed. The combination of the two cache layers and the network connection allows boosting the data throughput even higher while reducing the load on the storage resources and the external network connection of the TOPAS cluster. Performance benchmarks confirmed that the cluster is able to achieve a data throughput of about 80 MB s^{-1} per CPU slot at maximum if 50% of the input data on the SSD caches and the process is not limited by ROOT. The residual fraction is transferred from Grid storage provided by the GridKa Tier 1 computing center.




Figure 9.12: Submitting the selection step of the jet energy calibration workflow to the testing system allows us to estimate the benefits of distributed coordinated caching for typical data-intensive HEP analyses. Therefore, CPU efficiency and the processing time of the jobs is measured for certain fractions of cached files. The prediction of the processing time is calculated according to Equation 9.1. Graphs were first published in [119].

Performing the jet energy calibration workflow as a benchmark allows to estimate the benefits of caching for data-intensive HEP workflows. Figure 9.12 shows the benchmark results for caching various fractions of cached files per job. Even without caching any files, the design TOPAS cluster reduces the processing time of the jobs dramatically from about 53 minutes to 30 minutes comparing the results with the ones of the testing setup in Figure 9.7 (b) and Figure 9.12 (b). Since files are stored on GridKa Tier 1 Grid storage elements, jobs already profit from the fast network connection to the storage resources. Caching the input files on the NVMe SSDs allows reducing this even further to about 25 minutes. Here, data throughput is limited by the algorithm efficiency of the jobs and allows no further improvement in


the processing time. To profit from the even higher data throughput rates provided by the TOpAS cluster, the implementation of the workflow needs to be optimized. As data throughput is limited by the workflow, the slight decrease of the processing time can be explained by lower latency when accessing cached files. This increases the CPU efficiency shown in Figure 9.7 (a) to 95%. Here, the CPU efficiency reached the maximum, since the algorithm efficiency of the workflow itself limits the CPU efficiency to that value.

After benchmarking the SSD caches, the performance of the CephFS distributed file system is tested. The same jet energy calibration jobs are submitted to the TOpAS cluster and compare the measured results without and with the caching of input files. Figure 9.13 (a) shows the distribution of data throughput versus CPU efficiency for all jet energy calibration jobs used for benchmarking when accessing data from remote Grid storage elements. In contrast, Figure 9.13 (b) concentrates on jobs whose input files are cached on the distributed file system. In both benchmarks, two different clusters of points that are caused by processing recorded and simulated data sets are observed. The difference relies on the increased fraction of events that pass the selection within the simulated data sets compared to the recorded ones. The more events are rejected early in the selection procedure, the fewer calculations are required for the same amount of data. In both clusters, a linear dependency between data throughput and CPU efficiency is observed. The linear dependency represents a limitation of the workflow due to data access. Increasing the data throughput thus implies a higher CPU efficiency. Comparing the benchmark results with and without cached files, a shift towards higher CPU efficiency and data throughput by caching input files is clearly visible. The average CPU efficiency is increased by roughly 10% to about 93% and its spread is reduced. Hence, caching input files on the distributed file system also boosts the performance of data-intensive jobs.


Increasing the data throughput and reducing the latency when accessing files brought the jet energy calibration workflow to the maximum achievable performance. Any further optimization requires improving the code, e.g. by using modern techniques for efficient columnar data analysis. The processing time and, thus, the turnaround cycles are reduced by roughly 60%. This accelerates the processing of new intermediate results for the calibration of the jet energy and the jet momentum resolution. Hence, final scale factors of the calibration can more quickly be made accessible to other physical analyses.

The presented performance benchmarks utilize data on the Grid storage element at GridKa, which is connected to the TOpAS cluster with a high bandwidth. When accessing remote data provided by other WLCG computing centers or Tier 3 resources such as the Tier 2 computing centers in Aachen or Hamburg, the processing jobs are

bound to the shared external network connection of the computing center instead of exploiting the internal one. The bandwidth to those computing centers is limited to 100 Gbit s^{-1} or 10 Gbit s^{-1} respectively. Thus an even higher improvement of the performance of data-intensive workflows is expected when processing cached data.



(a) Calibration jobs accessing remote files



(b) Calibration jobs accessing cached files

Figure 9.13: Distribution of jet energy calibration jobs, shown as a diagram of data throughput versus CPU efficiency for two scenarios, accessing data from remote Grid storage at the top, (a) and cached on CephFS at the bottom, (b). Processing recorded data and simulated data resulted in two different clusters of points. A dependency between data throughput and CPU efficiency reflects a limitation of the workflow caused by data access.

Conclusion and outlook

The measurements performed at the CMS experiment are based on the precise reconstruction of the particles to be analyzed. At hadron colliders, the accurate measurement of the jets, which are crucial for almost any analysis, requires a special effort in calibration, since numerous effects bias their measurement.

This thesis concentrates on the entire data set recorded by the CMS detector in 2018 corresponding to 58.83 fb^{-1} integrated luminosity. Here, we observed that LHC run 2 with 13 TeV center-of-mass-energy and unprecedented instantaneous luminosity caused a massive contribution of pile-up, parallel proton-proton interactions, and a drastic increase of additional hadronic activity. Since those effects massively bias the reconstruction of compound objects such as jets, a high effort in the calibration of jets is required. At the CMS experiment, the calibration of jets comprises the work of several analysis groups, each correcting for a certain effect, to adjust the jet energy scale and the jet momentum resolution.

This thesis focuses on $Z + \text{jets}$ events in which the Z boson decays into two muons. By this, it exploits the important strength of the CMS detector in precisely reconstructing muons and, thus, accurately reconstructing Z bosons. All presented analysis procedures are also applicable to the $Z + \text{jets}$ channel in which the Z boson decays into two electrons. Balancing the jet that needs to be calibrated with the Z boson in the transverse plane enables calibrating jets in average. This includes both the jet energy scale and the jet momentum resolution.

The results of the corrections of the jet energy presented in this thesis made an important contribution to the data-driven correction of the jet energy scale. Two complementary balancing methods were used, the Missing E_T Projection Fraction




Figure 10.1: Jet momentum resolution scale factors derived with an advanced extraction procedure. These are extracted depending on $|\eta^{\text{jet}1}|$ due to heterogeneous detector layers in the longitudinal direction.

(MPF) and the p_{T} -balance method. Both allow to estimate the detector response for jets. The derived scale factors correct the jet energy scale for residual differences between the MC simulation. They provide precisely calibrated jets that are perfectly suited for further analysis while introducing only small uncertainties below the percentage level.

On top of that, this thesis proposes a new procedure to determine the jet momentum resolution based on the p_{T} -balance response. Contributions to the measured p_{T} -balance resolution coming from the Z boson measurement or an imbalance between the balancing objects on the particle-level have successfully been eliminated. Within the analysis, a strong influence on the measurement of the jet momentum resolution caused by pile-up events and jet-splitting due to additional hadronic activity was observed. As shown in this thesis, an advanced procedure to eliminate the influence of these effects made it possible to extract the actual jet momentum resolution. The resolution of the jet momentum extracted from data was compared to the one extracted from MC simulation. Figure 10.1 presents the derived scale factors that compensate for insufficient modeling of the detector resolution in the MC simulation.

Using Z + jets events, this analysis is most sensitive to the central detector region $|\eta| < 1.3$, which shows low statistical uncertainties and scale factors close to 1.2. Although muon reconstruction is limited to $|\eta| < 2.4$, scaling factors are also estimated for detector regions further forward. Within the forward region, however, the analysis suffers from the lack of events, which leads to larger statistical uncertainties of the scaling factors. The derived scale factors were successfully verified within MC closure tests. Concluding, the feasibility study presented in this thesis clarified that the

proposed procedure is suitable for deriving jet momentum scale factors using $Z + \text{jets}$ events.

Within the sensitive region, the presented jet momentum resolution scale factors agree well with the ones established by CMS and derived from dijet events. As a next step for integrating the new procedure into the CMS jet calibration approach, it is necessary to study systematic uncertainties and propagated them through the presented extraction procedure. This enables cross-checking the $Z + \text{jets}$ and dijet results and, finally, combining them to achieve even higher precision at simultaneously lower uncertainties.

However, the calibration tasks performed in the context of this thesis require large data sets to be processed iteratively for adjusting and optimizing the analysis. Especially, when results, e.g. the jet energy calibration, rely on the outputs of several analysis groups, short turnaround cycles are essential.

Therefore, besides the physics analysis, this thesis also focuses on improving the performance of data-intensive HEP end-user analysis workflows. A distributed coordinated caching concept was conceived that considers previous studies done at the ETP. It optimizes the overall data throughput for data-intensive workflows by caching input data on caches placed in the distributed HEP computing infrastructure.

For testing the advantages of the concept, a prototype system based on common HEP tools, an HTCONDOR batch system and an XROOTD caching infrastructure, was developed. A coordination service called NAVIX in between both parts developed in the context of this thesis made it possible to schedule individual jobs to computing resources that already provide cached input data. This also influences the placement of data in the caches. Benchmarking the prototype system showed that the coordination of jobs to already cached data succeed using a basic coordination logic. On the one hand, this allowed reducing the replication of data in limited cache volumes. On the other hand, the performance of the workflow consisting of multiple jobs was boosted. The processing time of the jet energy calibration workflow used as a benchmark was drastically reduced by roughly a factor of 5 when caching a certain fraction of input data. At this point, the job exploits the bandwidth of both, the cache and the network connection to external Grid storage resources, to reach the maximum data throughput. Besides increasing the efficiency in processing large amounts of data, the coordinated distributed caching system also releases pressure from the storage and network resources. Other jobs that do not profit directly from the caching benefit from computing resources that are freed earlier and from reduced congestion on storage and network resources.

Alongside this, the applicability of the concept for shared computing resources such as High-Performance Computing (HPC) clusters and cloud resources was studied. Tests at the NEMO HPC cluster and the OTC cloud offered a highly fluctuating performance due to resource utilization by other users. The caching concept developed in the context of this thesis supports the coordination of files and jobs in a distributed infrastructure to enable an efficient use of caches and optimized data processing. Thereby, it is capable of dynamically adapting coordination to the changing performance of shared computer resources. This unlocks HPC or cloud resources for data-intensive processing of HEP workflows.

The conclusions drawn from benchmarking the prototype setup were used for deploying a production system at the GridKa computing center optimized for high-throughput data analysis. This system brings the jet energy calibration workflow to the maximum achievable performance by reducing the latency in accessing data and providing highly performant caches and network connection to WLCG storage resources. Indeed, the caching of input data decreased the processing time enormously by roughly 60%.

In the context of the performed analysis, the distributed coordinated caching enables fast turnaround cycles and accelerates the release of jet calibration results for all CMS analyses. It is generally expected that the presented concept will have an impact on the processing of any generic analysis with large input data sets.

In summary, both parts of this thesis make a substantial contribution in the terms of concepts, realizations and results to current analyses and future challenges in HEP. The calibration performed in this thesis including the newly proposed procedure to determine the jet momentum resolution using $Z + \text{jets}$ events leads to a better understanding of jets for practically all CMS analyses. Alongside this, the developed concept of distributed coordinated caching accelerates data-intensive workflows by efficiently using the distributed computing infrastructure. This becomes crucial in future HEP experiments with regard to the drastically increasing amount of data to be processed.

Appendix

A.1 Technical information for analysis

This section gives some technical details used for performing the $Z + \text{jets}$ analysis.

A.1.1 Data sets

Lists of $Z + \text{jets}$ data samples used for jet energy scale and momentum resolution analysis are given by Table A.1, Table A.2, and Table A.3.

The determination of the jet energy scale factors uses data of the double-muon channel, compared to a LO MadGraph MC simulation. The extraction of the jet momentum resolution focus on proofing the feasibility of the analysis procedure. Therefore, an extended NLO MC sample generated by aMC@NLO is compared to the double-muon data sample.

The data sets are preselected using the following CMS HLT tags for double-muon events.

- HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ
- HLT_Mu19_TrkIsoVVL_Mu9_TrkIsoVVL_DZ
- HLT_Mu17_TrkIsoVVL_Mu8_TrkIsoVVL_DZ_Mass3p8

Data set name	global tag	number of events
/DoubleMuon/Run2018A-17Sep2018-v2/MINIAOD	102X_dataRun2_Sep2018Rereco_v1	75499908
/DoubleMuon/Run2018B-17Sep2018-v1/MINIAOD	102X_dataRun2_Sep2018Rereco_v1	35057758
/DoubleMuon/Run2018C-17Sep2018-v1/MINIAOD	102X_dataRun2_Sep2018Rereco_v1	34565869
/DoubleMuon/Run2018D-PromptReco-v2/MINIAOD	101X_dataRun2_Prompt_v11	169225355

Table A.1: Data samples recorded by the CMS detector preselected by double-muon trigger

Data set name	cross section [pb ⁻¹]	number of events
/DY1JetsToLL_M-50_TuneCP5_13TeV-madgraphMLM-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v2/MINIAODSIM	877.8	68898175
/DY2JetsToLL_M-50_TuneCP5_13TeV-madgraphMLM-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v2/MINIAODSIM	304.4	20456037
/DY3JetsToLL_M-50_TuneCP5_13TeV-madgraphMLM-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v2/MINIAODSIM	111.5	5652357
/DY4JetsToLL_M-50_TuneCP5_13TeV-madgraphMLM-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MINIAODSIM	44.03	2817812

Table A.2: MC simulated Z+jets data samples binned in number of jets used for extracting jet momentum resolution

Data set name	cross section [pb^{-1}]	number of events
/DY1JetsToLL_M-50_LHEZpT_50-150_TuneCP5_13TeV-amcnloFXFX-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MINIAODSIM	316.6	10968454
/DY1JetsToLL_M-50_LHEZpT_150-250_TuneCP5_13TeV-amcnloFXFX-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MINIAODSIM	9.543	14009712
/DY1JetsToLL_M-50_LHEZpT_250-400_TuneCP5_13TeV-amcnloFXFX-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MINIAODSIM	1.098	7531281
/DY1JetsToLL_M-50_LHEZpT_400-inf_TuneCP5_13TeV-amcnloFXFX-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MINIAODSIM	0.1193	1245138
/DY2JetsToLL_M-50_LHEZpT_50-150_TuneCP5_13TeV-amcnloFXFX-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MINIAODSIM	169.6	24255781
/DY2JetsToLL_M-50_LHEZpT_150-250_TuneCP5_13TeV-amcnloFXFX-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MINIAODSIM	15.65	42787613
/DY2JetsToLL_M-50_LHEZpT_250-400_TuneCP5_13TeV-amcnloFXFX-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MINIAODSIM	2.737	16704440
/DY2JetsToLL_M-50_LHEZpT_400-inf_TuneCP5_13TeV-amcnloFXFX-pythia8/RunIIAutumn18MiniAOD-102X_upgrade2018_realistic_v15-v1/MINIAODSIM	0.4477	2829789

Table A.3: MC simulated Z + jets data samples binned in number of jets and p_{T}^Z used for extracting jet momentum resolution

A.1.2 Data certification

Valid run periods are selected using the so-called Golden JSON file provided by the Physics Validation group at the CMS experiment:

`Cert_314472-325175_13TeV_PromptReco_Collisions18_JSON.txt` [65, 129]

This file marks 58.83 fb^{-1} as valid run periods and reduces the amount of events recorded by the CMS detector by $\sim 7.6\%$.

A.1.3 Muon Identification

A muon that passes the tight Muon ID fulfills the following requirements [68].

- The muon originates from the primary vertex within a transverse impact parameter less than 0.2cm and a longitudinal impact parameter less than 0.5cm.
- The muon track uses at least six layers of the inner tracker and at least one hit in the pixel detector.
- It must be reconstructed as a tracker muon and a global muon.
- As tracker muon, it needs to have a segment that matches at least two muon chambers.
- As global muon, the global muon fit must have a $\chi^2/ndf < 10$.

A.1.4 Jet identification


A jet that passes the tight Jet ID with lepton veto fulfills the selections [69] listed in Table A.4.

	$ \eta \leq 2.6$	$2.6 < \eta \leq 2.7$	$2.7 < \eta \leq 3.0$	$3.0 < \eta \leq 5.0$
Neutral hadron fraction	< 0.9	< 0.9	-	> 0.2
Neutral electromagnetic fraction	< 0.9	< 0.9	> 0.02 and < 0.99	< 0.9
Number of constituents	> 1	-	-	-
Muon fraction	< 0.8 (for lepton veto)	< 0.8 (for lepton veto)	-	-
Charged hadron fraction	> 0	-	-	-
Charged multiplicity	> 0	> 0	-	-
Charged electromagnetic fraction	< 0.8 (for lepton veto)	< 0.8 (for lepton veto)	-	-
Number of neutral particles	-	-	> 2	> 10

Table A.4: Restrictions for tight Jet ID with lepton veto for CMS data of 2018.

A.1.5 Characteristics of muons, Z boson, and jets for the MadGraph data set

Chapter 5 shows the characteristics of the muons, the Z boson, the leading jet, and the second-leading jet using the aMC@NLO simulation. For comparison, the same distributions are shown here for the MadGraph simulation. They mainly show a similar behavior as the aMC@NLO distributions.



(a) p_T distribution of positively charged muon (b) η distribution of positively charged muon

Figure A.1: A shape comparison of the MadGraph MC simulation and the 2018 data set for p_T and η of the positively charged muon. MC simulation and data distributions are normalized to each other for shape comparisons.






Figure A.2: A shape comparison of the MadGraph MC simulation and the 2018 data set for the reconstructed mass of the Z boson. MC simulation and data distributions are normalized to each other for shape comparison.



(a) p_T distribution of the leading jet

(b) η distribution of the leading jet

Figure A.3: A shape comparison of the MadGraph MC simulation and the 2018 data set for p_T and η of the leading jet. The distributions of MC simulation and data are normalized to each other for comparing the shapes.



(a) p_T distribution of the second-leading jet

(b) η distribution of the second-leading jet

Figure A.4: A shape comparison of the MadGraph MC simulation and the 2018 data set for p_T and η of the second-leading jet. The distributions of MC simulation and data are normalized to each other for comparing the shapes.




Figure A.5: A shape comparison of the MadGraph MC simulation and the 2018 data set for E_T^{miss} . The distributions of MC simulation and data are normalized to each other for comparing the shapes.




Figure A.6: A shape comparison of the MadGraph MC simulation and the 2018 data set for p_T -balance and MPF. The distributions of MC simulation and data are normalized to each other for comparing the shapes.




Figure A.7: A shape comparison of the MadGraph MC simulation and the 2018 data set for the number of jets per event. The distributions of MC simulation and data are normalized to each other for comparing the shapes.

A.1.6 Manually removed or tuned bins during the jet momentum resolution extraction

Bins filled with few events makes it difficult for the variance extraction procedure described in Section 7.2 to extract valid results. Thus, all bins used in the analysis are checked. Not recoverable fit results are removed from further analyzing, while bins that extract valid results after manual adjusting the fit range are kept. Table A.5 lists all removed bins, Table A.7 all manually tuned bins.

Response distribution	p_{T}^Z -bin [GeV]	$ \eta^{\text{jet1}} $ -bin	α -bin
p_{T} -balance (MC)	[1000, 1500]	inclusive	[0.05, 0.30]
	[500, 700]	[2.043, 2.322]	[0.20, 0.25]
	[400, 500]	[2.500, 2.853]	[0.10, 0.15]
p_{T} -balance (Data)	[500, 700]	[0.783, 1.131]	[0.05, 0.30]
	[500, 700]	[0.000, 0.522]	[0.20, 0.30]
	[400, 500]	[1.131, 1.305]	[0.25, 0.30]
	[400, 500]	[1.305, 1.740]	[0.05, 0.30]
	[300, 400]	[2.322, 2.500]	[0.05, 0.30]
	[300, 400]	[2.043, 2.322]	[0.05, 0.30]
	[300, 400]	[1.930, 2.043]	[0.05, 0.30]
	[300, 400]	[1.740, 1.930]	[0.05, 0.30]
	[230, 300]	[2.500, 2.853]	[0.05, 0.30]
	[175, 230]	[2.964, 3.139]	[0.05, 0.30]
	[175, 230]	[2.853, 2.964]	[0.05, 0.30]
	[230, 300]	[2.500, 2.650]	[0.05, 0.30]
	[230, 300]	[2.650, 2.853]	[0.05, 0.30]
[175, 230]	[2.500, 2.650]	[0.25, 0.30]	
PLI (MC)	[400, 500]	[2.322, 2.500]	[0.25, 0.30]

Table A.5: List of bins that are removed from the analysis due to not-recoverable fits

Response distribution	p_{T}^Z -bin [GeV]	$ \eta^{\text{jet1}} $ -bin	α -bin	Fit range	
p_{T} -balance (MC)	inclusive	[2.964, 3.139]	[0.20, 0.25]	[0.8, 1.5]	
	inclusive	[2.853, 2.964]	[0.10, 0.15]	[0.8, 1.5]	
	inclusive	[2.500, 2.650]	[0.15, 0.20]	[0.7, 1.2]	
	[500, 700]	[0.783, 1.131]	[0.20, 0.25]	[0.55, 1.35]	
	[230, 300]	[1.740, 1.930]	[0.20, 0.25]	[0.5, 1.2]	
	[175, 230]	[3.139, 5.191]	[0.25, 0.30]	[0.3, 1.2]	
	[175, 230]	[2.964, 3.139]	[0.05, 0.10]	[0.6, 1.4]	
	[175, 230]	[2.964, 3.139]	[0.25, 0.30]	[0.4, 1.2]	
	[130, 175]	[2.964, 3.139]	[0.20, 0.25]	[0.5, 1.5]	
	[130, 175]	[2.964, 3.139]	[0.25, 0.30]	[0.2, 1.5]	
	[130, 175]	[1.131, 1.305]	[0.05, 0.10]	[0.75, 1.3]	
	[130, 175]	[1.131, 1.305]	[0.25, 0.30]	[0.4, 1.2]	
	[105, 130]	[0.522, 0.783]	[0.05, 0.10]	[0.8, 1.2]	
	[105, 130]	[2.853, 2.964]	[0.10, 0.15]	[0.8, 1.5]	
	[105, 130]	[2.650, 2.853]	[0.20, 0.25]	[0.65, 1.5]	
	[105, 130]	[2.322, 2.500]	[0.10, 0.15]	[0.8, 1.5]	
	[105, 130]	[2.322, 2.500]	[0.15, 0.20]	[0.8, 1.5]	
	[105, 130]	[2.043, 2.322]	[0.20, 0.25]	[0.65, 1.5]	
	[105, 130]	[1.930, 2.043]	[0.25, 0.30]	[0.7, 1.25]	
	[85, 105]	[2.500, 2.650]	[0.25, 0.30]	[0.65, 1.5]	
	[85, 105]	[2.500, 2.650]	[0.10, 0.15]	[0.7, 1.2]	
	[85, 105]	[2.043, 2.322]	[0.20, 0.25]	[0.8, 1.5]	
	[60, 85]	[2.322, 2.500]	[0.10, 0.15]	[0.65, 1.25]	
	[60, 85]	[2.043, 2.322]	[0.10, 0.15]	[0.65, 1.25]	
	[60, 85]	[1.930, 2.043]	[0.10, 0.15]	[0.7, 1.2]	
	[60, 85]	[1.131, 1.305]	[0.10, 0.15]	[0.7, 1.5]	
	p_{T} -balance (Data)	inclusive	[2.853, 2.964]	[0.10, 0.15]	[0.5, 1.5]
		[130, 175]	[1.930, 2.043]	[0.05, 0.10]	[0.8, 1.2]

Table A.6: List of bins (part 1), where the fit range was manually tuned to be able to extract valid results

Response distribution	p_T^Z -bin [GeV]	$ \eta^{\text{jet1}} $ -bin	α -bin	Fit range
PLI (MC)	inclusive	[3.139, 5.191]	[0.05, 0.10]	[0.8, 1.2]
	inclusive	[2.853, 2.964]	[0.05, 0.10]	[0.8, 1.2]
	[700, 1000]	inclusive	[0.20, 0.25]	[0.5, 1.1]
	[500, 700]	[0.522, 0.783]	[0.25, 0.30]	[0.5, 1.2]
	[400, 500]	[2.322, 2.500]	[0.20, 0.25]	[0.2, 1.5]
	[300, 400]	[2.650, 2.853]	[0.05, 0.10]	[0.44, 1.2]
	[300, 400]	[1.930, 2.043]	[0.15, 0.20]	[0.6, 1.2]
	[230, 300]	[2.853, 2.964]	[0.15, 0.20]	[0.6, 1.2]
	[230, 300]	[2.043, 2.322]	[0.25, 0.30]	[0.4, 1.2]
	[175, 230]	[2.500, 2.650]	[0.15, 0.20]	[0.65, 1.5]
	[175, 230]	[2.650, 2.853]	[0.20, 0.25]	[0.45, 1.2]
	[175, 230]	[2.650, 2.853]	[0.25, 0.30]	[0.44, 1.2]
	[175, 230]	[1.740, 1.930]	[0.05, 0.10]	[0.7, 1.2]
	[175, 230]	[1.740, 1.930]	[0.25, 0.30]	[0.5, 1.2]
	[130, 175]	[2.964, 3.139]	[0.25, 0.30]	[0.5, 1.5]
	[130, 175]	[2.043, 2.322]	[0.25, 0.30]	[0.5, 1.2]
	[130, 175]	[2.043, 2.322]	[0.20, 0.25]	[0.55, 1.2]
	[130, 175]	[1.930, 2.043]	[0.20, 0.25]	[0.5, 1.2]
	[130, 175]	[1.930, 2.043]	[0.15, 0.20]	[0.8, 1.2]
	[105, 130]	[2.650, 2.853]	[0.25, 0.30]	[0.65, 1.5]
	[105, 130]	[2.500, 2.650]	[0.15, 0.20]	[0.8, 1.5]
	[105, 130]	[2.500, 2.650]	[0.25, 0.30]	[0.5, 1.5]
	[105, 130]	[1.740, 1.930]	[0.25, 0.30]	[0.75, 1.3]
	[105, 130]	[1.740, 1.930]	[0.05, 0.10]	[0.8, 1.2]
	[105, 130]	[1.131, 1.305]	[0.05, 0.10]	[0.8, 1.2]
	[105, 130]	[0.783, 1.131]	[0.25, 0.30]	[0.7, 1.2]
	[105, 130]	[0.783, 1.131]	[0.05, 0.10]	[0.5, 1.5]
	[85, 105]	[1.930, 2.043]	[0.10, 0.15]	[0.8, 1.3]
	[85, 105]	[1.740, 1.930]	[0.10, 0.15]	[0.8, 1.3]
	[60, 85]	[2.322, 2.500]	[0.15, 0.20]	[0.8, 1.2]
[60, 85]	[2.322, 2.500]	[0.10, 0.15]	[0.2, 1.7]	
Z boson momentum resolution (MC)	inclusive	[3.139, 5.191]	[0.10, 0.15]	[0.9, 1.1]
	inclusive	[2.853, 2.964]	[0.05, 0.10]	[0.9, 1.1]
	[1000, 1500]	inclusive	[0.20, 0.25]	[0.9, 1.1]
	[105, 130]	[2.500, 2.650]	[0.15, 0.20]	[0.9, 1.1]
	[85, 105]	[2.650, 2.853]	[0.10, 0.15]	[0.9, 1.1]

Table A.7: List of bins (part 2), where the fit range was manually tuned to be able to extract valid results

Response distribution	$p_{\mathbf{T}}^Z$ -bin [GeV]	$ \eta^{\text{jet1}} $ -bin	α -bin	Fit range
generated jet momentum resolution (MC)	inclusive	[2.853, 2.964]	[0.05, 0.10]	[0.8, 1.2]
	[700, 1000]	[0.000, 0.522]	[0.05, 0.10]	[0.8, 1.4]
	[700, 1000]	[0.000, 0.522]	[0.25, 0.30]	[0.8, 1.4]
	[130, 175]	[2.964, 3.139]	[0.25, 0.30]	[0.7, 1.5]
	[105, 130]	[2.500, 2.650]	[0.15, 0.20]	[0.8, 1.4]
	[105, 130]	[2.500, 2.650]	[0.25, 0.30]	[0.8, 1.4]
	[105, 130]	[2.322, 2.500]	[0.10, 0.15]	[0.8, 1.2]
	[105, 130]	[2.322, 2.500]	[0.20, 0.25]	[0.8, 1.2]
	[105, 130]	[1.930, 2.043]	[0.20, 0.25]	[0.8, 1.2]
	[85, 105]	[1.930, 2.043]	[0.10, 0.15]	[0.8, 1.2]
	[60, 85]	[2.043, 2.322]	[0.10, 0.15]	[0.7, 1.5]
[60, 85]	[1.930, 2.043]	[0.10, 0.15]	[0.8, 1.4]	

Table A.8: List of bins (part 3), where the fit range was manually tuned to be able to extract valid results

A.1.7 Additional plots for jet momentum resolution MC closure test




Figure A.8: Full 2-dimensional MC closure test simultaneously binned in $|\eta^{\text{jet}1}|$ and p_T^Z . Therefore, the pull of the MC scale factor from 1.0 is calculated, which shows whether the deviations are within the statistical uncertainties.

Figure 7.7 within Section 7.5 concentrates on the 2-dimensionally binned pull of the MC scale factors. The individual values of the scale factor binned in $|\eta^{\text{jet}1}|$ and p_T^Z are shown as 2-dimensional map in Figure A.8. Their corresponding uncertainties are shown in Figure A.9. Bins with larger deviation from the expected value 1.0 are mostly covered by larger uncertainties as well. The closure of the method is, thus, successful, which validates the procedure.

A.1.8 Additional plots for derived jet momentum scale factors




Figure A.10: Illustration of the uncertainties for the 2-dimensional jet momentum resolution scale factors simultaneously binned in $|\eta^{\text{jet}1}|$ and p_T^Z .

Figure 7.10 in Section 7.6 shows the extracted scale factor values that correct the jet momentum resolution in MC simulation to data. The corresponding uncertainties are illustrated in Figure A.11 and the calculated pull is shown in Figure A.11. While the uncertainties increase with $|\eta^{\text{jet}1}|$, the pull is mostly uniformly distributed with a few bins that exceed 2σ deviation. Few events in most of the bins lead to rather large uncertainties compared to the scale factor values. For a more conclusive study that achieves more accurate results in two dimensions, more recorded events are required.

Figure A.12 shows the comparison between the official jet momentum resolution scale factors that are based on the second half of the CMS data set of 2018 (Run D) and the scale factors derived by the extraction procedure using Z + jets events as proposed in this thesis. A comparison with the first half of the data set (Run A, B, and C) can be found in Figure 7.11. However, the official scale factors of both parts of the data set show similar trends. Differences are covered by the uncertainties.




Figure A.11: Illustration of the pull for the 2-dimensional jet momentum resolution scale factors simultaneously binned in $|\eta^{\text{jet}1}|$ and p_T^Z .




Figure A.12: Comparison of the jet momentum resolution scale factors derived from Z + jets events as presented in this thesis and the official results from the dijets channel. Both are binned in $|\eta^{\text{jet}1}|$. The jet momentum resolution scale factors derived from Z + jets events are based on the full 2018 data set. For comparison, the official jet momentum resolution scale factors that are based on the second half of the CMS data set of 2018 (Run D) are used. A comparison with the first half of the data set (Run A, B, and C) can be found in Figure 7.11. However, the official scale factors of both parts of the data set show similar trends. Differences are covered by the uncertainties.

A.2 Computing

A.2.1 Draft for a general coordination service

The presented realization of the coordinated distributed caching concept using the coordination service NAVIX is based on an HTCONDOR batch system and an XROOTD caching infrastructure. While both systems are commonly used in HEP, restricting the approach to those tools limits its flexibility.

Thus, this thesis proposes a more flexible design for building a coordination service. As illustrated in Figure A.13, it consists of a sequence of multiple so-called proxies that




Figure A.13: Flexible design for a general purpose coordination service to realize distributed coordinated caching on multiple computing infrastructures. The batch system as well as the caching mechanism can be chosen freely.

connect a caching system with a batch system. Each proxy takes input and calculates a new output that is given to the next proxy. In the beginning, a BATCH SYSTEM ADAPTER triggers the processing of information when registering the submission of a job. Furthermore, it translates between any batch system and the internal namespace of the coordination service. This allows the BATCH SYSTEM ADAPTER to provide job attributes such as a list of input files to the coordination service to the sequence of proxies. The job attributes are then passed through all proxies to the last one, the CACHE INFO PROXY. It uses the CACHE STORAGE ADAPTER to create a priority list of caches per job, which evaluates the suitability of caches according to already cached input files. This requires the CACHE STORAGE ADAPTER to be able to get file attributes, such as the location or size of files, and cache attributes, such as the performance or occupancy of caches. Afterward, the NAME-SPACE PROXY translates the list of suitable caches into the corresponding batch system resources that are served by the individual caches. This step is necessary, as the batch system just knows

about the computing resources it manages and nothing about caches. Besides, an EXTERNAL INFO PROXY can use external information, such as political or financial aims, to adjust the priority list of resources. The required information can e.g. come from external monitoring services or other tools. Finally, the JOB INFO PROXY uses meta-data provided by the batch system to adjust the priority calculation for resources. For example, this allows avoiding computing resources that are already occupied or will not be available for the expected job runtime. This makes it possible to adapt the coordination decision or resource priority list respectively to a volatile computing environment where computer resources are constantly being added and removed. The final priority list is given to the BATCH SYSTEM ADAPTER to influence the scheduling decision of the batch system.

This design of a coordination service for realizing the distributed coordinated caching concept is independent of the batch system and cache infrastructure. For each batch system, a dedicated BATCH SYSTEM ADAPTER can be implemented that translates between the internal namespace of the coordination service and the batch system namespace. The same applies to the caching infrastructure. A dedicated CACHE STORAGE ADAPTER is implemented for each caching infrastructure that enables the exchange of a certain stack of meta-data. Besides, the algorithms within the individual proxies can dynamically be adjusted to the surrounding computing infrastructure, the type of usage, and the political or financial aims of the collaboration. This allows the concept to be applied to any type of infrastructure and opens up new fields for the allocation of computing resources for data-intensive HEP workflows.

List of Figures

2.1	Overview of the fermions, the gauge bosons, and the Higgs boson within the Standard Model of particle physics [1]. The gauge bosons reflect three of the four fundamental forces. Gravity is still missing in the model. Fermions are classified into quarks and leptons according to the forces to which they are responsive.	10
2.2	An illustration of a proton-proton scattering with a hard interaction. Two partons, each of them carries a momentum fraction of the proton according to the PDF, interact within a hard scattering. The resulting gluons and quarks cause collinear or soft radiation that lead to parton showers, hadronization to derive color-neutral compounds, and the decay of unstable hadrons. Besides, color-charged proton remnants are created that also hadronize. The graph is based on [3].	13
2.3	An illustration of a ISR and FSR, where a gluon is radiated by a quark in either the initial or the final state.	14
2.4	Feynman-graph of the production of a Z boson from a quark-antiquark annihilation decaying into a charged lepton pair.	15
2.5	Examples of Feynman-graphs of $Z + 1$ jet process via different production channels. The quark-gluon production channel either via an intermediate quark (a) or exchange of a quark (b), the quark-antiquark production channel (c), and the gluon production channel (d). Except for the latter, all the channels shown can have a complementary process in which an anti-quark instead of a quark interacts. While (d) represents an NLO process, the others are LO production channels.	17

3.1	Schematic view of the LHC accelerator complex [14]. A sequence of different accelerators prepare and pre-accelerate the hadron bunches before filling them into the actual LHC accelerator ring. It finally reaches up to 6.5 GeV kinetic energy per proton. Alongside, several adjacent experiments are supplied with hadron bundles of different energies.	20
3.2	Schematic overview, adapted from [25, 26], of the CMS detector shows the different subdetectors: The inner most tracking system, the ECAL, the HCAL, the magnetic coil, and the muon chambers within the magnetic return yoke. While the layout is almost uniform in radial plane, the detector is split into the central region, the outer end-caps, and the forward components at both ends.	22
3.3	Visualization of the conventions for the CMS coordinate system taken from [28]. The collision point in the center of the detector defines the origin, whereas the orientation of x-, y-, and z-axis is towards the surface, the center of the LHC ring, and the Jura mountains. Besides the Cartesian coordinate system, the rotational symmetries of the cylindrical detector allow to define a spherical coordinate system based on the angles θ and ϕ and the radius.	23
3.4	Beside the hard process at the primary vertex, the beam remnants can interact, the so-called underlying event. Additional proton-proton collisions are called in-time pile-up. Effects in the detector material from previous collisions are summarized as out-of-time pile-up. In-time pile-up, underlying event, out-of-time pile-up, and the beam remnants themselves superimpose the hard process distort a proper detection.	25
3.5	Side view of a cross section through one quadrant of the CMS detector [29]. It visualizes the different layers of subdetectors in size and spacial coverage in η : From inside to outside, there are the tracking system (light blue), the ECAL (green), the HCAL (yellow), and the muon chambers (red). The superconducting solenoid (dark blue) encloses all subdetectors except the muon chambers, which are located in the return yoke (gray).	27
4.1	Example of collinear and infrared unsafe jets. Collinear safety requires jet clustering to be independent of splitting input constituent. In the collinear unsafe case, the algorithm does not reconstruct a jet if the momenta of the split constituents are below the minimum threshold. Infrared safety prevents jet clustering result to change when adding soft object. Illustration inspired by [29].	35

4.2	An ideal event topology balances a precisely measured reference object with a jet in the transverse plane. Assuming that the transverse momentum is conserved in the event allows for calibrating the latter using a balancing method.	37
4.3	The factorized approach to calibrate the jet energy scale extracts correction factors in multiple stages. Each stage is either applied to MC simulated or recorded data. Illustration inspired from [47].	39
5.1	Balancing a precisely measurable reference object, in this case, a Z boson, with a jet allows calibrating the latter using the p_T -balance or MPF method. To exploit the strength of the CMS detector, this analysis concentrates on Z bosons reconstructed from two muons.	50
5.2	The workflow for calibrating the jet energy scale and momentum resolution is processed in different steps. While the amount of data is reduced during each step, the number of times each step is processed increases.	52
5.3	A shape comparison of the aMC@NLO MC simulation and the 2018 data set for p_T and η of the positively charged muon. MC simulation and data distributions are normalized to each other for shape comparisons.	60
5.4	A shape comparison of the aMC@NLO MC simulation and the 2018 data set for the reconstructed mass of the Z boson. MC simulation and data distributions are normalized to each other for shape comparison.	60
5.5	A shape comparison of the aMC@NLO MC simulation and the 2018 data set for p_T and η of the leading jet. The distributions of MC simulation and data are normalized to each other for comparing the shapes.	61
5.6	A shape comparison of the aMC@NLO MC simulation and the 2018 data set for p_T and η of the second-leading jet. The distributions of MC simulation and data are normalized to each other for comparing the shapes.	62
5.7	A shape comparison of the aMC@NLO MC simulation and the 2018 data set for E_T^{miss} . The distributions of MC simulation and data are normalized to each other for comparing the shapes.	63
5.8	A shape comparison of the aMC@NLO MC simulation and the 2018 data set for p_T -balance and MPF. The distributions of MC simulation and data are normalized to each other for comparing the shapes.	65
5.9	A shape comparison of the aMC@NLO MC simulation and the 2018 data set for the number of jets per event. The distributions of MC simulation and data are normalized to each other for comparing the shapes.	66

5.10	Mean of the balancing method versus p_{T}^{Z} shown for the MadGraph MC simulation and the 2018 data taking periods.	68
5.11	Mean of the balancing method versus $ \eta^{\text{jet1}} $ shown for the MadGraph MC simulation and the 2018 data taking periods.	69
6.1	Linear extrapolation of the mean values of both balancing methods, p_{T} -balance and MPF, comparing data and MC simulation. The ratio at $\alpha = 0$ corresponds to the absolute residual correction factor, which corrects for remaining differences between the MC and data distributions. While this plot exemplarily shows inclusive results, the official correction factors are binned in $ \eta^{\text{jet1}} $ and p_{T}^{Z}	72
6.2	Combination of p_{T}^{Z} dependent absolute residual correction factors for the jet energy scale done via a global fit. The results of the $\text{Z} + \text{jets}$ channel decaying into muons and electrons were pre-combined into $\text{Z} + \text{jet}$. While the solid black line indicates the extracted correction factors, the yellow band shows the absolute jet energy scale uncertainty. The graph is taken from [51].	74
6.3	Linear extrapolation of the mean values of both balancing methods, p_{T} -balance and MPF, comparing data and MC simulation. The absolute residual corrections derived within the global fit are applied and the closure between MC simulation and data at $\alpha = 0$ validates the results.	75
6.4	Mean of the p_{T} -balancing method versus p_{T}^{Z} and $ \eta^{\text{jet1}} $ shown for the MadGraph MC simulation and the 2018 data taking periods. The absolute residual corrections derived within the global fit are applied and MC simulation and data closes dramatically better compared to Figures 5.10 and 5.11. Iterating the analysis procedure and gradually eliminating distorting influences allows the closure to be further improved. Systematic uncertainties, which are likely to be in the order of the statistical uncertainties, may cover the remaining differences.	76
7.1	Response distributions for the aMC@NLO MC simulation and the data of 2018 used for extracting their variance and calculate the jet momentum resolution. As example, the distributions are shown inclusively in p_{T}^{Z} for the bin $0.25 < \alpha < 0.3$, $ \eta^{\text{jet1}} < 0.522$. We observe non-Gaussian tails that originate from FSR in all shown distributions. Additionally, the p_{T} -balance and PLI distributions show an excess at the left side of the peak caused by the jet splitting, pile-up contribution, and a cut-off at low values due to the applied selection criteria. Their influence on the variance extraction is reduced by choosing $[\text{Max} - \text{RMS}_{\text{corr}}, \text{Max} + 2 \cdot \text{RMS}_{\text{corr}}]$ as the range for fitting a Gaussian function.	79

7.2	Generated jet momentum response distribution from the aMC@NLO MC simulation used for comparison with the extracted jet momentum resolution. As example, the distributions are shown inclusively in p_{T}^{Z} for the bin $0.25 < \alpha < 0.3$, $ \eta^{\text{jet1}} < 0.522$. We observe non-Gaussian tails that originate from FSR. Their influence on the variance extraction is reduced by choosing $[\text{Max} - \text{RMS}_{\text{corr}}, \text{Max} + 2 \cdot \text{RMS}_{\text{corr}}]$ as the range for fitting a Gaussian function.	80
7.3	Illustration of a two-sided truncation of a Gaussian distribution with the mean μ and the variance σ . a marks the lower truncation threshold, b the upper one, respectively.	81
7.4	As an example, the dependency of the extracted variances that go into the input response distributions versus the second jet activity is shown inclusively in p_{T}^{Z} for the bin $ \eta^{\text{jet1}} < 0.522$	83
7.5	As an example, the dependency of the generated and extracted jet momentum resolution on second-jet activity is shown inclusively in p_{T}^{Z} for the bin $ \eta^{\text{jet1}} < 0.522$	84
7.6	The closure of the MC scale factor with 1.0 validates the method. We, therefore, study its dependency on $ \eta^{\text{jet1}} $	85
7.7	Full 2-dimensional MC closure test simultaneously binned in $ \eta^{\text{jet1}} $ and p_{T}^{Z} allows searching for detector regions that provide a p_{T} dependent jet momentum resolution. Therefore, the pull of the MC scale factor from 1.0 is calculated, which shows whether the deviations are within the statistical uncertainties.	86
7.8	As an example, the dependency of the jet momentum resolution extracted from data and MC simulation on second-jet activity is shown inclusively in p_{T}^{Z} for the bin $ \eta^{\text{jet1}} < 0.522$	87
7.9	The dependency of the jet momentum resolution scale factor on $ \eta^{\text{jet1}} $	87
7.10	Illustration of 2-dimensional jet momentum resolution scale factors simultaneously binned in $ \eta^{\text{jet1}} $ and p_{T}^{Z} . This allows searching for detector regions that provide a p_{T} dependent jet momentum resolution.	88
7.11	Comparison of the jet momentum resolution scale factors derived from $\text{Z} + \text{jets}$ events as presented in this thesis and the official results from the dijets channel. Both are binned in $ \eta^{\text{jet1}} $. The jet momentum resolution scale factors derived from $\text{Z} + \text{jets}$ events are based on the full 2018 data set. For comparison, the official jet momentum resolution scale factors that are based on the first half of the CMS data set of 2018 (Run A,B, and C) are used. A comparison with the second half of the data set (Run D) can be found in the appendix in Figure A.12. However, the official scale factors of both parts of the data set show similar trends. Differences are covered by the uncertainties.	89

8.1	Estimated demand for storage and computing resources during LHC run 2 and the HL-LHC run [75] required by the CMS experiment. The required CPU hours are given in terms of HS06, a HEP specific performance measure [76]. The amount of required resources will increase dramatically in 2026 while preparing for HL-LHC. Although computing and storage technologies are constantly improving over time, this development does not cover such a strong increase in demand.	92
8.2	The WLCG is organized in a tiered structure. While the Tier 0 is responsible for collecting and storing raw data, the Tier 1 computing centers focus on providing computing capacities for event reconstruction and Grid storage for distribution of data to the end-users. The Tier 2 computing centers provide storage and CPU capacities for MC simulation and analyses of end-users. Outside the strict WLCG structure, the main part for end-user analysis processing is provided by institutes, summarized in the so-called Tier 3 layer.	93
9.1	XROOTD management server allows the connection of multiple data servers within a hierarchical structure. This enables the creation of a global namespace for all data stored on the data servers. Analogously, XROOTD allows designing a highly scalable caching infrastructure that combines multiple caching proxy servers. Each management server is capable of handling up to 64 data servers, caching proxy servers, or further management servers. Data accesses and metadata requests to any management server are redirected to its subordinate data servers or caching proxy servers.	108
9.2	The NAVIX coordination service connects an XROOTD based caching infrastructure with the HTCONDOR batch system to introduce data locality into job scheduling. By this, it influences the placement of data in multiple caches avoiding replication of data and matching jobs to suitable worker-nodes that already provide cached input data. The graph was previously published in [117–119].	110
9.3	XROOTD and HTCONDOR already support hierarchical scaling. This allows to combine multiple local pools of computing resources into regional superimposed systems. Several regional resource pools are merged into a national one. NAVIX considers only the data locality in the granularity of the direct subsystems, i.e. the computing resources one level down. This allows for scaling the complete approach easily.	113

9.4	Small-scale testing setup with three worker-nodes connected to an HTCONDOR batch system. Each of them is equipped with a SSD-cache via its own XROOTD caching proxy server. Jobs are matched to the worker-node that already provides cached files via the coordination service NAVIX.	114
9.5	While a step-wise increasing fraction of cached files in multiple benchmark runs, the duplicity of files in all three caches is monitored. Each bar in the plot monitors a time window that corresponds to one step of the benchmark that covers a certain fraction of cached files. In each step, the number of files accessed and their duplicity in the caches is measured. Since there is no file cached more than once, the coordination of jobs to already cached data appears to succeed. Otherwise, processing jobs on different worker-nodes would increase file duplicity. The graph was first published in [119].	116
9.6	Within the performance benchmark, the data transfer rate and the processing time of the jobs are monitored while step-wise increasing the fraction of already cached files. Using Equation 9.1 allows us to predict the influence of the cached fraction of files on the processing time. Comparing prediction and measurement shows quite a good agreement. In both graphs shown, the error bars indicate the variance of the measured set of jobs. Graph (b) was first published in [117].	117
9.7	Submitting the selection step of the jet energy calibration workflow to the testing system allows us to estimate the benefits of distributed coordinated caching for typical data-intensive HEP analyses. Therefore, the CPU efficiency and the processing time of the jobs is measured per given fraction of cached files. The prediction of the processing time is calculated according to Equation 9.1. In both graphs shown, the error bars indicate the variance of the measured set of jobs. Graphs were first published in [119].	119
9.8	Caching structure that is applicable for shared resources such as HPC clusters or cloud resources, where the worker-nodes do not provide a suitable storage volume for caching purposes. Instead, a dedicated caching proxy server that is connected to a suitable cache volume serves as a global cache for all worker-nodes within the computing cluster.	121

9.9	A maximum performance benchmark was executed on the HPC center NEMO and at OTC resources. The average processing time of the jobs was measured while step-wise increasing the fraction of cached files. The benchmarks are repeated after a week to estimate the time dependency of the results. All results confirm a fluctuating performance of the shared computing resources. In both graphs shown, the error bars indicate the variance of the measured set of jobs. Parts of the measurements were first published in [119].	123
9.10	Design of the TOPAS cluster dedicated to high-throughput data analysis. Two cache layers allow us to boost data throughput and reduce the load on Grid storage and network resources. A distributed file system spread across all 11 worker-nodes serves as 1 PB cache. Additionally, one high-performance NVMe SSD per worker-node provides a fast accessible 1 TB cache volume. Both caches can be accessed via XROOTD caching proxy servers internally within the worker-node. .	124
9.11	Analysis-like benchmarks that apply a basic selection on a data sample were performed to determine the maximum performance when accessing the TOPAS cache volumes. The NVMe SSDs performed slightly better than the distributed file system CephFS. However, both show a relatively stable, high data transfer rate up to 560 parallel processes, which corresponds to the number of logical cores delivered by the setup. For the CephFS benchmarks, a classical network setup was used without the modern optimization RoCE that allows remote direct memory access via network. The average CPU usage of the performance benchmark represents its CPU efficiency since no advanced calculations are performed with data. The performance benchmarks applied to a set of 10 TOPAS worker-nodes were published in [128].	125
9.12	Submitting the selection step of the jet energy calibration workflow to the testing system allows us to estimate the benefits of distributed coordinated caching for typical data-intensive HEP analyses. Therefore, CPU efficiency and the processing time of the jobs is measured for certain fractions of cached files. The prediction of the processing time is calculated according to Equation 9.1. Graphs were first published in [119].	126
9.13	Distribution of jet energy calibration jobs, shown as a diagram of data throughput versus CPU efficiency for two scenarios, accessing data from remote Grid storage at the top, (a) and cached on CephFS at the bottom, (b). Processing recorded data and simulated data resulted in two different clusters of points. A dependency between data throughput and CPU efficiency reflects a limitation of the workflow caused by data access.	129

10.1	Jet momentum resolution scale factors derived with an advanced extraction procedure. These are extracted depending on $ \eta^{\text{jet1}} $ due to heterogeneous detector layers in the longitudinal direction.	132
A.1	A shape comparison of the MadGraph MC simulation and the 2018 data set for p_T and η of the positively charged muon. MC simulation and data distributions are normalized to each other for shape comparisons.	140
A.2	A shape comparison of the MadGraph MC simulation and the 2018 data set for the reconstructed mass of the Z boson. MC simulation and data distributions are normalized to each other for shape comparison.	141
A.3	A shape comparison of the MadGraph MC simulation and the 2018 data set for p_T and η of the leading jet. The distributions of MC simulation and data are normalized to each other for comparing the shapes.	141
A.4	A shape comparison of the MadGraph MC simulation and the 2018 data set for p_T and η of the second-leading jet. The distributions of MC simulation and data are normalized to each other for comparing the shapes.	142
A.5	A shape comparison of the MadGraph MC simulation and the 2018 data set for E_T^{miss} . The distributions of MC simulation and data are normalized to each other for comparing the shapes.	142
A.6	A shape comparison of the MadGraph MC simulation and the 2018 data set for p_T -balance and MPF. The distributions of MC simulation and data are normalized to each other for comparing the shapes. . .	143
A.7	A shape comparison of the MadGraph MC simulation and the 2018 data set for the number of jets per event. The distributions of MC simulation and data are normalized to each other for comparing the shapes.	143
A.8	Full 2-dimensional MC closure test simultaneously binned in $ \eta^{\text{jet1}} $ and p_T^Z . Therefore, the pull of the MC scale factor from 1.0 is calculated, which shows whether the deviations are within the statistical uncertainties.	148
A.9	Uncertainties of the full 2-dimensional MC closure test simultaneously binned in $ \eta^{\text{jet1}} $ and p_T^Z . Therefore, the pull of the MC scale factor from 1.0 is calculated, which shows whether the deviations are within the statistical uncertainties.	149
A.10	Illustration of the uncertainties for the 2-dimensional jet momentum resolution scale factors simultaneously binned in $ \eta^{\text{jet1}} $ and p_T^Z	150
A.11	Illustration of the pull for the 2-dimensional jet momentum resolution scale factors simultaneously binned in $ \eta^{\text{jet1}} $ and p_T^Z	151

A.12 Comparison of the jet momentum resolution scale factors derived from Z + jets events as presented in this thesis and the official results from the dijets channel. Both are binned in $ \eta^{\text{jet1}} $. The jet momentum resolution scale factors derived from Z+jets events are based on the full 2018 data set. For comparison, the official jet momentum resolution scale factors that are based on the second half of the CMS data set of 2018 (Run D) are used. A comparison with the first half of the data set (Run A, B, and C) can be found in Figure 7.11. However, the official scale factors of both parts of the data set show similar trends. Differences are covered by the uncertainties.	151
A.13 Flexible design for a general purpose coordination service to realize distributed coordinated caching on multiple computing infrastructures. The batch system as well as the caching mechanism can be chosen freely.	152

List of Tables

5.1	Overview of the data reduction, the average processing time, and the frequency of reprocessing the analysis steps. All values are estimated for a single run of the jet momentum resolution analysis for the muon channel. This includes the recorded double-muon data set and the aMC@NLO MC sample mentioned in section 5.1. The frequency is roughly estimated and depends highly on the activity within the CMS collaboration.	52
A.1	Data samples recorded by the CMS detector preselected by double-muon trigger	136
A.2	MC simulated Z + jets data samples binned in number of jets used for extracting jet momentum resolution	136
A.3	MC simulated Z + jets data samples binned in number of jets and p_T^Z used for extracting jet momentum resolution	137
A.4	Restrictions for tight Jet ID with lepton veto for CMS data of 2018.	139
A.5	List of bins that are removed from the analysis due to not-recoverable fits	144
A.6	List of bins (part 1), where the fit range was manually tuned to be able to extract valid results	145
A.7	List of bins (part 2), where the fit range was manually tuned to be able to extract valid results	146
A.8	List of bins (part 3), where the fit range was manually tuned to be able to extract valid results	147

Bibliography

- [1] D. Galbraith and C. Burgard, *Standard model of physics*, 2020, URL: <http://www.texample.net/tikz/examples/model-physics/> (visited on 04/06/2020).
- [2] B.R. Webber, “Fragmentation and hadronization”, *eConf C990809* (2000) 577–606, DOI: 10.1142/S0217751X00005334, [arXiv:hep-ph/9912292].
- [3] K. S. Grogg, “Jets produced in association with W-bosons in CMS at the LHC”, 2011, URL: <https://cds.cern.ch/record/1376067>.
- [4] K. Rabbertz, *Jet Physics at the LHC: The Strong Force beyond the TeV Scale*, vol. 268, Springer Tracts in Modern Physics, Springer, 2017, DOI: 10.1007/978-3-319-42115-5.
- [5] B. Andersson et al., “Parton fragmentation and string dynamics”, *Physics Reports* **97** (1983) 31–145, DOI: [https://doi.org/10.1016/0370-1573\(83\)90080-7](https://doi.org/10.1016/0370-1573(83)90080-7).
- [6] B. Andersson, G. Gustafson and B. Söderberg, “A general model for jet fragmentation”, *Zeitschrift für Physik C Particles and Fields* **20** (1983) 317–329, DOI: 10.1007/BF01407824.
- [7] T. Sjöstrand, “The merging of jets”, *Physics Letters B* **142** (1984) 420–424, DOI: [https://doi.org/10.1016/0370-2693\(84\)91354-6](https://doi.org/10.1016/0370-2693(84)91354-6).
- [8] G. Arnison et al., “Experimental Observation of Lepton Pairs of Invariant Mass Around 95-GeV/c**2 at the CERN SPS Collider”, *Phys. Lett. B* **126** (1983) 398–410, DOI: 10.1016/0370-2693(83)90188-0.
- [9] L. Di Lella and C. Rubbia, “The Discovery of the W and Z Particles”, *Adv. Ser. Dir. High Energy Phys.* **23** (2015) 137–163, DOI: 10.1142/9789814644150_0006.
- [10] M. Tanabashi et al., “Review of Particle Physics”, *Phys. Rev. D* **98** (2018) p. 030001, DOI: 10.1103/PhysRevD.98.030001.
- [11] The CMS collaboration, *Measurements of inclusive and differential Z boson production cross sections in pp collisions at $\sqrt{s} = 13$ TeV*, tech. rep., 2016, URL: <https://cds.cern.ch/record/2140105>.

- [12] R. D. Ball et al., “Parton distributions for the LHC run II”, *Journal of High Energy Physics* (2015), DOI: 10.1007/jhep04(2015)040.
- [13] A. Hermann et al., *History of CERN, I: Volume I - Launching the European Organization for Nuclear Research*, Elsevier Science, 1987, ISBN: 9780444870377.
- [14] C. Lefèvre, “The CERN accelerator complex. Complexe des accélérateurs du CERN”, 2008, URL: <http://cds.cern.ch/record/1260465>.
- [15] L. Evans and P. Bryant, “LHC Machine”, *Journal of Instrumentation* **3** (2008) S08001–S08001, DOI: 10.1088/1748-0221/3/08/s08001.
- [16] The CMS Collaboration et al., “The CMS experiment at the CERN LHC”, *Journal of Instrumentation* **3** (2008) S08004–S08004, DOI: 10.1088/1748-0221/3/08/s08004.
- [17] The ATLAS Collaboration et al., “The ATLAS Experiment at the CERN Large Hadron Collider”, *Journal of Instrumentation* **3** (2008) S08003–S08003, DOI: 10.1088/1748-0221/3/08/s08003.
- [18] S. Chatrchyan et al., “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”, *Phys. Lett. B* **716** (2012) 30–61, DOI: 10.1016/j.physletb.2012.08.021, [arXiv:1207.7235].
- [19] G. Aad et al., “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, *Phys. Lett. B* **716** (2012) 1–29, DOI: 10.1016/j.physletb.2012.08.020, [arXiv:1207.7214].
- [20] The LHCb Collaboration et al., “The LHCb Detector at the LHC”, *Journal of Instrumentation* **3** (2008) S08005–S08005, DOI: 10.1088/1748-0221/3/08/s08005.
- [21] The ALICE Collaboration et al., “The ALICE experiment at the CERN LHC”, *Journal of Instrumentation* **3** (2008) S08002–S08002, DOI: 10.1088/1748-0221/3/08/s08002.
- [22] M. Ressegotti, *Overview of the CMS detector performance at LHC Run 2*, tech. rep., 2018, DOI: 10.3390/universe5010018.
- [23] P. Lujan and on behalf of the CMS collaboration, *CMS Luminosity - Public Results*, 2019, URL: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/LumiPublicResults> (visited on 28/02/2020).
- [24] G. Apollinari et al., *High Luminosity Large Hadron Collider HL-LHC*, 2015, DOI: 10.5170/CERN-2015-005.1, [arXiv:1705.08830].
- [25] L. Taylor, *CMS detector design*, 2011, URL: <http://cms.web.cern.ch/news/cms-detector-design> (visited on 15/05/2020).
- [26] The CMS collaboration, *About CMS*, 2020, URL: <https://cms.cern/detector> (visited on 08/06/2020).

-
- [27] D. Contardo et al., *Technical Proposal for the Phase-II Upgrade of the CMS Detector*, tech. rep., 2015, URL: <https://cds.cern.ch/record/2020886>.
- [28] D. Haitz, “Precision Studies of Proton Structure and Jet Energy Scale with the CMS Detector at the LHC”, PhD thesis, Karlsruhe Institut für Technologie (KIT), 2016, DOI: 10.5445/IR/1000055680.
- [29] J. Berger, “Search for the Higgs Boson Produced via Vector-Boson Fusion in the Decay Channel H to Tau Tau”, PhD thesis, Karlsruhe Institut für Technologie (KIT), 2014, DOI: 10.5445/IR/1000042221.
- [30] A. Dominguez et al., *CMS Technical Design Report for the Pixel Detector Upgrade*, tech. rep., 2012, URL: <https://cds.cern.ch/record/1481838>.
- [31] S. Chatrchyan et al., “Description and performance of track and primary-vertex reconstruction with the CMS tracker”, *JINST* **9** (2014) P10009, DOI: 10.1088/1748-0221/9/10/P10009, [arXiv:1405.6569].
- [32] S. Chatrchyan et al., “Performance and Operation of the CMS Electromagnetic Calorimeter”, *JINST* **5** (2009) T03010. 39 p, DOI: 10.1088/1748-0221/5/03/T03010.
- [33] T. Dimova, *Calibration and performance of the CMS electromagnetic calorimeter during the LHC Run II*, tech. rep., 2018, DOI: 10.1007/978-3-030-21970-3_22.
- [34] T. Mudholkar and on behalf of the CMS collaboration, “CMS ECAL monitoring and calibration in LHC Run II”, *Journal of Physics: Conference Series* **1162** (2019) p. 012002, DOI: 10.1088/1742-6596/1162/1/012002.
- [35] The CMS collaboration, “The CMS hadron calorimeter project: Technical Design Report” (1997), URL: <https://cds.cern.ch/record/357153>.
- [36] S. I. Cooper, “Phase I Upgrade of the CMS Hadron Calorimeter”, *Nuclear and Particle Physics Proceedings* **273-275** (2016) 1002–1007, DOI: <https://doi.org/10.1016/j.nuclphysbps.2015.09.157>.
- [37] M. Chadeeva and N. Lychkovskaya, “Calibration of the CMS hadron calorimeter in Run 2”, *Journal of Instrumentation* **13** (2018) C03025–C03025, DOI: 10.1088/1748-0221/13/03/c03025.
- [38] L. Guiducci, *CMS muon system towards LHC Run 2 and beyond*, tech. rep., 2014, DOI: 10.1016/j.nuclphysbps.2015.09.159.
- [39] N. Pozzobon, *The CMS Muon Systemperformance during the LHC Run-2*, tech. rep., 2019, DOI: 10.1088/1748-0221/14/11/C11031.
- [40] J. T. Boyd, *LHC Run-2 and Future Prospects*, 2020, [arXiv:2001.04370].

- [41] M. Felcini, “The Trigger System of the CMS Experiment”, *Nucl. Instrum. Meth. A* **598** (2009) 312–316, DOI: 10.1016/j.nima.2008.08.126, [arXiv:0806.2540].
- [42] T. Bawej et al., “The New CMS DAQ System for Run-2 of the LHC”, *IEEE Trans. Nucl. Sci.* **62** (2015) 1099–1103, DOI: 10.1109/RTC.2014.7097437.
- [43] S. Chatrchyan et al., “Determination of Jet Energy Calibration and Transverse Momentum Resolution in CMS”, *JINST* **6** (2011) P11002, DOI: 10.1088/1748-0221/6/11/P11002, [arXiv:1107.4277].
- [44] A.M. Sirunyan et al., “Particle-flow reconstruction and global event description with the CMS detector”, *JINST* **12** (2017) P10003, DOI: 10.1088/1748-0221/12/10/P10003, [arXiv:1706.04965].
- [45] M. Cacciari, G. P. Salam and G. Soyez, “The anti-ktjet clustering algorithm”, *Journal of High Energy Physics* (2008) 063–063, DOI: 10.1088/1126-6708/2008/04/063.
- [46] The CMS collaboration, *Pileup Removal Algorithms*, tech. rep., 2014, URL: <https://cds.cern.ch/record/1751454>.
- [47] V. Khachatryan et al., “Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV”, *JINST* **12** (2017) P02014, DOI: 10.1088/1748-0221/12/02/P02014, [arXiv:1607.03663].
- [48] D. Bertolini et al., “Pileup Per Particle Identification”, *JHEP* **10** (2014) p. 059, DOI: 10.1007/JHEP10(2014)059, [arXiv:1407.6013].
- [49] The CMS collaboration, *Jet Performance in pp Collisions at 7 TeV*, tech. rep., 2010, URL: <http://cds.cern.ch/record/1279362>.
- [50] The CMS collaboration, “Jet energy scale and resolution performance with 13 TeV data collected by CMS in 2016” (2018), URL: <http://cds.cern.ch/record/2622157>.
- [51] The CMS collaboration, “Jet energy scale and resolution performance with 13 TeV data collected by CMS in 2016-2018” (2020), URL: <https://cds.cern.ch/record/2715872>.
- [52] The CMS collaboration, *Jet Energy Resolution in CMS at $\sqrt{s} = 7$ TeV*, tech. rep., 2011, URL: <https://cds.cern.ch/record/1339945>.
- [53] J. Alwall et al., “MadGraph 5: going beyond”, *Journal of High Energy Physics* (2011), DOI: 10.1007/jhep06(2011)128.
- [54] T. Sjöstrand et al., “An Introduction to PYTHIA 8.2”, *Comput. Phys. Commun.* **191** (2015) 159–177, DOI: 10.1016/j.cpc.2015.01.024, [arXiv:1410.3012].

-
- [55] J. Alwall et al., “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”, *Journal of High Energy Physics* (2014), DOI: 10.1007/jhep07(2014)079.
- [56] V. Khachatryan et al., “The CMS trigger system”, *Journal of Instrumentation* **12** (2017) P01020–P01020, DOI: 10.1088/1748-0221/12/01/p01020.
- [57] L. Cadamuro, “The CMS Level-1 trigger system for LHC Run II”, *Journal of Instrumentation* **12** (2017) C03021–C03021, DOI: 10.1088/1748-0221/12/03/c03021.
- [58] T. Virdee, A. Petrilli and A. Ball, *CMS High Level Trigger*, tech. rep., 2007, URL: <https://cds.cern.ch/record/1043242>.
- [59] F. Stober, “Measurement of the three-jet mass cross-section at $\sqrt{s}=7\text{TeV}$ ”, PhD thesis, KIT Karlsruhe, 2012, DOI: 10.5445/IR/1000034422.
- [60] KIT CMS group, *Karlsruhe Package for Physics Analysis*, 2020, URL: <https://github.com/KIT-CMS/Kappa> (visited on 20/05/2020).
- [61] KIT CMS group, *Excalibur*, 2020, URL: <https://github.com/KIT-CMS/Excalibur> (visited on 20/05/2020).
- [62] J. Berger et al., *ARTUS - A Framework for Event-based Data Analysis in High Energy Physics*, 2015, [arXiv:1511.00852].
- [63] KIT CMS group, *Artus*, 2020, URL: <https://github.com/KIT-CMS/Artus> (visited on 20/05/2020).
- [64] D. Savoiu, *Karma*, 2020, URL: <https://github.com/dsavoIU/Karma> (visited on 20/05/2020).
- [65] A. K. Kalsi, *Collisions18 13 TeV data: Final JSON file - PromptReco 314472-325175*, 2018, URL: <https://hypernews.cern.ch/HyperNews/CMS/get/physics-validation/3259.html> (visited on 01/05/2020).
- [66] A. M. Sirunyan et al., “Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s}=13\text{TeV}$ ”, *Journal of Instrumentation* **13** (2018) P06015–P06015, DOI: 10.1088/1748-0221/13/06/p06015.
- [67] The CMS collaboration, “Muon identification and isolation efficiencies with 2017 and 2018 data” (2018), URL: <http://cds.cern.ch/record/2629364>.
- [68] The CMS collaboration, “Muon reconstruction performance during Run II” (2019), URL: <https://cds.cern.ch/record/2682902>.
- [69] M. Diamantopoulou et al., *Jet Identification for the 13 TeV data Run2018*, 2020, URL: <https://twiki.cern.ch/twiki/bin/view/CMS/JetID13TeVRun2018%7D> (visited on 04/03/2020).

- [70] M. Diamantopoulou, N. Saoulidou and E. Tziaferi, “Performance of the Particle-Flow jet identification criteria using proton-proton collisions at 13 TeV” (2016), URL: <http://cms.cern.ch/iCMS/user/noteinfo?cmsnoteid=CMS%20AN-2015/269>.
- [71] H. Lattaud and V. Sordini, “Jet transverse momentum resolution measurement using photon+jet events at 13 TeV” (2018), URL: <http://cms.cern.ch/iCMS/user/noteinfo?cmsnoteid=CMS%20AN-2017/322>.
- [72] N. Balakrishna and C. D. Lai, “Univariate Distributions”, *Continuous Bivariate Distributions: Second Edition*, Springer New York, 2009, 1–32, DOI: 10.1007/b101765_1.
- [73] N. L. Johnson, S. Kotz and N. Balakrishnan, “Wiley series in probability and mathematical statistics: Applied probability and statistics: Truncated Normal Distributions”, *Continuous Univariate Distribution*, vol. 1, Wiley & Sons, 1994, 156–162, ISBN: 978-0-471-58495-7.
- [74] K. Bos et al., “Technical Design Report LCG: LHC computing Grid: Technical Design Report. Version 1.06 (20 Jun 2005)”, **1** (2005) 156–162, URL: <https://cds.cern.ch/record/840543>.
- [75] J. Albrecht et al., “A Roadmap for HEP Software and Computing R&D for the 2020s”, *Computing and Software for Big Science* **3** (2019) 156–162, DOI: 10.1007/s41781-018-0018-8.
- [76] HEPIX Benchmarking Working Group, *HEP-SPEC06 (HS06)*, 2020, URL: <https://w3.hepik.org/benchmarking.html> (visited on 02/06/2020).
- [77] D. Lange et al., “CMS Computing Resources: Meeting the demands of the high-luminosity LHC physics program”, *EPJ Web Conf.* **214** (2019) 03055. 7 p, DOI: 10.1051/epjconf/201921403055.
- [78] I. Bird et al., *Update of the Computing Models of the WLCG and the LHC Experiments*, tech. rep., 2014, 156–162, URL: <https://cds.cern.ch/record/1695401>.
- [79] I. Bird, “Computing for the Large Hadron Collider”, *Annual Review of Nuclear and Particle Science* **61** (2011) 99–118, DOI: 10.1146/annurev-nucl-102010-130059, eprint: <https://doi.org/10.1146/annurev-nucl-102010-130059>.
- [80] R. Brun, F. Carminati and G.G. Carminati, *The Frontiers Collection: From the Web to the Grid and Beyond: Computing Paradigms Driven by High-Energy Physics*, vol. 1, Springer Berlin Heidelberg, 2012, 156–162, DOI: 10.1007/978-3-642-23157-5.
- [81] Andreas Haupt and Yves Kemp, “The NAF: National Analysis Facility at DESY”, *J. Phys. Conf. Ser.* **219** (2010) p. 052007, DOI: 10.1088/1742-6596/219/5/052007.

-
- [82] K. Anampa et al., “Deploying and extending CMS Tier 3s using VC3 and the OSG Hosted CE service”, *EPJ Web of Conferences* **214** (2019) p. 03035, DOI: 10.1051/epjconf/201921403035.
- [83] C. Heidecker et al., “Dynamic Resource Extension for Data Intensive Computing with Specialized Software Environments on HPC Systems”, *Proceedings of the 5th bwHPC Symposium HPC Activities in Baden-Württemberg, Freiburg - September 24/25 2018*. Ed.: M. Janczyk, vol. 1, 51.01.01; LK 01, Tübingen Library Publishing, Tübingen, 2019, 159–161, DOI: 10.15496/publikation-29051.
- [84] Computing Project Board der Helmholtz Allianz „Physics at the Terascale“, *Computing in der Hochenergiephysik in Deutschland: Bestandsaufnahme und Aussicht*, 2013, URL: https://www.ketweb.de/content/e199639/e223266/GridPB_HEP_Computing_211113.pdf (visited on 02/06/2020).
- [85] J. Dost et al., “Site in a box: Improving the Tier 3 experience”, *Journal of Physics: Conference Series* **898** (2017) p. 082050, DOI: 10.1088/1742-6596/898/8/082050.
- [86] ErUM-Data ITD collaboration, *Innovative Digital Technologies for Research on Universe and Matter*, 2020, URL: <https://www.erum-data-idt.de> (visited on 02/06/2020).
- [87] The HTCCondor team, *Computing with HTCCondor*, 2020, URL: <https://research.cs.wisc.edu/htcondor> (visited on 04/06/2020).
- [88] M. Litzkow, M. Livny and M. Mutka, “Condor - A Hunter of Idle Workstations”, *Proceedings of the 8th International Conference of Distributed Computing Systems*, vol. 8, Institute of Electrical and Electronics Engineers, 1988, 104–111, DOI: 10.1109/DCS.1988.12507.
- [89] D. Thain, T. Tannenbaum and M. Livny, “Distributed computing in practice: the Condor experience”, *Concurrency - Practice and Experience* **17** (2005) 323–356, DOI: 10.1002/cpe.938.
- [90] XRootD Collaboration, *XRootD*, 2020, URL: <https://xrootd.slac.stanford.edu> (visited on 04/06/2020).
- [91] A. Dorigo et al., “XROOTD/TXNetFile: A Highly Scalable Architecture for Data Access in the ROOT Environment”, *Proceedings of the 4th WSEAS International Conference on Telecommunications and Informatics*, vol. 1, World Scientific, Engineering Academy and Society (WSEAS), 2005, 156–162, ISBN: 960-8457-11-4.
- [92] J. Nabrzyski, J. Schopf and J. Węglarz, *Grid Resource Management: State of the Art and Future Trends*, vol. 64, World Scientific, Engineering Academy and Society (WSEAS), 2004, 156–162, DOI: 10.1007/978-1-4615-0509-9.

- [93] K. Meier et al., “Dynamic provisioning of a HEP computing infrastructure on a shared hybrid HPC system”, *J. Phys. Conf. Ser.* **762** (2016) p. 012012, DOI: 10.1088/1742-6596/762/1/012012.
- [94] M. J. Schnepf et al., “Dynamic Integration and Management of Opportunistic Resources for HEP”, *EPJ Web of Conferences* **214** (2019) p. 08009, DOI: 10.1051/epjconf/201921408009.
- [95] Martin Gasthuber, Helge Meinhard and Robert Jones, “HNSciCloud - Overview and technical challenges”, *J. Phys. Conf. Ser.* **898** (2017) p. 052040, DOI: 10.1088/1742-6596/898/5/052040.
- [96] João Fernandes et al., “HNSciCloud, a hybrid cloud for science”, *EPJ Web Conf.* **214** (2019) p. 09006, DOI: 10.1051/epjconf/201921409006.
- [97] The Apache Software Foundation, *The Hadoop Project*, 2020, URL: %5Curl%7Bhttps://hadoop.apache.org/%7D (visited on 15/05/2020).
- [98] Q. Huang et al., “Using Hadoop for High Energy Physics Data Analysis”, *Proceedings of the 4th WSEAS International Conference on Telecommunications and Informatics*, vol. 1, World Scientific, Engineering Academy and Society (WSEAS), 2019, 146–153, DOI: 10.1007/978-3-030-28061-1_16.
- [99] S. A. Russo, M. Pinamonti and M. Cobal, “Running a typical ROOT HEP analysis on Hadoop MapReduce”, *Journal of Physics: Conference Series* **513** (2014) p. 032080, DOI: 10.1088/1742-6596/513/3/032080.
- [100] W. Bhimji, T. Bristow and A. Washbrook, “HEPDOOP: High-Energy Physics Analysis using Hadoop”, *Journal of Physics: Conference Series* **513** (2014) p. 022004, DOI: 10.1088/1742-6596/513/2/022004.
- [101] D. Smith et al., *Sharing server nodes for storage and compute*, tech. rep., 2018, 156–162, DOI: 10.1051/epjconf/201921408025.
- [102] L. A. T. Bauerdick et al., “XRootd, disk-based, caching proxy for optimization of data access, data placement and data replication”, *Journal of Physics: Conference Series* **513** (2014) p. 042044, DOI: 10.1088/1742-6596/513/4/042044.
- [103] E. Fajardo et al., “A federated Xrootd cache”, *Journal of Physics: Conference Series* **1085** (2018) p. 032025, DOI: 10.1088/1742-6596/1085/3/032025.
- [104] T. Li, R. Currie and A. Washbrook, “A data caching model for Tier 2 WLCG computing centres using XCache”, *EPJ Web Conf.* **214** (2019) p. 04047, DOI: 10.1051/epjconf/201921404047.
- [105] S. C. Skipsey et al., “Caching technologies for Tier-2 sites: A UK perspective”, *EPJ Web Conf.* **214** (2019) p. 04002, DOI: 10.1051/epjconf/201921404002.

-
- [106] M. Fischer et al., “Tier 3 batch system data locality via managed caches”, *Journal of Physics: Conference Series* **608** (2015) p. 012018, DOI: 10.1088/1742-6596/608/1/012018.
- [107] M. Fischer et al., “High Performance Data Analysis via Coordinated Caches”, *Journal of Physics: Conference Series* **664** (2015) p. 092008, DOI: 10.1088/1742-6596/664/9/092008.
- [108] M. Fischer et al., “Data Locality via Coordinated Caching for Distributed Processing”, *Journal of Physics: Conference Series* **762** (2016) p. 012011, DOI: 10.1088/1742-6596/762/1/012011.
- [109] M. Fischer, “Coordinated Caching for High Performance Calibration using $Z \rightarrow \mu\mu$ Events of the CMS Experiment”, PhD thesis, KIT, Karlsruhe, 2016, 156–162, DOI: 10.5445/IR/1000067354.
- [110] M. Fischer et al., “Opportunistic data locality for end user data analysis”, *Journal of Physics: Conference Series* **898** (2017) p. 052034, DOI: 10.1088/1742-6596/898/5/052034.
- [111] C. Heidecker, *Absolute jet energy scale determination at the CMS detector for LHC run 2 at $\sqrt{s} = 13$ TeV using an optimized processing setup*, 2017, URL: <https://publish.etp.kit.edu/record/21318>.
- [112] C. Heidecker et al., “Provisioning of data locality for HEP analysis workflows”, *Journal of Physics: Conference Series* **1085** (2018) p. 032005, DOI: 10.1088/1742-6596/1085/3/032005.
- [113] T. Feßenbecker, *Modeling of distributed coordinated caching for LHC data analysis*, 2020.
- [114] E. Kuehn et al., *tfesenbecker/lapis: Initial release*, version v0.1, 2020, DOI: 10.5281/zenodo.3819501.
- [115] E. Kuehn et al., *MatterMiners/lapis: v0.4.0*, version v0.4.0, 2020, DOI: 10.5281/zenodo.3822379.
- [116] R. Brun and F. Rademakers, “ROOT: An object oriented data analysis framework”, *Nucl. Instrum. Meth. A* **389** (1997) 81–86, DOI: 10.1016/S0168-9002(97)00048-X.
- [117] R. Caspart et al., “Advancing throughput of HEP analysis work-flows using caching concepts”, *EPJ Web Conf.* **214** (2019) p. 04007, DOI: 10.1051/epjconf/201921404007.
- [118] C. Heidecker et al., “Boosting Performance of Data-intensive Analysis Work-flows with Distributed Coordinated Caching”, *Journal of Physics: Conference Series* **1** (to be published) 156–162.

- [119] M. Sauter, *Increasing efficiency of HEP workflows by coordinated distributed caching*, 2019, URL: <https://publish.etp.kit.edu/record/21966>.
- [120] D. Bradley et al., “An update on the scalability limits of the Condor batch system”, *Journal of Physics: Conference Series* **331** (2011) p. 062002, DOI: 10.1088/1742-6596/331/6/062002.
- [121] C. Heidecker and M. Sauter, *NaviX*, 2020, URL: <https://gitlab.etp.kit.edu/ETP-HTC/Navix> (visited on 04/06/2020).
- [122] R. F. von Cube et al., “Federation of compute resources available to the German CMS community”, *Journal of Physics: Conference Series* **1** (to be published) 156–162.
- [123] Steinbuch Centre for Computing, Karlsruhe Institute of Technology, *Informationen über GridKa*, 2020, URL: <http://www.gridka.de> (visited on 13/06/2020).
- [124] H. Marten, K.-P. Mickel and R. Kupsch, *Project Grid Computing: A Grid Computing Centre at Forschungszentrum Karlsruhe - Response on the Requirements for a Regional Data and Computing Centre in Germany*, 2001, URL: <http://www.gridka.de/RDCCG-answer-v8.pdf> (visited on 13/06/2020).
- [125] M. J. Schnepf et al., “HEP Analyses on Dynamically Allocated Opportunistic Computing Resources”, *Journal of Physics: Conference Series* **1** (to be published) 156–162.
- [126] M. J. Schnepf et al., “Dynamic Integration and Management of Opportunistic Resources for HEP”, *EPJ Web Conf.* **214** (2019) p. 08009, DOI: 10.1051/epjconf/201921408009.
- [127] M. J. Schnepf et al., “Mastering Opportunistic Computing Resources for HEP”, *Journal of Physics: Conference Series* **1085** (2018) p. 032056, DOI: 10.1088/1742-6596/1085/3/032056.
- [128] R. Caspart et al., “Setup and commissioning of a high-throughput analysis cluster”, *EPJ Web Conf.* **1** (to be published) 156–162.
- [129] CMS Physics Validation, *Golden JSON Certificate for run periods 314472-325175 at 13TeV for PromptReco Collisions 2018*, 2018, URL: https://cms-service-dqm.web.cern.ch/cms-service-dqm/CAF/certification/Collisions18/13TeV/PromptReco/Cert_314472-325175_13TeV_PromptReco_Collisions18_JSON.txt (visited on 01/05/2020).

Acronyms

Kappa KARLSRUHE PACKAGE FOR PHYSICS ANALYSES.

p_T^{miss} missing transverse momentum.

ALICE A Large Ion Collider Experiment.

ATLAS A Toroidal LHC ApparatuS.

BEER Batch on EOS Extra Resources.

BOOSTER Proton Synchrotron Booster.

CALO jet Calorimeter jet.

CephFS Ceph file system.

CERN Conseil européen pour la recherche nucléaire.

CHS Charged Hadron Subtraction.

CMS Compact Muon Solenoid.

DAQ Data Acquisition System.

DESY Deutsches Elektronen-Synchrotron.

ECAL electromagnetic calorimeter.

ETP Institute of Experimental Particle Physics.

FSR Final State Radiation.

HCAL hadronic calorimeter.

HDD Hard Disk Drive.

- HEP** High Energy Physics.
- HL-LHC** High-Luminosity LHC.
- HLT** High-Level Trigger.
- HNSciCloud** Helix Nebula Science Cloud.
- HPC** High-Performance Computing.
- ISR** Initial State Radiation.
- JPT jet** Jet-Plus-Track jet.
- KIT** Karlsruhe Institute of Technology.
- KSETA** Karlsruhe School of Elementary Particle and Astroparticle Physics:
Science and Technology.
- L1 Trigger** Level-1 Trigger.
- LAN** Local Area Network.
- LEP** Large Electron-Positron Collider.
- LHC** Large Hadron Collider.
- LHCb** Large Hadron Collider beauty.
- LINAC** Linear accelerator 2.
- LO** Leading-Order.
- MC** Monte-Carlo.
- MPF** Missing E_T Projection Fraction.
- NAF** National Analysis Facility.
- NEMO** bwForCluster for Elementary Particle Physics, Neuroscience, and
Microsystems Engineering.
- NLO** Next-To-Leading-Order.
- NVMe** Non-Volatile Memory express.

- OTC** OpenTelekom Cloud.
- PDF** Parton Distribution Function.
- PF** Particle-Flow.
- PF jet** Particle-Flow jet.
- PLI** Particle Level Imbalance.
- PS** Proton Synchrotron.
- PUPPI** Pileup Per Particle Identification.
- QCD** Quantum Chromodynamics.
- RMS** Root Mean Square.
- ROOT** Data analysis tool.
- RWTH Aachen University** Rheinländisch-Westfälische Technische Hochschule Aachen.
- SPS** Super Proton Synchrotron.
- SRM** Storage Resource Manager.
- SSD** Solid State Drive.
- TOpAS** Throughput Optimized Analysis System.
- WAN** Wide Area Network.
- WLCG** Worldwide LHC Computing Grid.

Danksagung

An dieser Stelle möchte ich mich bei allen bedanken, die mir während der meiner Promotion zur Seite standen, mich unterstützt und motiviert haben. Mein größter Dank gilt Professor Günter Quast für das Ermöglichen meiner Promotion am ETP, insbesondere für seine begeisternde und weitsichtige Leitung der Gruppe, seine individuelle Förderung und die fachliche und organisatorische Unterstützung. Dazu gehört auch die Gelegenheit gleich zwei Interessensgebiete zu studieren, Teilchenphysik und Big Data Computing.

Besonderer Dank gilt auch meinen Betreuern des jeweiligen Themengebiets, Priv.-Doz. Klaus Rabbertz und Dr. Manuel Giffels. Ohne ihre fachkundige Unterstützung während des ganzen Zeitraums, die mir immer eine große Hilfe war, wäre diese Arbeit nur schwer möglich gewesen.

Des Weiteren bedanke ich mich für die tatkräftige Unterstützung im Computing-Bereich vonseiten Rene Caspart, Eileen Kühn und Max Fischer.

In diesem Zusammenhang dürfen auch meine langjährigen Kollegen im Bereich Computing Florian von Cube und Matthias Schnepf nicht fehlen, die in Diskussionen und mit Feedback viel zum Gelingen der Arbeit beigetragen haben. Vor allem Matthias Schnepf als langjährigem Begleiter während des Studiums und der Promotion gilt ein besonderer Dank.

Zusätzlich danke ich Daniel Savoiu und Thomas Berger für die analysespezifischen Erörterungen und die Hilfe bei programmiertechnischen Fragen.

Ein großer Dank gilt auch Martin Sauter, mit dessen Hilfe ich das entworfene Caching Konzept realisierte und ausführlich testete, und Tabea Feßenbecker, die das Projekt mit der Entwicklung einer Simulation bereit für zukünftige Weiterentwicklungen machte.

An alle übrigen Mitglieder des ETP geht ein zusätzlicher Dank, insbesondere den Korrekturlesern für die investierte Zeit und Mühe, und den Administratoren, die immer für einen reibungslosen Betrieb der komplexen Infrastruktur sorgten.

Besonders danke ich auch meiner Freundin Vanessa und meiner Familie, die mich während meiner ganzen Studienzeit und der anschließenden Promotion unterstützten und immer ein offenes Ohr hatten.

Zuletzt möchte ich den Institutionen danken, welche die für diese Arbeit benötigten Ressourcen bereitgestellt haben. Das beinhaltet das CMS Experiment mit dazugehöriger Kollaboration und die WLCG Partner mit Fokus auf das Tier 1 Rechenzentrum GridKa am KIT. Außerdem gehören dazu das vom Land Baden-Württemberg finanzierte und vom Rechenzentrum der Universität Freiburg betriebene HPC cluster NEMO, sowie die in Kooperation mit dem GridKa-Team im Rahmen des HNSciCloud Projekts getesteten OTC Ressourcen. Es beinhaltet des Weiteren die Karlsruhe School of Elementary Particle and Astroparticle Physics: Science and Technology (KSETA) für die finanzielle Unterstützung beim Besuch internationaler Konferenzen, Workshops und Sommerschulen.

Erklärung der selbständigen Anfertigung der Dissertation

Hiermit erkläre ich, dass ich die Dissertation mit dem Titel

*» Jet Momentum Resolution
for the CMS Experiment
and Distributed
Data Caching Strategies «*

selbständig und unter ausschließlicher Verwendung der angegebenen Hilfsmittel angefertigt habe.

Christoph Alexander Heidecker
Karlsruhe, den 24.06.2020