SCUOLA NORMALE SUPERIORE
PERFEZIONAMENTO IN FISICA

Dissertation

The CMS Tracker: contributions to hardware integration, software development and first data taking

by

Boris Mangano

Advisor

Professor **Lorenzo Foà**

Contents

Introduction

The Standard Model is a relativistic quantum field theory that describes all the particles observed in nature and their interactions, with the exception of gravitation. The predictions of this theory are in agreement with observations at the smallest scales and highest energies accessible to current experiments, making the Standard Model one of the most successful descriptions of the interactions among the components of matter. However, the Standard Model is an incomplete theory. The mechanism to give mass to the particles is introduced ad-hoc in the theory. Moreover, such a mechanism requires the existence of a new scalar particle, the Higgs Boson, which is currently unobserved. Theoretical considerations and the recent experimental evidences of neutrino oscillations and dark matter lead to consider the Standard Model as a low-energy effective theory. New physics is expected to emerge at higher energy scales than those accessible at current experiments.

The Large Hadron Collider (LHC) is a proton-proton collider that has been built at the CERN laboratory, near Geneva, Switzerland and started operations in September 2008. With about an order of magnitude increase in energy (14 TeV) compared to previous accelerators, the collider provides the possibility to probe an unexplored range of high energies and analyze a new domain of processes that may shed light on many of the open questions of physics at both sub-atomic and cosmological scales.

One of the two multi-purpose apparatuses installed at the LHC is the CMS detector, which is committed primarily to search for the Higgs boson and explore the physics beyond the Standard Model. A vital component of CMS is the silicon tracking system. The charged particle trajectories that are reconstructed by the Tracker provide crucial information for several key observables (such as muons, electrons and particle-flow objects) used in the analyses of LHC data.

This thesis describes the contribution of the author toward the construction of the *strip* sub-components of the Tracker and the development of the software used to reconstruct the trajectories of charged particles. The data recorded by the *complete* tracking system in its final configuration during

the CMS global runs of 2008 has been reconstructed using this software and analyzed by the author.

The chapters of the thesis are structured as follows.

- Chapter 1 introduces the Standard Model and the open questions of particle physics that motivated the construction of the LHC. Afterwards, the salient features of the collider are discussed and all the components of the CMS detector are described in detail, with the exception of the silicon Tracker.

- Chapter 2 is devoted to the silicon tracking system of CMS. After describing the physics requirements that drove the design of the Tracker, the chapter describes the pixel and the silicon detector modules employed within the tracking system. One section documents also the *module test procedure* that was implemented to select highly reliable detector modules before their installation on the Tracker sub-structures. The integration activities carried out by the author include the qualification of strip modules and their assembly on the four innermost silicon strip layers of the Tracker, which were constructed in the INFN laboratories of San Piero, Pisa.

- Chapter 3 describes the software which is used for the reconstruction of charged particles within the silicon tracking system. After an introduction to the general principles of track-finding and track-fitting, the implementation of the tracking algorithm used in CMS is analyzed in detail. The author reviews the performance of the reconstruction software on different samples of isolated muons, electrons and pions and, eventually, on realistically simulated LHC collisions.

- Chapter 4 explains the main extensions and improvements of the tracking software that were implemented with the help of the author and are now part of the official software of CMS.

- Chapter 5 explains how the tracking software, that was originally optimized for the reconstruction of particles emerging from LHC collisions, has been modified to reconstruct also cosmic muons. The cosmic data from the 2008 CMS global runs are analyzed by the author, measuring the track reconstruction efficiency and the transverse momentum resolution.

Chapter 1

The LHC collider and the CMS experiment

Starting from the description in Section 1.1 of the physics program that motivated the construction of the Large Hadron Collider (LHC) [1], the salient features of the machine built at CERN, Geneva, are discussed in Section 1.2. The last Section 1.3 introduces the Compact Muon Solenoid (CMS) [2, 3, 4]. All the sub-systems of the detector are described in details, with the exception of the silicon tracking system that is discussed in the next chapter. While the description of some components of CMS, such as the electromagnetic and hadronic calorimeters, is included in this Chapter only for completeness, the description of the detector solenoid and of the muon tracking system is preparatory for Chapter 4 and 5, respectively.

1.1 Overview of the LHC physics program

The Standard Model (SM) [5, 6] of particle physics provides an unified framework that describes both the Electroweak theory of electromagnetic and weak interactions between quarks and leptons [5, 6] and the Quantum Chromodynamics (QCD) theory of strong interactions among colored quarks [7, 8, 9, 10]. The model has the advantage of being perturbative at sufficiently high energies and is renormalizable [11, 12, 13]. A cornerstone of the SM is the mechanism of spontaneous ElectroWeak Symmetry Breaking (EWSB) that not only "generates" the masses of the intermediate vector bosons [14, 15, 16, 17], but also allows the model to be renormalized and unitary.

The Standard Model is in excellent agreement with many experimental results and a spectacular confirmation of its predictive power was the discovery of the Z^0 and $W^{\pm}$ intermediate vector bosons at the $Sp\bar{p}S$ accelerator of

CERN. Those particles were detected for the first time in 1983 [18, 19, 20, 21, 22] and their measured masses and production cross sections confirmed accurately the values predicted by the SM. About one decade later, the discovery of *top* quark at Tevatron [23] completed the third quark generation. The SM and its EWSB mechanism also predict an additional new particle, the scalar Higgs boson, that has not been observed so far by any experiment. Unlike the masses of the intermediary vector bosons, the mass of the Higgs boson is a free parameter in the SM and its value cannot be predicted. The Higgs boson could be so heavy that none past and current accelerators

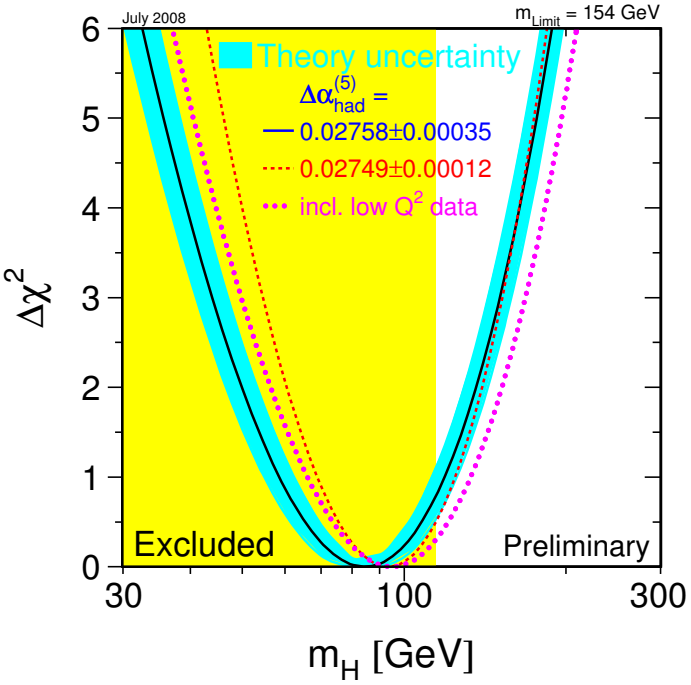

Figure 1.1: $\Delta = \chi^2 - \chi^2_{min}$ curve (black) derived from the fit of precision electroweak measurements, as a function of the Higgs boson mass, assuming the Standard Model to be the correct theory of nature. The preferred value for the Higgs boson mass, corresponding to the minimum of the curve, is at $m_H = 84$ GeV/c^2, with an experimental uncertainty of +34 and -26 GeV/c^2 (68% confidence level). Red and purple dashed lines are obtained using in the global fit a different set of measurements. The yellow band represents the values of m_H below 114 GeV/c^2, which have been excluded by direct search at the LEP collider [24].

collected enough center of mass (c.m.) energy to produce it. From direct

searches at LEP [25], we know that, if the Higgs exists, its mass m_H is greater than 114.4 GeV/c^2 (Fig. 1.1). On the other hand, theoretical arguments show that amplitudes of some scattering processes don't violate the unitarity only if m_H is smaller than about 700-900 GeV/c^2 [26]. The LHC and the CMS experiment are designed to explore the full range of possible Higgs masses between ~ 110 GeV/c^2 and ~ 1 TeV/c^2 and to unambiguously confirm or reject the existence of the Higgs boson.

While the investigation of the Higgs sector is one of the primary goals of the LHC, other new physics phenomena are expected to appear at the 1 TeV scale [27, 28]. The models based on Sypersymmetry [29] are among the most appealing theories beyond the Standard Model which are expected to manifest themselves at the energy accessible by the LHC. Assuming the existence of a super-partner for each known particle of the current SM, supersymmetric models can solve the "hierarchy problem" avoiding a fine-tuning of the theory parameters [30]. In addition, within the framework of the Minimal Supersymmetric extension of the Standard Model (MSSM), the three coupling "constants" of the fundamental interactions evolve, as a function of the renormalization scale, to a single common value (Fig. 1.2). This last aspect could be a step forward towards a Grand Unified Theory (GUT) of electroweak and strong forces, if Supersymmetry is confirmed by the LHC

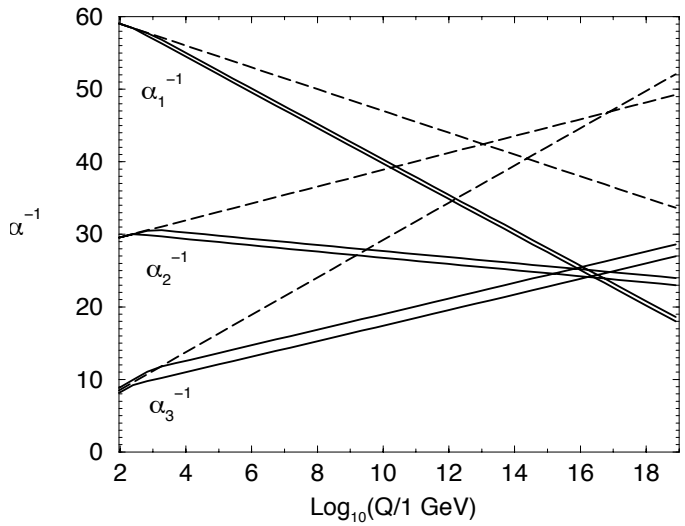

Figure 1.2: Evolution of the gauge couplings of the three Standard Model (SM) interactions as a function of the interaction energy scale. The dashed lines are calculated using the Standard Model, while the solid ones refer to its Minimal Supersymmetric formulation (MSSM). In the latter case, pairs of lines represent the theoretical uncertainty due to the unknown values of the sparticles masses [29].

The search for Supersymmetry signatures at the LHC is also motivated by an important puzzle in modern cosmology, that is the existence of Dark Matter. According to many versions of Supersymmetry, one of the superparteners is an heavy stable particle (e.g. the neutralino) which could serve as a WIMP (Weakly Interacting Massive Particle) and could explain, from a particle physics point of view, the existence of Dark Matter [31, 32].

The LHC could possibly shed light also on other models beyond the Standard Model that, inspired by the string theories, postulate the existence of extra dimensions aiming to unify all the fundamental forces, including gravitation [33].

The LHC is expected to produce such a large quantity of b-hadrons that further investigation on the violation of symmetry between matter and antimatter can be carried out. Finally, during the short runs of the accelerator with heavy-ion beams at energies over 30 times higher than at present day machines [34], quark-gluon-plasma could be observed and studied for the first time.

The Compact Muon Solenoid is, along with ATLAS, one of the two multipurpose experiments positioned in one of the four collision points of LHC. While CMS aims to cover the full LHC physics program described above, its design has been optimized to confirm or reject the existence of the Higgs boson and to elucidate the nature of electroweak symmetry breaking.

1.2 LHC accelerator and collider machine

The LHC machine boosts two beams of protons in opposite directions and makes them collide at an energy of 14 TeV in the c.m., seven times higher than the previous most powerful accelerator (Tevatron [35]). The hard scattering processes don't occur directly between the protons, but involve the quarks and gluons confined inside. For every new collision, the proton components can carry different fractions of the hadron energy, as described by the parton distribution functions (pdf): the energy in the c.m. of the hard scattering system is not constant at the LHC, but has a distribution related to the pdf of 7 TeV protons. Hence, as an hadron collider, the LHC can use the same configuration of beams to explore simultaneously different scattering processes within a wide range of energies and is an ideal machine for discovering particles of unknown mass. On the other hand, the energy available in the c.m. of the collision among quarks and gluons is rarely greater than $\sim 30\%$ of the initial energy of the protons system. For this reason, in order to entirely explore the range of energy up to about 1 TeV, the LHC has been designed to have two beams of 7 TeV protons.

Unlike other hadron machines like $Sp\bar{p}S$ and Tevatron, the LHC employs protons for both beams. Because of the much higher energy involved, the

dominant hard scattering processes at the LHC energy scale are mostly collisions between two gluons, as opposed to collisions between quarks and anti-quarks. Therefore the use of an anti-proton beam would not provide an advantage in term of cross section and, on the other hand, it would be an obstacle in achieving a high luminosity because of the difficult task of producing and collecting a high quantity of anti-matter. The use of two protons beams is one of the key features that allows the LHC to have a nominal luminosity of $10^{34} cm^{-2} s^{-1}$, two orders of magnitude greater than the Tevatron.

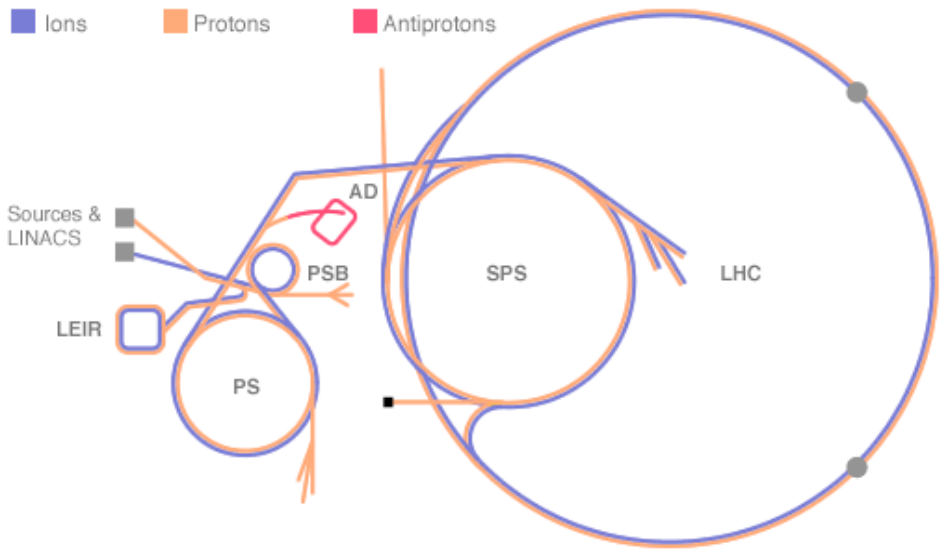

Figure 1.3: LHC Injection Chain. By means of a liner accelerator (Linac) and three synchrotrons (PSB, PS and SPS), a beam of protons is boosted to 450 GeV before it is finally injected in both rings of the LHC and accelerated to the nominal energy of 7 TeV.

1.2.1 The LHC layout

The decision to build the Large hadron Collider at CERN was strongly influenced by the cost saving made by re-using the existing injection chain and the 26.7 km tunnel that was constructed between 1984 and 1989 for the LEP machine.

The LHC layout, following the geometry of the LEP tunnel, consists of eight arcs and eight straight sections. Each straight section and the two half-arcs on both sides form one of the eight octants of the LHC ring (Fig. 1.4). Each arc is about 2.8 km long and contains the main dipoles that are used to bend the proton beams into the accelerator bores. Each straight section is 528

m long and can serve as an experimental or utility area. The ATLAS [36] and CMS experiments are respectively located at Interaction Point 1 (IP1) and Interaction Point 5 (IP5), along the two diametrically opposite straight sections where the beams cross at the highest luminosity. The other two experimental areas are located at IP2 (ALICE detector [37]) and IP8 (LHC-B detector [38]) along the straight sections of the second and eighth octants, respectively.

There are four utility regions in the LHC. The straight section 4 contains the Radio Frequency (RF) equipment that is used to accelerate the beams. The section 6 houses the beam dump system, which is used to extract the two beams from the machine at the end of each run or in case of abort. The third and seventh sections (IR3 and IR7 respectively) contain the beam

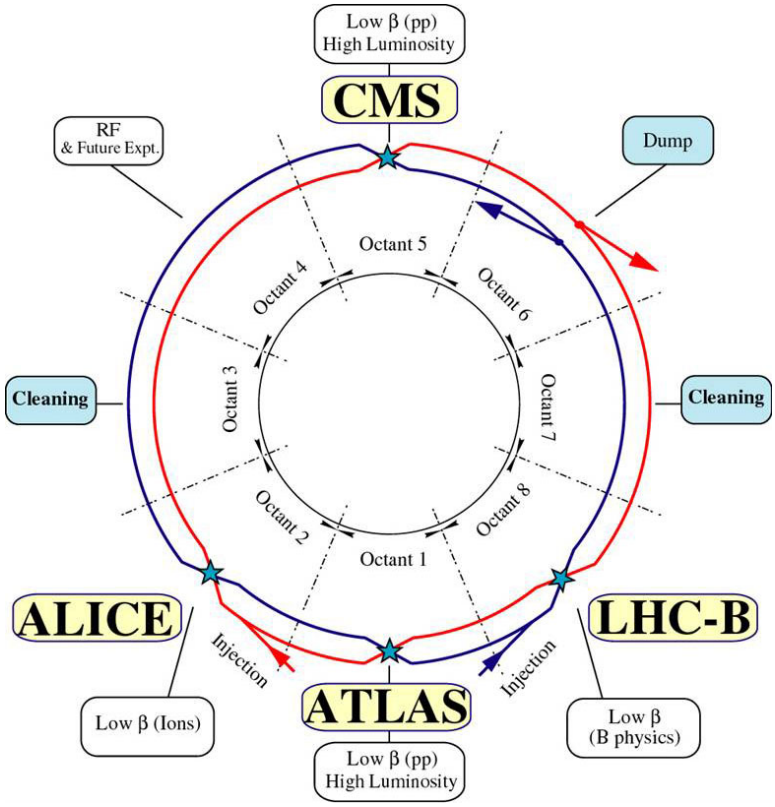

Figure 1.4: Layout of the LHC ring. It consists of 8 octants, each containing a straight sector which is completed by 2 half-arcs on both sides. Four sectors house the detectors of CMS, ATLAS, ALICE and LHC-B experiments, while the remaining octants enclose the cleaning, beam-dump and radio-frequency systems of the accelerator.

cleaning system: particles with a large momentum offset are scattered by the primary collimator in IR3 and particles with a large betatron amplitudes

are scattered by the primary collimator in IR7. In both cases, the scattered particles are absorbed by secondary collimators.

All the arcs and the straight sections contain quadrupole magnets which are used to focus the beams by reducing their emittance.

The second and eighth straight sections include also the injection system necessary to acquire into LHC the two beams coming from the pre-accelerator complex (Fig.1.3).

Bunches of protons are produced and accelerated up to 50 MeV in the Linac2 linear accelerator and then boosted to 1.4 GeV in the Proton Synchrotron Booster (PSB). The injection chain continues with the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS) that accelerate the protons to 26 and 450 GeV respectively. As a first step the bunches of protons are accumulated inside the two rings of LHC; subsequently they are accelerated to the nominal energy of 7 TeV.

Due to the eight straight sections placed between the arcs of the LHC and because of straight quadrupoles and other linear components placed along the arc sections themselves, the radius of curvature of the beam-bending dipole magnets is smaller than the average radius of the tunnel (i.e. 2 804 m and 4 249 m, respectively). Given that the highest magnetic field produced safely by the LHC dipoles is 8.33 T, the maximum energy of the circulating beams cannot exceed[1] 7 TeV.

1.2.2 Magnets

Since the LHC circulates particles with the same charge in both beams, two separate magnetic fields have to be used for two de-coupled accelerator bores. Nevertheless, the tunnel of LEP has an internal diameter of only 3.7 m, which prevents the installation of two physically distinct rings. This hard limit on space led to the adoption of a twin-bore magnet design: the LHC consists of one single closed loop where two sets of coils and two beam pipes are placed within the same mechanical structure and cryostat. Unfortunately the disadvantage of the twin-bore design is a magnetic and mechanical coupling between the two rings. The overall design of the mag-

[1] In absence of an electric field, the relation among the transverse momentum p_t of a charged particle, the curvature of its trajectory and the value of a perpendicular magnetic field B is given by the formula:

$$p_t = 0.3 \cdot R \cdot B$$

where R is the radius of curvature measured in meters, B is expressed in Tesla and p_t is in GeV/c.

nets is necessarily more challenging [2].

Given that the curvature of the LHC ring is not an adjustable parameter, but is constrained by the anatomy of the existing LEP tunnel, the proton beams can be correctly bent inside the machine's radius and accelerated to unprecedented energies only using a new generation of dipole magnets. The 1 232 superconducting dipoles of LHC have coils of Niobium-titanium alloy (NbTi), which are cooled down to 1.9 K and can produce fields of 8.33 T along all the arcs of the accelerator rings.

During LHC operation, the magnets are cooled to cryogenic temperatures and the NbTi coils become superconducting: the electrical resistance of the alloy drops exactly to zero and currents of up to several thousand amperes can flow inside the coil windings. In general, superconductive magnets are able to produce much stronger magnetic fields than ordinary iron-core electromagnets and they are cheaper to operate, since no power is lost in ohmic resistance.

The temperature below which a material becomes superconductive is called the critical temperature. The NbTi critical temperature is 10 K and magnets based on this alloy have been already used in other experiments (Tevatron-FNAL, HERA-DESY and RHIC-BNL): cooling the dipoles to 4.2 K by means of supercritical helium, fields of about 5 T were obtained. The LHC dipoles are cooled below 2 K using superfluid helium and they can safely operate at fields above 8 T.

In order to keep the two beams inside the accelerator, at the same time as the protons are boosted, the magnetic field of the dipoles has to be turned up in proportion (synchrotron concept). While the field changes, the coil windings may move and the consequent friction can create hot-spots that "quench" the magnet out of its cold superconducting state and cause the beams to be lost. For this reason, before starting operation, all the magnets of LHC have been "trained" by ramping up and down their fields many times in order to remove any small wrinkle from the coils and allow the magnets to run in the nominal configuration without quenching.

In addition to the dipoles, the LHC rings contain about 850 focusing quadrupoles based on the same superconducting NbTi technology and another ∼6200 magnets, including superconducting and ordinary ones, used to adjust the beams' trajectories.

[2] By comparison, due to the less stringent civil engineering constraints, the now-canceled proton-proton Superconducting Super Collider (SSC) was originally designed to have two separate rings [39]

1.2.3 Machine Luminosity

The number of events generated per time unit in LHC collisions is given by the formula:

$$N_{\text{event}} = L \cdot \sigma_{\text{process}} \tag{1.1}$$

where σ_{process} is the cross section for the process under study and L the machine instantaneous luminosity. The luminosity depends on the beam and machine parameters and can be written, for a Gaussian beam distribution, as:

$$L = \frac{N_b^2 \cdot n_b \cdot f_{rev} \cdot \gamma_r}{4\pi \cdot \varepsilon_n \cdot \beta*} \cdot F \tag{1.2}$$

where N_b is the number of particles per bunch, n_b the number of bunches per beam, f_{rev} the revolution frequency, γ_r the relativistic gamma factor, ε_n the normalized beam emittance, $\beta*$ the beta function at the interaction point and F the geometric luminosity reduction factor due to the crossing angle between the two beams. An expression for F is:

$$F = \left(1 + \left(\frac{\theta_c \sigma_z}{2\sigma*}\right)^2\right)^{-\frac{1}{2}} \tag{1.3}$$

where θ_c is the full crossing angle at the IP, σ_z the bunch length RMS and $\sigma*$ the transverse beam size RMS.

Number of particles per bunch	N_b	$1.15 \cdot 10^{11}$	[part/bunch]
number of bunches per beam	n_b	2808	
revolution frequency	f_{rev}	11.245	[kHz]
relativistic gamma	γ_r	7461	
normalized transverse emittance	ε_n	3.75	[μm]
beta-function	$\beta*$	0.55	[m]
geometric reduction factor	F	0.836	
crossing angle at IP1 & IP5	θ_c	285	μrad
RMS bunch length	σ_z	7.55	cm
RMS beam size	$\sigma*$	16.7	μm

Table 1.1: Beam and machine parameters relevant for the estimation of the peak luminosity at interaction points 1 (ATLAS) and 5 (CMS).

Using the nominal parameters of the LHC (Table 1.1), the formula 1.2 gives a peak luminosity of $\mathcal{L} = 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$. In the early phase of the LHC

machine operation, part of the beam dump and collimation systems are staged and the beam current is not allowed to exceed 50% of its nominal value. Hence the peak luminosity is limited to $\mathcal{L} = 2 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

1.2.4 Beam lifetime and Integrated Luminosity

The peak luminosity defined in Eq. 1.2 is not constant over a physics run, but it decays due to the degradation of intensities and emittances of the circulating beams. The main cause of the decay during nominal LHC operation is the beam loss from collisions. The initial decay time of the bunch intensity, due to this effect, is:

$$\tau_{nuclear} = \frac{N_{tot,0}}{L_0 \cdot \sigma_{tot} \cdot k} \tag{1.4}$$

where $N_{tot,0}$ is the initial beam intensity, L_0 the initial luminosity, σ_{tot} the total cross section at 14 TeV and k the number of interaction points. Assuming an initial peak luminosity of $L_0 = 10^{34} cm^{-2} s^{-1}$ and two high luminosity experiments, the above expression yields an initial decay time $\tau_{nuclear}$ = 44.85 h for each beam. Consequently the instant luminosity as a function of time can be expressed as:

$$L(t) = \frac{L_0}{(1 + t/\tau_{nuclear})^2} \tag{1.5}$$

and the time required to reach 1/e of the initial value L_0 is given by

$$\tau_L = (\sqrt{(e)} - 1) \cdot \tau_{nuclear} \simeq 29 \text{ h} \tag{1.6}$$

Because of other secondary contributions affecting the beam intensity and focus, the actual luminosity lifetime is expected to be approximately τ_L = 14.9 hours.

Filling the proton bunches of the LHC requires about 16 minutes and the minimum time required for ramping the beam energy from the injection value of 450 GeV to the final value of 7 TeV is approximately 20 minutes. After a beam abort at the 7 TeV energy, it takes approximately 20 minutes to ramp the magnets down to the field value corresponding to 450 GeV beams. Therefore, the theoretical limit of the total *turnaround time* for the LHC is about 1 hour. From the experience of other hadron machines, it is reasonable to assume a factor 6-7 between the theoretical and the final average turnaround time: as a result, the turnaround time of the LHC is expected to be around 7 hours.

Integrating the luminosity over one run yields:

$$L_{int} = L_0 \cdot \tau_L \cdot (1 - e^{-T_{\rm run}/\tau_L}) \tag{1.7}$$

where $T_{\rm run}$ is the total length of the run. Given an average turnaround time of 7 hours (with respect to 1 hour) and a luminosity lifetime τ_L of 14.9 hours, the optimal run time is around 12 hours (5.5 h). Assuming that the accelerator can be operative for 200 days per year, the maximum total integrated luminosity per year is of the order of 80 fb^{-1} (120 fb^{-1}).

1.3 Description of the CMS detector

The design of the CMS apparatus was driven by three major guidelines:

- all the sub-detectors have to be precise, efficient and fast enough so that CMS can extract interesting events from the huge amount of experimental data produced by the LHC and successfully process them. Two important challenges for CMS are set by the short time interval between two consecutive bunch crossings (25 ns) and the large number of superimposed proton-proton collisions in the same crossing (about 20).
- most of the detector components must be resistant to the huge dose of hadronic radiation coming from the LHC collisions and they have to operate safely for ∼10 years of data taking;
- the total cost of the experiment has to be affordable and budget constraints should not compromise the completion of the detector construction and its subsequent operation and maintenance.

In the rest of this chapter and in chapter 2, major components of CMS are described using the coordinate system adopted by CMS. This system has the origin centered at the nominal collision point inside the experiment, the y-axis pointing vertically upward, and the x-axis pointing radially inward toward the center of the LHC. Thus, the z-axis points along the beam direction toward the fourth octant from IP5. The azimuthal angle ϕ is measured from the x-axis in the $x - y$ plane and the radial coordinate in this plane is denoted by r. The polar angle θ is measured from the z-axis. Pseudorapidity is defined as $\eta = -\ln\tan(\theta/2)$.

1.3.1 General Detector layout

CMS is a 4π "hermetic" detector (Fig. 1.5). It is developed around a 13-m-long, 6-m-inner-diameter, 4 T superconducting solenoid providing a large

bending power (12 Tm). The bore of the magnet coil is large enough to accommodate both the inner tracking system and the calorimeters inside it. Outside the magnet, a system of muon chambers is embedded in the return field yoke.

The tracking volume is given by a cylinder of 5.8 m length and 2.6 m diameter, which hosts sensitive layers exclusively based on radiation-hard silicon sensors. Strips and pixels modules provide the granularity and precision that are necessary to cope with the high multiplicity of charged particles produced in LHC collisions and to achieve, in combination with the 4 T magnetic field, an accurate measurement of particle transverse momenta and impact parameters. The silicon tracking system is hermetic up to $|\eta| \lesssim 2.5$.

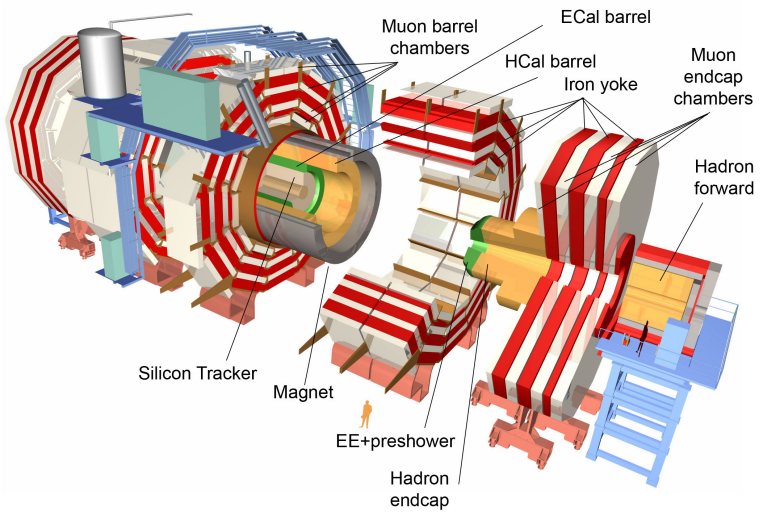

Figure 1.5: Exploded view of the CMS apparatus.

The calorimetric system consists of both electromagnetic and hadronic detectors. The electromagnetic calorimeter (ECAL), along with the tracking system, is devoted to the identification and reconstruction of electrons and photons. It uses lead tungstate (PbWO4) scintillation crystals organized in a barrel and two end-caps structures. A preshower system is installed in front of the two end-cap to facilitate the discrimination between prompt photons and neutral pions. The ECAL sub-detector has an uniform coverage in ϕ and in pseudorapidity up to $|\eta| < 3.0$, but two small gaps around $|\eta| = 1.479$ between the barrel and the end-caps that are necessary to anchor the inner tracking system and route its services.

Between the ECAL volume and the solenoid magnet, sits the hadronic calorimeter (HCAL) also with a coverage up to $|\eta| < 3.0$. Unlike the electromagnetic calorimeter, the hadronic calorimeter is a sampling calorimeter, which consists of alternating layers of brass and scintillator. Outside the

magnet, the HCAL is completed by the Outer Hadron calorimeter (HO), also known as "tail-catcher", which assures the full containment within the calorimeter of the most energetic jets of particles. Finally, the forward Cerenkov calorimeter (HF) guarantees full geometric coverage of the calorimetric system up to $|\eta| < 5.0$.

The return field of the magnet is saturated in 1.5 meters of iron that hosts stations for the identification and reconstruction of muons: layers of drift tubes (DT) are set in the barrel region, while cathode strip chambers (CSC) are used in the end-caps. For redundancy and in order to have a fully efficient Trigger, stations based on resistive plate chambers (RPC) are placed both in the barrel and in the endcaps.

Altogether the CMS detector is 21.6 m long, has a diameter of 14.6 m and a total weight of about 12 500 tons. The ECAL has a thickness of more than 25 radiation lengths (X_0), while the thickness of the HCAL system varies in the range of 10-15 interaction lengths (λ_I) depending on η.

1.3.2 Magnet

The superconducting magnet of CMS [40] has been designed to reach a 4 T field in a free bore of 6 m diameter and 12.5 m length, for a stored energy at full current of 2.6 GJ[3]. The field flux is returned through a 10 000 ton iron yoke consisting in a barrel structure made of three concentric cylinders divided in five wheels, and two end-caps structures composed of three disks each (Fig. 1.5). These elements can be moved quite easily to facilitate the assembly of the sub-detectors. In order to facilitate the matching between trajectories reconstructed in the inner tracker and corresponding energy deposits in the calorimeter, the radial extent of the coil (ΔR) has to be preferably small. The coil of CMS has a small $\Delta R/R$ ratio of ~ 0.1 due to an innovative "self-supporting" design. The conductor itself includes a structural material that limits the magnetic and mechanic stress level inside the winding and prevents a premature quench.

Due to the much higher permeability of iron compared to air, the yoke provides a path of low reluctance for the magnetic field flux. Figure 1.6 shows that the magnetic inductance B of the return flux in the barrel is between 1 and 2 Tesla within the iron elements. In the regions occupied by the muon chambers the B field is usually below 0.1 T, except for peaking at 0.3 T near the air gaps in between the yoke wheels. In the end-caps the magnetic field lines aren't completely contained within the yoke disks and the B field is up to about 1 T in the air gaps occupied by the end-cap muon detectors. As it is discussed in Section 1.3.3, the significant value of the B field motivated the design of a different type of muon detectors for the end-caps, with respect to the Drift tubes employed in the barrel. Inside the volume of the solenoid bore the magnetic inductance varies between 4 T (central region) to approximately 3 T (sides). Although the field is relatively uniform within the volume occupied by the Tracker, it is not exactly constant. Chapter 4 will describe how the small inhomogeneities of the field in the end-caps of the Tracker have to be accounted for to achieve the ultimate resolution in measuring charged particles momenta.

1.3.3 Muon Tracking system

The Muon Tracking system has three primary functions:

- identification of muon particles.
- measurement of muon momenta.
- bunch crossing identification and triggering.

[3]For comparison, the energy stored in a 8.33 T dipole of LHC is 6.93 MJ.

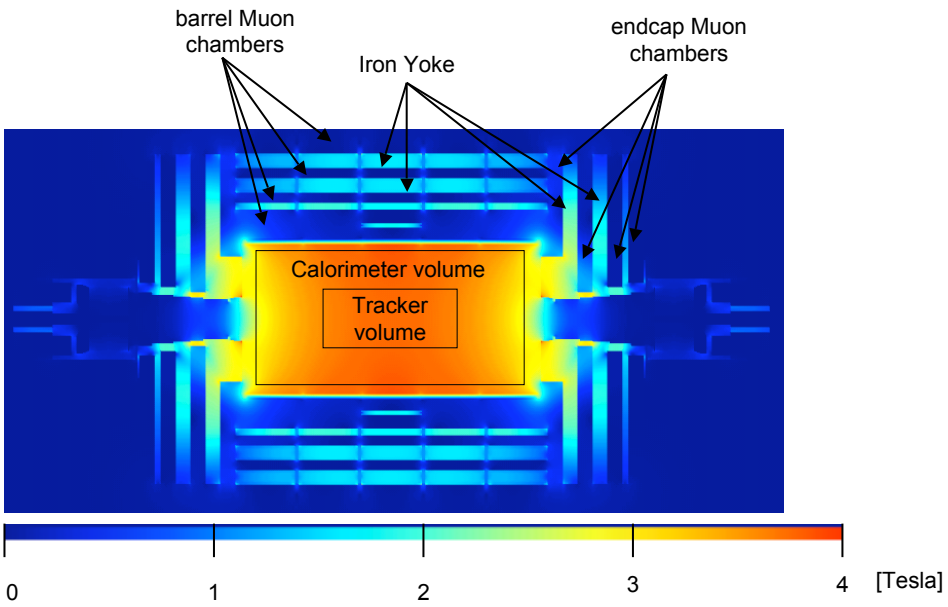

Figure 1.6: Longitudinal R-z view of CMS showing the module of the magnetic inductance B within different subsystems of the detector.

A good standalone measurement of the muon momentum is assured by the long lever arm of the muon system (about 3.4 m) and by the high field in the flux-return yoke of the magnet. The iron yoke also serves as a hadron absorber for the identification of muons: a particle reaching the outermost muon chambers crosses between 20 and 25 interaction lengths of material (Fig. 1.7). About 25 000 m^2 of detection planes are covered by muon chambers in CMS. The technology used is inexpensive, but, at the same time, reliable and robust. Despite the short time interval between consecutive LHC collisions, the muon system is sufficiently fast to be used for triggering. In the following the three types of muon detectors implemented in CMS are described.

Drift Tubes

In the barrel region, the neutron-induced background is small, the muon rate is low and the 4 T magnetic field is uniform and mostly contained in the iron yoke. In this region drift chambers with standard rectangular drift cells are used.

The Drift Tube (DT) chambers cover the range of $|\eta| < 1.2$ and are organized into four stations interspersed among the three layers of the flux-return wheels (Fig. 1.8). Each wheel is divided in 12 ϕ sectors containing stations

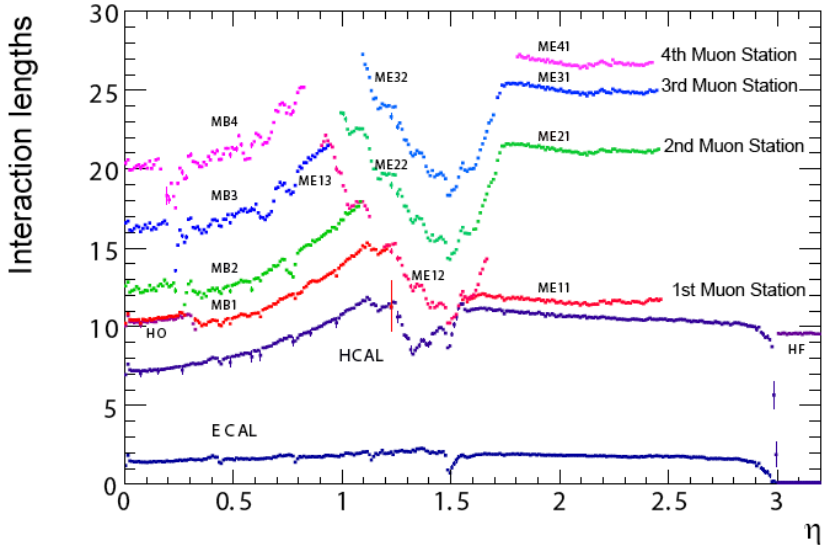

Figure 1.7: Material thickness of the CMS detector expressed in interaction lengths at various depths, as a function of pseudorapidity.

labeled as MBxy, where x = 1-4 and y is the number of the ϕ sector.

The yoke supports that are between two stations of two consecutive sectors generate 12 unavoidable dead zones in the ϕ coverage. However, the supports are placed so that they do not overlap and muons coming from LHC collisions usually don't cross more than one dead zone.

Each of the first three stations contains two set of chambers (two super-layers), which measure the muon coordinate in the $r-\phi$ bending plane, and an additional super-layer, which provides a measurement in the z direction along the beam line (Fig.1.9). The fourth station does not contain the z-measuring planes. The length of the chambers arranged along the beam line direction is limited to 2.4 m by the longitudinal segmentation of the barrel yoke wheels. The transverse dimension of the drift cells is 21 mm (Fig.1.10) and corresponds to a drift time of 380 ns using the gas mixture of 85% Ar + 15% CO_2. This drift time is small enough to produce a negligible occupancy and the cell size is sufficiently large to limit the number of active channels to an affordable value. The drift cells of each chamber are offset by a half-cell width with respect to their neighbors to eliminate dead spots in the efficiency and also to allow the time-alignment of the muon chambers. Indeed, staggered adjacent cells are expect to collect the signal from the same ionizing particle with a fixed delay which corresponds to the time of drift along half-cell width.

One Super-Layer (SL) of 4 staggered drift tube planes gives excellent time-

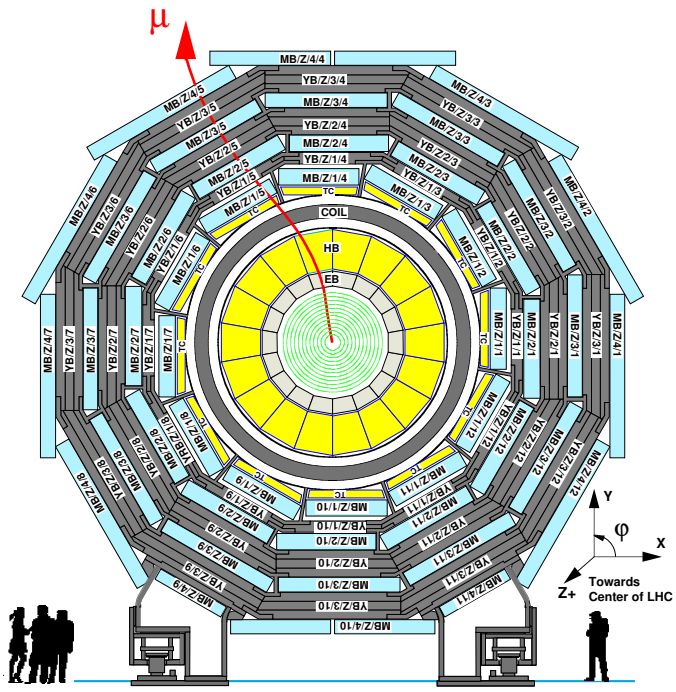

Figure 1.8: Transverse view of the Muon Barrel system. It consists of 12 ϕ-sectors, each containing four stations of muon chambers interposed among the three iron-yoke layers.

tagging capability with a time resolution of a few nanoseconds. This capability provides stand-alone and efficient bunch crossing identification.

The electronics of the Drift tubes is heavily integrated: the chamber gas volume contains the front-end electronics (signal amplifier, shaper and comparator), the readout (signal digitizer and buffer) and also part of the L1 trigger boards. Independent track segments are reconstructed locally from each chamber SL and subsequently transmitted through a fast network to the 70-m-far underground control room, where the global L1 trigger algorithms are run.

An important components of the DT system is the circuit that provides fresh Ar/CO_2 gas mixture to all drift cells. A return line for each barrel wheel brings samples of the gas back to the gas room where they are analyzed. The oxygen and humidity content of the gas mixture are monitored constantly in order to avoid variations in the drift velocity inside the DT cells.

According to test beam results [41], DT cells have an average spatial resolution of about 170 μm. In the presence of stray magnet field outside the iron yoke, the cell resolution is slightly deteriorated, but still better than the

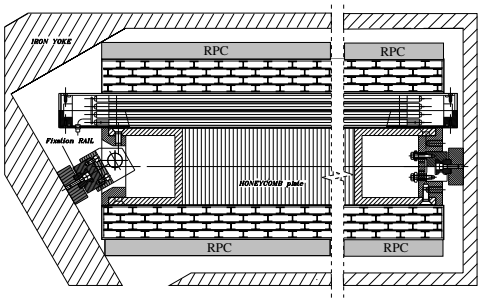

Figure 1.9: Transverse view of a barrel muon station. It consists of two Drift Tube (DT) chambers with anode wires parallel to the z-axis, one central DT chamber with wires perpendicular to z and two Resistive Plate chambers.

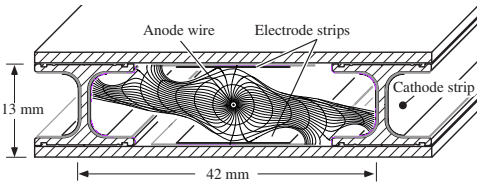

Figure 1.10: Transverse section of a Drift Tube cell. Electrode strips compensate the distortion of the drift lines (shown in figure) due to the residual magnetic field inside the cell volume.

250 μm design requirement. Combining the 8 measurements of two superlayers, a DT chamber can determine the ϕ-coordinate of a muon with an accuracy finer than 100 μm. Single cells detect muon with an efficiency higher than 99% and have a bunch crossing identification efficiency of about 92%. Combining the 8 or 12 layers of cells in a chamber, bunch crossings can be identified with an efficiency around 99.7%.

Cathode Strip Chambers

In the two end-caps of CMS the muon rate and radiation level are higher than in the barrel. Besides, the magnetic field is significantly different from zero and not uniform outside the return yoke (see Sec. 1.3.2). In these regions, the muon system consists of Cathode Strip Chambers (CSCs).

The CSCs are multi-wire proportional chambers containing 6 anode wire planes interleaved among 7 cathode panels (Fig. 1.11). The cathode strips of each chamber run radially outward and provide a precision measurement of the $r-\phi$ coordinate, obtained interpolating charges induced on near strips. The anode wires run approximately perpendicular to the strips and are also read out in order to provide measurements of r and of the beam-crossing time.

Compared with DTs, the ionization charges of CSCs drift on a very short distance and such chambers can operate at high rate and inside a large and non-uniform field. The gas mixture doesn't require a precise control of temperature or pressure.

For each end-cap of CMS, there are four stations with CSC chambers positioned perpendicular to the beam line and in between the flux-return plates.

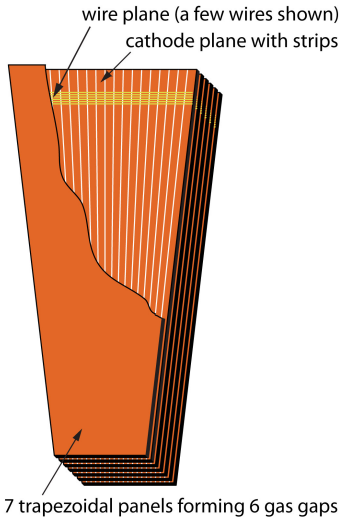

Figure 1.11: A Cathode Strip Chamber is a trapezoidal structure containing 6 planes of anode wires interposed among 7 cathode surfaces.

The region of pseudorapidity covered is $0.9 < |\eta| < 2.4$.

Figure 1.12 shows how the CSC chambers are arranged on three rings on the innermost disk and two rings for each of the three remaining disks. Each ring contains 36 or 72 trapezoidal chambers covering 20 or 10 degrees in ϕ. Except for the outer ring of the inner disk (ME1/3), the chambers overlap and provide a contiguous ϕ-coverage. There is a total of 468 CSCs for each end-cap, being the 72 chambers of ring ME4/2 (see Fig. 1.12) initially staged.

The mechanical structure of a CSC is based on seven 16.2-mm-thick trapezoidal panels. Six of the panels have a pattern of 80 strips milled on one side. Strips, being radial, have a pitch that varies from 8.4 mm at the narrow chamber end to 16 mm at the wide end.

Three of the internal panels are so-called anode panels (panels 2, 4 and 6) and the anode wires are wound around them. The wire spacing is about 3.2 mm, but groups of 16 wires make 1 single anode read-out channel with a width of about 5 cm. The 72 ME1/1 chambers have a slightly different mechanical design with respect to the other CSCs. The gas gap is 7 mm (compared with 9.5 mm), the wire spacing is 2.5 mm (3.2 mm) and, most notably, the anode wires of these chambers are not azimuthal, but tilted (Fig. 1.13). Unlike the other CSCs, the ME1/1 chambers are inside the CMS solenoid and see an approximately uniform axial field of 4 T. The wire tilt compensates for the Lorentz angle, so that electrons keep drifting parallel to the strips. Hence, a simply weighted mean of the strips signals provides an unbiased measurement of the ϕ coordinate of the transversing muon.

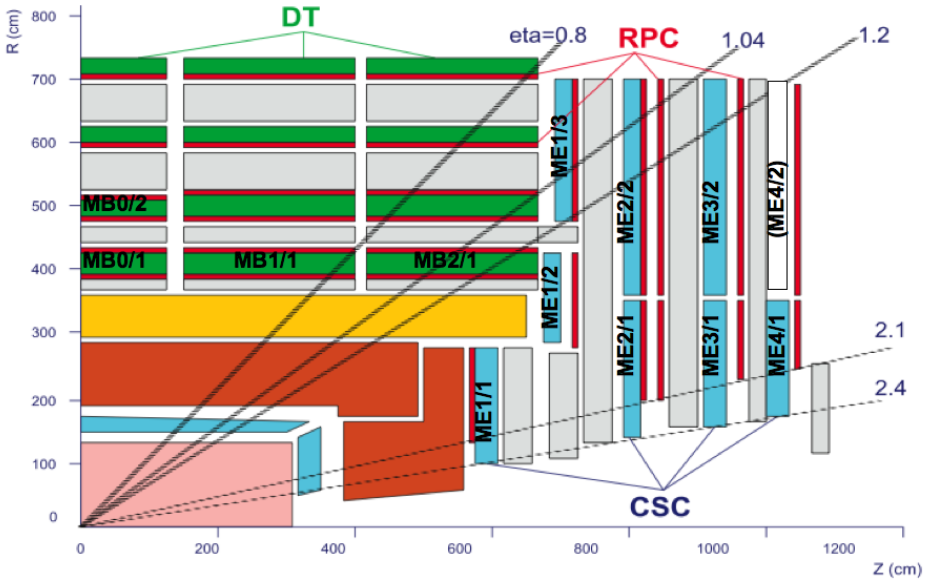

Figure 1.12: Longitudinal R-z view (one quadrant) of the Muon system. In the barrel, the DT and RPC chambers are arranged on four muon stations interposed among the three rings of the barrel iron yoke. Each end-cap consists of four disks supporting chambers arranged on two or three rings of CSC and RPC detectros. The outer CSC ring of the fourth disk (M4/2) has been staged.

Similar to the DTs, the CSC stations have integrated electronics to identify patterns of anode wire measurements that are compatible with a muon track coming from the collision point. Track segments are reconstructed independently also for cathode strip measurements. Finally, strip and anode candidates are combined in 3-dimensional track segments that are subsequently processed by the L1 Trigger system.

Each anode plane of a CSC detects the passage of mouns with a resolution of 11 ns. Combining all 6 planes of a chamber, the resolution improves to 5 ns, allowing the bunch crossings to be identified with $\sim$ 99% efficiency. Combining the measurements of all the strip cathode planes, the offline spatial resolution on the ϕ-coordinate is about 33 μm for the chambers with the finer pitch and $\sim$ 80 μm for all the others.

Resistive Plate Chambers

The Resistive Plate Chambers (RPC) provide a complementary system for triggering and bunch crossing identification. They are positioned both in

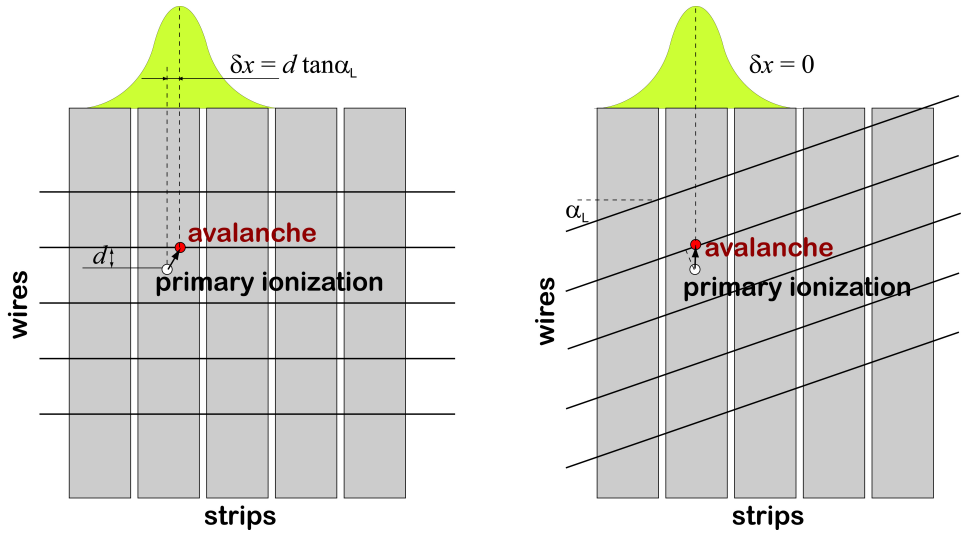

Figure 1.13: Formation of ionization signal on the strips of a Cathode Strip Chambers (CSC) installed within the magnetic field volume. Because of the Lorentz drift, an arrangement with anode wires placed perpendicular to the strips (left) is prone to produce a bias in the reconstructed position of the primary ionization. However, the bias is canceled if the wires are slightly tilted (right), as it is implemented for the CSC chambers of the ME1/1 ring.

the barrel and in the two end-caps, covering a large portion of the $|\eta| < 1.6$ pseudorapidity range. A total of 6 layers of RPCs are embedded in the barrel muon system: 2 layers in each of the first 2 stations (Fig. 1.9), and 1 in each of the last 2 stations. The redundancy in the first 2 stations allows the trigger algorithm to work even for low-p_T muons that may be absorbed in the iron yoke before reaching the two outer stations. In the end-caps, rings of RPCs are coupled to the 3 CSC disks that are closer to the IP (Fig.1.12). A third inner RPC ring can be installed extending the geometrical acceptance of the RPC system up to $|\eta| = 2.1$, but it is staged for the early operation of CMS.

Each RPC is a parallel-plate detector containing two 2-mm gaps filled with gas and operated in avalanche mode. Read-out strips are placed in between the two gaps (Fig. 1.14), so that they can collect a signal which is the sum of the two single-gap signals. The RPCs have an efficiency higher than 97% and time response and time resolution performance similar to those of scintillators. The position measurement is less accurate than DT's or CSC's, but still adequate.

The barrel RPCs are rectangular-shaped and have strips running along the z-axis. The end-cap RPCs have a trapezoidal shape and their strips are

radially arranged. In both cases, the coordinate which is measured is ϕ.

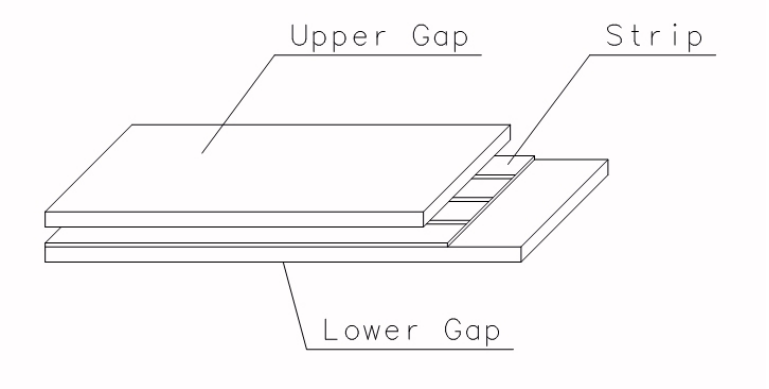

Figure 1.14: Schematic view of a Resistive Plate Chamber installed in the barrel region. A plate with read-out strips is inserted between two gaps filled with gas.

Laboratory tests showed that barrel (end-cap) RPC chambers have an average efficiency around 97% (95%).

Muon system performance

Combining the information provided by DT, CSC and RPC, the muon system can measure the transverse momentum of muons from LHC collisions with a p_t resolution of about 10% in the barrel (Fig. 1.15) and about 20-30% in the end-caps (Fig. 1.16). The precision of this measurement is dominated by multiple scattering up to p_t values of about 200 GeV/c, when the chambers spatial resolution starts to dominate. Without the obstacle of the multiple scattering effects due to the calorimeters, the magnet and the return-flux yoke, the inner silicon Tracker generally provides better p_t measurements than the muon system. Nevertheless, because of the longer lever arm, the combination of inner tracking system and muon system provides the best measurements, in particular for high momentum muons (Fig. 1.15 and Fig. 1.16).

1.3.4 Electromagnetic calorimeter

The Electromagnetic Calorimeter (ECAL) is a hermetic homogeneous calorimeter made of 61 200 lead tungstate ($PbWO_4$) scintillation crystals mounted in a central barrel structure, closed by other 7 324 crystals in each of the two end-caps (Fig. 1.20). The high density (8.28 g/cm^3), short radiation length ($X_0 = 0.89$ cm) and small Molière radius (2.2 cm) make the $PbWO_4$

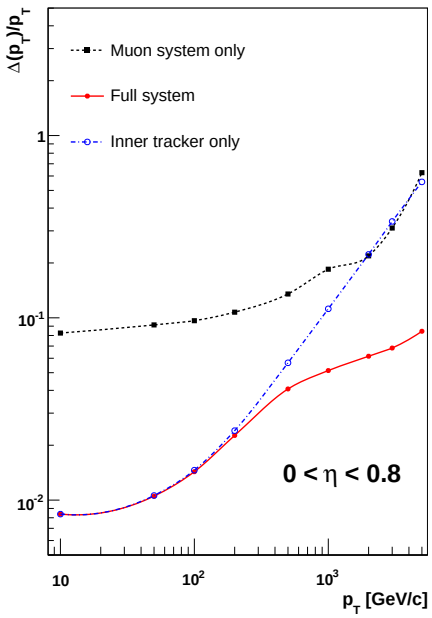

Figure 1.15: Resolution on transverse momentum for muons reconstructed in the barrel region of CMS using only the muon system (black), only the tracker (blue) or both (red).

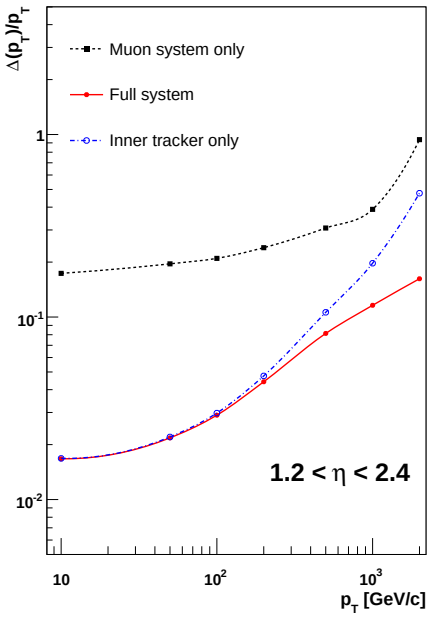

Figure 1.16: Resolution on transversemomentum for muons reconstructed in the end-cap region of CMS using only the muon system (black), only the tracker (blue) or both (red).

crystal an ideal material for the identification and energy measurement of electrons and photons by means of a compact and fine-granularity detector.

The barrel part of the ECAL (EB) covers the pseudorapidity range $|\eta| < 1.479$ with a matrix of 360×170 crystals uniformly arranged in the ϕ-η plane. The crystal cross-section corresponds to approximately 0.0174×0.0174 in ϕ-η or 22×22 mm^2 at the front crystal face, while the rear one is 26×26 mm^2. The centers of the front faces are at a radius of 1.29 m from the beam line. The crystal length is 23 cm, which corresponds to 25.8 X_0. The nominal crystal to crystal distance is 0.35 mm inside a submodule (5x5 crystals structure) and 0.5 mm between submodules. Each barrel crystal is mounted in such a way that its axis makes a small angle ($\sim$ 3 degrees) with respect to the vector pointing from the crystal's face toward the CMS origin (Fig. 1.19). This arrangement is necessary to avoid dead zones aligned with particle trajectories and assures that the calorimeter is fully hermetic. However, the drawback is that all the barrel crystals have to be shaped slightly differently to implement this quasi-projective geometry.

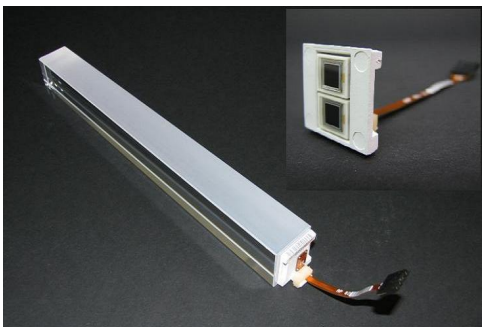

Figure 1.17: Picture of an ECAL barrel crystal and the corresponding Avalanche PhotoDiode, which is used to collect the scintillation light.

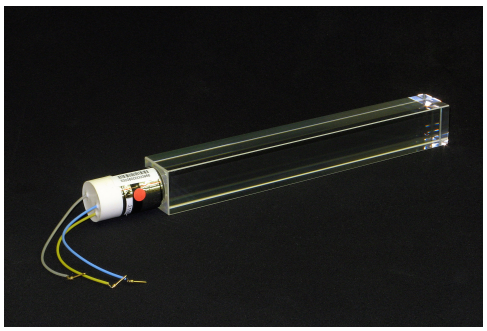

Figure 1.18: Picture of an ECAL end-cap crystal and the corresponding Vacuum PhotoTriode, which is used to collect the scintillation light.

The crystals of the two ECAL end-caps (EE) cover the pseudorapidity range $1.479 < |\eta| < 3.0$. They have a length of 22 cm (24.7 X_0) and a front cross-section of 28.6×28.6 mm^2. The longitudinal distance between the interaction point and the end-cap envelope is 315.4 cm, taking into account the estimated shift towards the interaction point (1.6 cm) when the 4 T magnetic field is switched on. To avoid dead zones, the axes of EE crystals don't point toward the center of CMS, but they are off-pointing by 2-8 degrees (Fig. 1.19).

The ECAL crystals have to withstand the radiation levels and particle fluxes throughout the duration of the experiment. While the quality of $PbWO_4$ has been progressively improved [42, 43, 44], ionizing radiation in the crystal still produces absorption bands that reduce the light transmission. This kind of damage is constantly monitored and corrected for by injecting laser light into the crystals and measuring the optical transparency.

The scintillation light produced by the passage of of particles in the crystals is collected by Avalanche PhotoDiodes (APDs) in the barrel (Fig. 1.17) and by Vacuum PhotoTriodes (VPTs) in the end-caps, the latter being tailored to work with an higher level of radiation (Fig. 1.18).

The number of scintillation photons emitted by the $PbWO_4$ depends on the temperature of the crystals and similarly the gain of the APDs is temperature and voltage dependent. Hence the electromagnetic calorimeter uses very steady power supply and a cooling system capable of keeping the temperature stable within 0.05°C.

For each bunch crossing, the ECAL data, in the form of trigger primitives, are sent to the Level-1 calorimeter trigger processor. While the high-granularity data are buffered in the front-end electronics, the trigger primitives consist in the summed transverse energy deposited in a trigger tower,

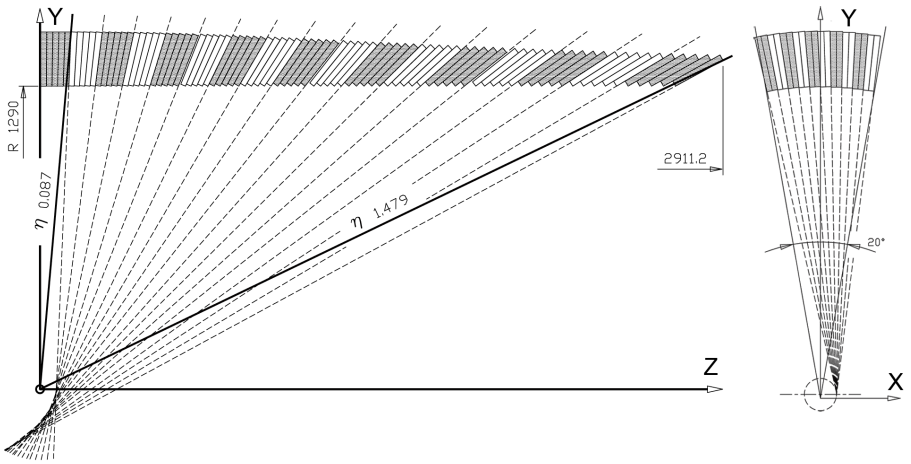

Figure 1.19: Quasi-projective geometry of electromagnetic calorimeter crystals in the barrel. This arrangements reduces dead-areas in the calorimeter avoiding that a photon or an electron, produced at the Interaction Point, can go through the small passive cavities between two consecutive crystals.

i.e. the matrix of 5×5 crystals that match a HCAL tower. Each trigger primitive contains also an additional bit that characterizes the lateral profile of the electromagnetic shower. In case of Level-1 acceptance, the full high-granulartity signal is retrieved from the buffers and used in the High-level trigger or offline reconstruction.

Preshower

The end-cap crystals don't have enough granularity to distinguish two adjacent showers, initiated by the decay products of a π^0, from single-photon showers. The principal aim of the CMS Preshower is to identify π^0 particles by means of a higher granularity silicon strip detector placed in front of the ECAL end-cap crystals. It consists of a sampling calorimeter with two layers: lead radiators initiate electromagnetic showers from incoming photons and electrons whilst silicon strip sensors placed after each of the two radiators measure the deposited energy and the transverse shower profiles. The geometrical coverage of the detector is $1.653 < |\eta| < 2.6$ (Fig. 1.20).

The material thickness of the preshower traversed at $\eta = 1.653$ before reaching the first sensor plane is 2 X_0. It is followed by a further radiation length before reaching the second plane. Thus about 95% of single incident photons start showering before the second sensitive plane.

Each silicon sensors has an active area of $61 \times 61\,\mathrm{mm}^2$, a thickness of 320 μm and 32 strips with a pitch of 1.9 mm. The strips are oriented vertically in the

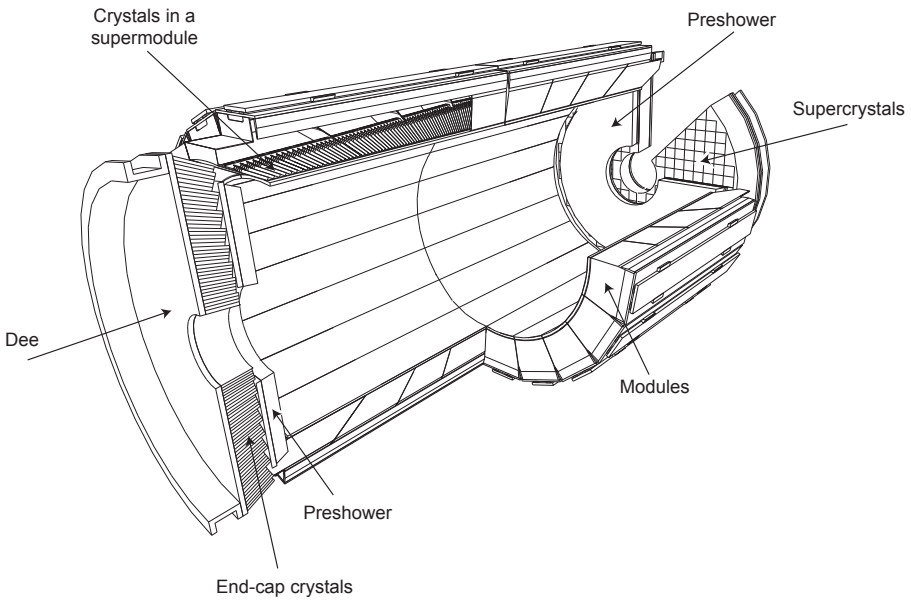

Figure 1.20: Cross section of the electromagnetic calorimeter system including the barrel part and the two end-caps complete with preshowers.

first plane and they measure the shower x-coordinate; the strip orientation is orthogonal in the second silicon plane.

CMS does not require a preshower in the barrel region because its value would be very little: the average energy of neutral pions in the barrel is smaller than in the end-caps and, consequently, the separation between the two π^0-originated photons is larger. Usually the barrel crystals granularity is sufficient to detect them.

ECAL Performance

For energies below about 500 GeV, for which shower leakage from the rear of the calorimeter starts to become significant, the energy resolution of the ECAL can be parametrized as:

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{S}{\sqrt{E}}\right)^2 + \left(\frac{N}{E}\right)^2 + C^2 \tag{1.8}$$

where S is the stochastic term, N the noise term, and C the constant term.

The stochastic term S is determined by different independent effects: a contribution of about 1.5-2% is given by event-to-event fluctuations in the lateral shower containment and one of about 2% is due to fluctuations in the

gain of photomultipliers. The preshower in the end-caps, which is not a homogeneous calorimeter, contributes with another 5% share. The Noise term N is dominated by electronics and digitization noise and it is around 120-150 MeV for measurements done with 3×3 crystal matrices. The constant term C receives contributions from the non-uniformity of the longitudinal light collection (less than 0.3% contribution) and from the uncertainty on the intercalibration of crystals.

The resolution the ECAL barrel was measured in a test beam [45, 46] and found to be consistent with design expectations (Fig.1.21).

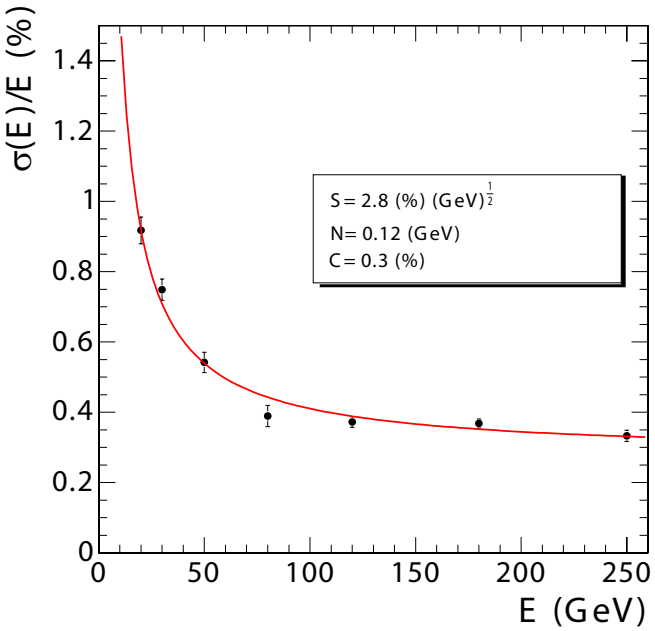

Figure 1.21: Energy resolution of the electromagnetic calorimeter, $\sigma(E)/E$, as a function of electron energy. Results obtained in a beam test.

1.3.5 Hadron calorimeter

The hadron calorimeters (HCAL) aim primarily to measure the energy of hadrons in jets. In addition, due to its hermetic closure up to high value of pseudorapidity, the system of hadron calorimeters can infer the energy of weakly interacting particles (such as neutrinos or neutralinos) from the balance of the energy of all the other interacting particles. A fundamental ingredient in the measurement of the missing transverse energy is the fact that, before the collision, the momentum of the the two initial protons system does not have a transverse component.

The principal components of the hadron calorimeter system are a barrel and two end-caps structures located in the region between the outer extent of the electromagnetic calorimeter (R=1.77 m) and the inner extent of the magnet coil (R=2.95 m). To fully absorb hadronic showers in the central section around $|\eta| = 0$, an additional calorimeter (HO) is placed immediately outside the solenoid. Moreover a forward hadron calorimeter is placed, on both side, at 11.2 m from the interaction point to extend the pseudorapidity coverage down to $|\eta| = 5.2$

Hadron Barrel and Hadron End-caps

The Hadron Barrel calorimeter (HB) is a sampling calorimeter, made of brass absorber tiles alternated with plastic scintillator, that covers the pseudorapidity range $|\eta| < 1.3$. It consists of two half-barrel sections (HB+ and HB−), each one containing 18 identical azimuthal wedges. Every wedge is segmented into four ϕ sectors, being the azimuthal aperture of each sector $\Delta\phi = 2\pi/18/4 = 0.087$.

The plates of passive material are bolted together in a staggered geometry resulting in a configuration that contains no projective dead material for the full radial extent of a wedge. The active scintillator is divided into 16 η sectors for each half-barrel (Fig.1.22), resulting in a final segmentation of $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ HB towers that match geometrically an ECAL trigger tower of 5×5 crystals (Sec. 1.3.4).

The total absorber thickness at $|\eta| = 0$ is 5.82 interaction lengths. The effective thickness increases with polar angle as $1/\sin\theta$, resulting in 10.6 λ_I at $|\eta| = 1.3$

Most of the HB towers have a single longitudinal read-out; only the 2 sets of towers closest to the end-cap transition region (the 15th and 16th η sectors) have a minimal longitudinal segmentation, i.e. two group of scintillators are read separately (Fig.1.22).

Similar to HB, the hadron calorimeter end-caps (HE) consist of brass absorbers alternated with trapezoidal plastic scintillators. The pseudorapidity region which is covered is $1.3 < |\eta| < 3$ and the sampling granularity is $\Delta\eta \times \Delta\phi = 0.087 \times 0.087$ for $|\eta| < 1.6$ and $\Delta\eta \times \Delta\phi = 0.17 \times 0.17$ for $|\eta| > 1.6$. The total length of the calorimeter, including electromagnetic crystals, is about 10 interaction lengths. The design of the absorber is driven by the need to minimize the cracks between HB and HE, rather than optimize single-particle energy resolution: the precision of HE in reconstructing jets energies is limited by pileup and magnetic field effects [47, 48]. Unlike HB, most of the HE towers are longitudinally segmented, the choice being motivated by the different radiation environment. Indeed, to restore energy

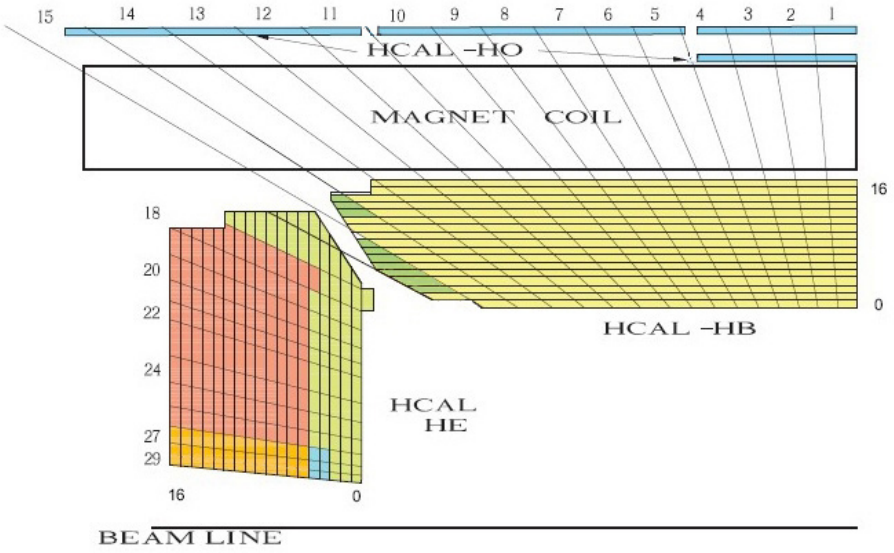

Figure 1.22: Longitudinal view (top-left quadrant) of the hadron calorimeter system, excluding the hadron forward detector. The drawing illustrates also the longitudinal segmentation implemented for most of the hadron towers in the end-cap and two towers in the barrel.

resolution, separate corrections to the calibration coefficients of the scintillators can be applied to compensate the effects of radiation. The two set of towers nearest to the beam line have three longitudinal divisions, the other have two longitudinal readouts (Fig. 1.22).

Both HB and HE have also an important function in CMS as support structures. Each pair of consecutive wedges, among the 18 ones that form the barrel HCAL, support a 20° ϕ-sector of the barrel ECAL. Similarly the two electromagnetic end-cap calorimeters are fixed to the two hadronic end-caps (Fig. 1.23). A pair of mechanical brackets per z-side transfer the weight of the whole inner tracker system to the HCAL barrel. Similarly the preshowers are connected to the HCAL end-caps.

Outer Hadron Calorimeter

In the central pseudorapidity region, the combined stopping power of EB plus HB does not provide sufficient containment for highly energetic hadronic showers. To ensure an adequate sampling depth for $|\eta| < 1.3$, the hadron calorimeter is extended outside the solenoid with the Outer Hadron calorimeter (HO).

In the barrel region, the return yoke of the magnetic field is divided, along

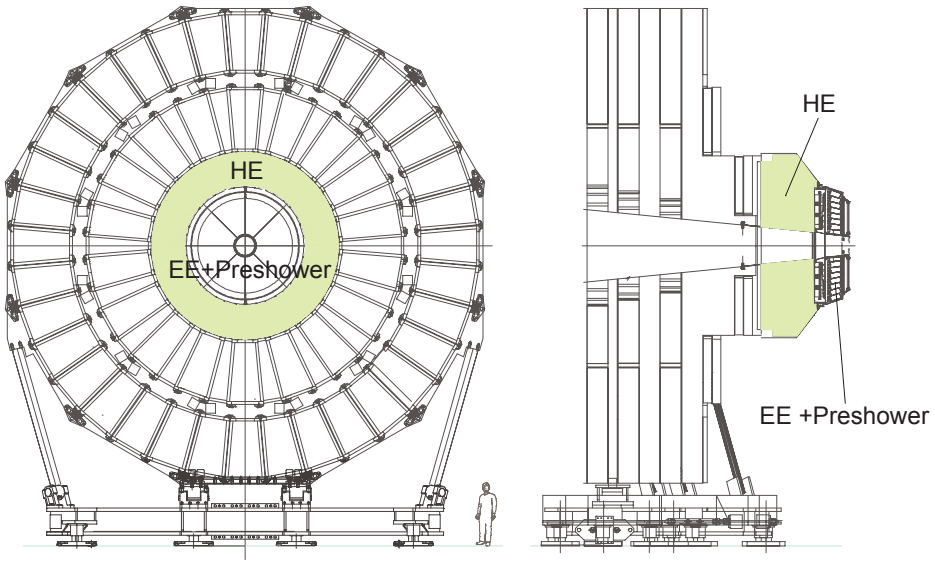

Figure 1.23: Transverse (left) and longitudinal (right) view of a CMS end-cap. The hadron calorimeter is attached to the innermost disk of the return-flux yoke and supports the electromagnetic end-cap and its preshower.

the z-axis, into five 2.536-meter-long wheels (see Fig. 1.12), each wheel being composed of 12 ϕ sectors (see Fig. 1.8). The HO is placed as the first sensitive layer in each wheel of the yoke, following closely the 5-fold-z and 12-fold-ϕ arrangement of the muon DTs (Sec. 1.3.3).

The HO utilizes a 19.5-cm-thick iron absorber and also the solenoid coil as an additional absorber equal to $1.4/\sin\theta$ interaction lengths. The total depth of the EB-HB-HO calorimeter system is thus extended to a minimum of 11.8 λ_I, except at the barrel-endcap boundary region. Plastic scintillator tiles are positioned after the iron absorber in front of the first layer of DT detectors. The sizes and positions of the tiles in HO are arranged to roughly map the HB towers with granularity 0.087 x 0.087 in η-ϕ.

Hadron Forward

The Hadron Forward calorimeter (HF) is composed of two cylindrical steel structures, located at 11.2 m from the interaction point, that cover the pseudorapidity range of $3.0 \lesssim |\eta| \lesssim 5.0$ on both sides of CMS (Fig. 1.5). The detectors have an outer radius of 130.0 cm and a small cylindrical hole of radius 12.5 cm that allows the passage of the beam pipe. According to simulations, the forward calorimeter is expected to experience unprecedented particle fluxes. On average, 760 GeV per proton-proton interaction is de-

posited into the two forward calorimeter structures, compared to only 100 GeV for the rest of the detector.

This hostile environment presents a considerable challenge to calorimetry, and the design of the HF calorimeter was first and foremost guided by the necessity to survive in these harsh conditions for about a decade. The HF is a sampling calorimeter: the absorber consist of 5 mm thick plates of steel, for a full depth of 165 cm (10 λ_I); the active medium consists in quartz fibers that were chosen for their radiation hardness. The signal is generated when charged particles above the Cherenkov threshold generate Cherenkov light [49].

The fibres are routed along the z-direction and are bundled to form $\Delta\eta \times \Delta\phi = 0.17 \times 0.17$ towers. Half of the fibers run over the full depth of the absorber, while the other half starts at a depth of 22 cm from the front of the detector. These two sets of fibers are read out separately in order to distinguish showers generated by electrons and photons (which deposit a large fraction of their energy in the first 22 cm) from those generated by hadrons (which on average produce nearly equal signals in both calorimeter segments).

The Cerenkov process seen by the phototubes in the forward calorimeter produce a signal pulse which is only 10 ns wide. Hence, thanks to its very fast time response, the HF is subject only to in-time pile-up and can cope with the high particle activity of the forward region of CMS.

Performance

Like ECAL, the energy resolution of the HCAL can be parametrized as the sum of three terms (eq. 1.8). Figure 1.24 shows the resolution on the measurement of the jet transverse energy as evaluated on simulated events. The stochastic term S, dominated by fluctuations in the energy deposited inside the passive absorbers, is of the order of 125%; the noise term N, dominated by pile-up and underlying event effects, is around 50-60%; the constant term C is less than 3%. Although the electromagnetic crystals are calibrated on photons energies, they absorb a substantial amount of the jet energy carried by charged pions. For this reason, the raw reconstructed jet energy is usually underestimated by about 20%. Simulation-based corrections are applied to reduce this bias.

1.3.6 Trigger and data acquisition

A single LHC collision "event" corresponds approximately to 1.5 MBytes of zero-suppressed data from the CMS read-out system. Given that the nominal collision rate of LHC is 40 MHz, in order to store all the events

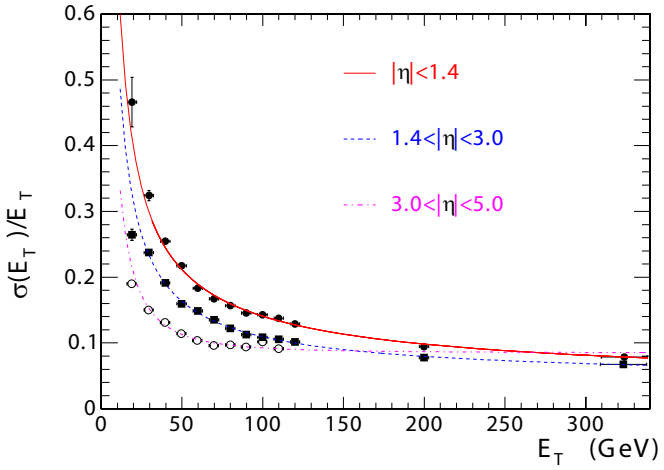

Figure 1.24: Jet transverse-energy resolution as a function of the jet transverse energy for barrel jets ($|\eta| < 1.4$), endcap jets ($1.4 < |\eta| < 3.0$), and very forward jets ($3.0 < |\eta| < 5.0$). Results obtained from simulation.

CMS would need a bandwidth of approximately 40 TBytes/s. In reality the rate sustainable by the offline storage system of CMS is about 5 order of magnitude smaller. The task of the trigger system is therefore to filter the huge amount of data received from the detector read-out and return to the storage and analysis systems only the most interesting events.

In CMS, the online filtering process is implemented in two steps:

- The Level-1 (L1) Trigger consists of custom-designed, largely programmable electronics that performs a first selection of the events on the base of coarse data received from a subset of the CMS subdetectors. The high-resolution data is hold in pipelined memories in the front-end electronics. The output rate of the L1 Trigger is limited to 100 kHz, which corresponds to a rate reduction factor of about 10^3.

- The High-Level Trigger (HLT) consists of a software system implemented in a filter farm of about one thousand commercial processors. The HLT has access to full-granularity data from all the CMS subdetectors and runs faster versions of the same algorithms used for the offline analysis. In order to match the input rate of the storage system, the HLT has to reduce the event rate approximately by another factor 10^3.

In the following, the major features of the L1 Trigger and of the HLT and Data Acquisition system (DAQ) are described.

L1 Trigger

The L1 Trigger uses coarse data from the muon stations and the calorimeter systems. The volume of data produced by the inner tracking read-out is too large to be transferred to and processed by the L1 electronics, and therefore is used only by the HLT. The L1 Trigger consists of local, regional and global components.

- The Local Triggers, also called Trigger Primitive Generators (TPG), are based on energy deposits in calorimeter trigger towers and track segments or hit patterns in muon chambers.
- Regional Triggers combine local information and use pattern logic to determine ranked trigger objects such as electron or muon candidates in limited spatial regions. The rank is usually determined by energy or momentum and quality factors.
- The Global Calorimeter and Global Muon Triggers determine the highest-rank calorimeter and muon objects across the entire detector and transfer them to the Global Trigger, the top entity of the Level-1 hierarchy (Fig. 1.25). The latter takes the decision to reject an event or to accept it on the base of algorithms that combine Muon and Calorimeter information. In addition to regional objects likes muons and jets, also global objects are used. Example of global quantities are the transverse missing E_T or the scalar sum of all calorimetric energy deposits.

The Global Trigger decision takes into account the status of the sub-detector front-ends that buffer the full granularity data and the readiness of the DAQ to receive it. Both conditions are communicated by the Trigger Control System (TCS). Similarly, the Level-1 Accept (L1A) decision is communicated to the sub-detectors through the Timing,Trigger and Control (TTC) system. The total deadtime estimated at the maximum L1 output rate of 100 kHz is estimated to be below 1%.

Due to the limited depth of the front-end pipelines that store the full granularity data, the allowed L1 Trigger latency between a given bunch crossing and the distribution of the trigger decision is 3.2 μs. To reduce to about 2 μs the time necessary to transfer the trigger primitives to the L1 electronics and send the L1A signal back to the detector, the L1 Trigger system is housed in an underground room located at a distance of approximately 90 m from the experimental cavern. Similarly, most of the trigger boards responsible for the local trigger reconstruction are located directly on the sub-detector read-outs. The time interval that remains available to the Trigger system for taking the L1 decision is about 1μs.

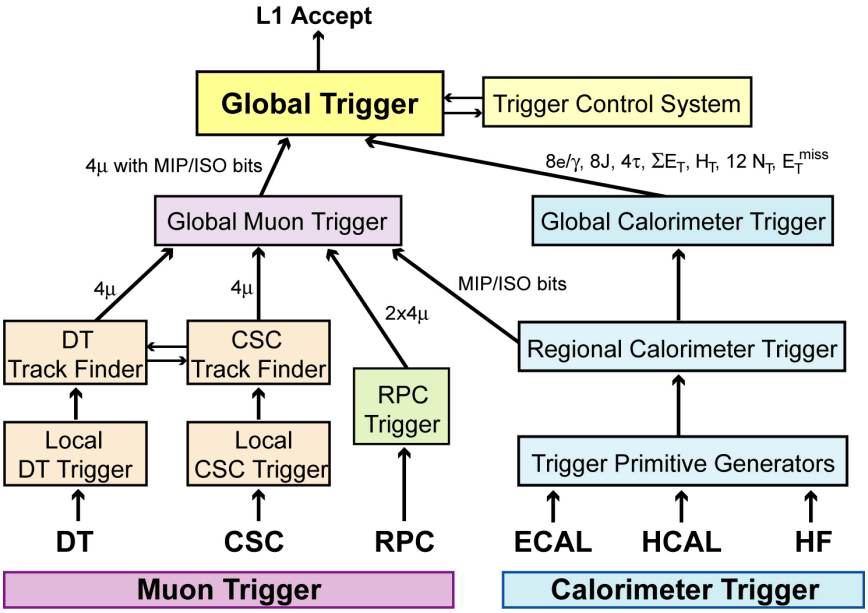

Figure 1.25: Architecture of the L1 trigger system.

Data Acquisition

The Front-End Systems (FES) of various CMS sub-detectors store data continuously in 40-MHz pipelined buffers. Upon arrival of a synchronous L1 trigger (3.2 μs latency) via the TTC system, the corresponding data are extracted from the front-end buffers and pushed into the DAQ system by the Front-End Drivers (FEDs) (Fig. 1.26). The design of the FED is sub-detector specific, however, a common interface from the FED to the central DAQ system has been implemented. The FED encapsulates the data received from the front-end electronics in a common data structure by adding a header and a trailer that mark the beginning and the end of an event fragment. Each event consists in 512 fragments with an approximate size of 2 KByte.

The event builder assembles the event fragments belonging to the same L1-accepted event from all FEDs into a complete event and transmits it to one of the Filter Units (FU) for the final HLT processing. The event builder is split into two stages: the FED-builder and the RU-builder (Fig. 1.26), The advantage of a two-stage event building is that the full size system can be deployed progressively, following the evolution of the LHC luminosity.

The FED-builder is responsible for transporting all the 512 event fragments from the underground electronic room (USC) to the surface building (SCX), where the fragments are merged into 72 super-fragments with an average size of 16 KByte. The super-fragments are stored in large buffers inside Read-

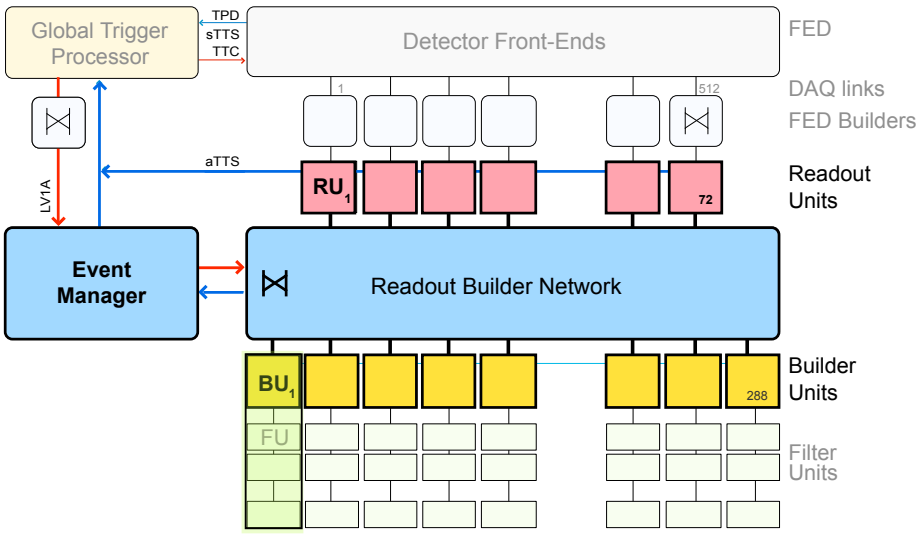

Figure 1.26: Architecture of the High-Level Trigger and Data Acquisition system.

out Units (RU), waiting for the second stage of event building. By means of a Gigabit ethernet network, the RU-builder sends all the super-fragments corresponding to one event to a single Builder Units (BUs), where the whole full-granularity RAW event is assembled for the first time and formatted into C++ objects. Both the RU and the BU nodes are commercial multi-core PCs running custom CMS software on a Linux operative system.

The Filter Unit is run on the same PC of the Builder Unit that assembled the event. The HLT filtering consists in a sequence of C++ modules that analyze the RAW event, produce and attach higher-level objects to the event data structure and eventually accept or reject the trigger candidate. When a decision is reached, the accepted events, comprising raw data and the additional information produced by the HLT, are transfered to the Storage Manager (SM). During operation, trigger thresholds and pre-scales will be optimized to fully utilize the available DAQ and HLT throughput capacity.

Chapter 2

The Silicon Tracker of CMS

This chapter analyzes the silicon tracking system in details. In Section 2.1 the physics requirements that drove the design of the Tracker are discussed, while Section 2.2.1 describes the final layout of the system. The last Section 2.3 explains how the silicon strip modules have been tested and qualified prior to their installation on the detector.

2.1 Physics requirements and design goals

The inner tracking system based on silicon sensors plays an essential role for a general-purpose detector like CMS, which is designed to address a wide range of physics processes. Indeed, the charged particle trajectories that are reconstructed within the Tracker constitute a vital ingredient in the measurement of almost all the observables used for physics analyses [50]:

- In synergy with the muon system and the calorimeters, the tracking system is crucial to identify and precisely reconstruct muons and electrons. High-p_t leptons are expected to be of primary importance at LHC since leptonic decays of gauge bosons can provide an especially clean experimental signal for a wide range of physics phenomena, including the Higgs boson production [51].

- Along with the calorimeters, the Tracker is used to estimate the isolation of charged particles or collimated jets like those produced by energetic taus. The reliable evaluation of particle isolation is another important aspect in the detection of gauge bosons that decay leptonically.

- After jets are reconstructed in the calorimeters, the innermost detectors of the tracking system are fundamental to recognize those jets

that originated from heavy flavor hadrons. The identification is indeed performed through the reconstruction of secondary vertices and the measurement of the impact parameter of the charged particles contained inside the jets. B-hadrons and the corresponding jets are essential tags for top quark physics and are expected to be critical signatures for a broad spectrum of new processes beyond the Standard Model. Besides, detailed studies on B-physics may shed light on the mechanism responsible for CP-violation.

- With an efficient and precise reconstruction of charged particles momenta in a wide range of energies and directions, the tracking system is used to correct the energy and direction of the calorimetric jets. In the same way, the Tracker measurements are employed to improve the estimate of the missing E_T. A particle flow approach that combines tracking and calorimetric information seems particularly promising for the reconstruction of hadronic τ-jets [52].

- The reconstruction with the inner tracking system of electrons from Z and W decays or from photon conversions, allows the in-situ monitoring and calibration of single ECAL cells. As discussed in Section 1.3.4, a steady calibration of the ECAL crystals is fundamental to maintain the high resolution of the electromagnetic calorimeter.

- Most of the Tracker information which is used for the offline analysis is also heavily exploited by the algorithms that run in the High Level Trigger [53].

The physics program aimed by CMS sets the minimum level of performance that the Tracker has to match. The required efficiency and resolution of the tracking system was evaluated since the early stages of the detector design [2, 3]:

1. A transverse momentum resolution $\Delta p_T/p_T \approx 0.1\text{-}0.15 \cdot p_T$ (p_T in TeV/c) for charged particles in the central region, degrading to $\Delta p_T/p_T \approx 0.6 \cdot p_T$ for $|\eta|$ approaching 2.5.

2. Correct charge assignment for leptons of energy up to 2 TeV.

3. Reconstruction efficiency higher than 95% for isolated charged particles of momentum above 1-2 GeV/c and efficiency of the order of 90% for particles in the same momentum range, but collimated inside jets.

4. Measurement of the longitudinal impact parameter with a minimal resolution of 2 mm to associate particles to the pp collision vertices they originated from. Resolution around 20 μm for the transverse impact parameter and around 100 μm on the longitudinal one are required to implement valuable b-tagging.

5. Detector occupancy below 1% to facilitate the pattern recognition in the congested environment of LHC collisions.

6. At least in the barrel region of the Tracker, the amount of material budget should be less than 0.15-0.30 radiation lengths. This requirement is necessary to avoid that the reconstruction of photons and electrons is spoiled because of conversions and bremsstrahlung radiation. In the same way, an excessive material budget deteriorates the resolution on the trajectory parameters of low energy particles because of multiple scattering.

7. At the LHC design luminosity of $\mathcal{L} = 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, about 1000 particles from more than 20 overlapping interactions are expected to traverse the tracker volume every 25 ns. The tracking system is required to have a sufficiently fast time response in order to associate correctly the detected signals to the corresponding bunch crossings.

In addition to the previous list, it is imperative that the components of the Tracking system are sufficiently radiation tolerant to maintain, without major interventions, the design performance for the full extent of the CMS operations at LHC.

2.2 Tracker Design

A tracker system entirely based on silicon detector technology can match all the requirements on resolution, speed and radiation hardness [54] described in the previous section. Nevertheless, while a high number of readout channels is desirable to increase the granularity of the system, it is a major challenge to control the amount of material used for sensors, electronics, support structures and services. In the following the final layout of the silicon tracking system is described.

2.2.1 Tracker Layout

The tracking system of CMS contains a very high granularity pixel detector in the innermost region around the collision point where the particle density is higher and the best detector resolution is desirable. The intermediary and outer regions of the Tracker consist of silicon strip sensors with varying strip pitches and lengths to obtain the best compromise among measurement accuracy, material budget and cost. Three pixel barrel layers (PXB) are installed at radii between 4.4 cm and 10.2 cm and a silicon strip tracker with 10 barrel detection layers (TIB and TOB) extends outwards to a radius of

1.1 m. Both systems are completed on each side of the barrel by two end-caps consisting in 2 disks of pixels detectors (PXF) and 3 small (TID) plus 9 big (TEC) disks of strip detectors. Altogether the acceptance of the tracker extends up to a pseudorapidity of about 2.5 (Fig. 2.1).

With 15 148 strip (1 440 pixel) silicon modules, 9.4 M (66 M) read-out channels and 198 m^2 (1 m^2) of active area, the CMS tracker is the first apparatus, built for a collider experiment, to employ only silicon sensors over such a wide area. This novel approach was made possible by the reduction in the cost of larger (6 inches) silicon wafers, the use of industry standards for the front-end and read-out chips and the heavy automation of module assembly and wire bonding procedures [54].

This chapter gives the main features of the silicon detector modules while their exact arrangement on pixel and strip sub-detectors is given in Appendix A.

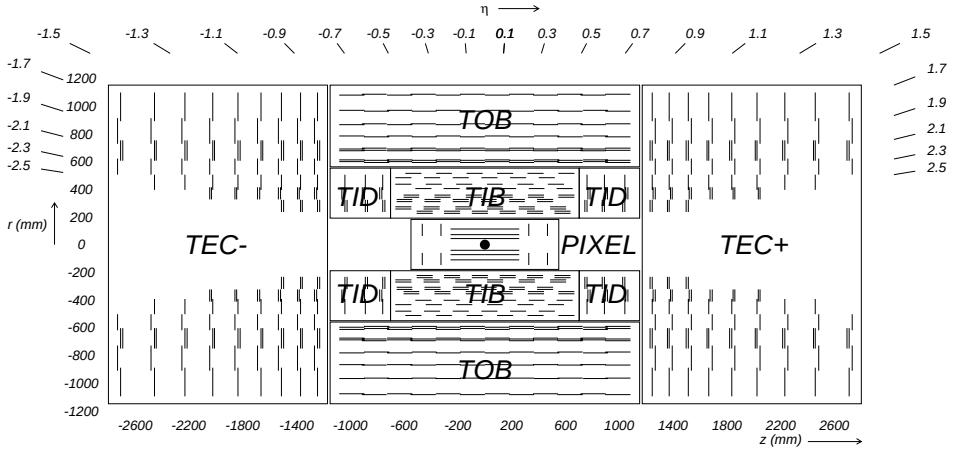

Figure 2.1: Longitudinal view of the all-silicon tracking system of CMS. The inner pixel detector consists of 3 cylindrical layers completed on both sides by 2 disks. The silicon strip detector includes 4 inner cylindrical layers (TIB) and 6 outer layers (TOB) installed in the barrel region, which are extended longitudinally by 3×2 small disks (TID) and 9×2 bigger disks (TEC). Single (pair of) lines denote single-sided (doble-sided) strip modules.

Pixel detectors

All the pixel sensors of CMS have a cell size of 100×150 μm^2 (Fig. 2.2 left), which was driven by the desired impact parameter resolution and limits the

occupancy of the detector[1] to the order of 10^{-4} per LHC bunch crossing.

The pixel cells are $n+$ implants on a n-substrate and the charge carriers that produce the signal are electrons. In spite of the higher costs due to double sided processing [55], this design is desirable because the use of electrons as charge carrier ensures a high signal at moderate bias voltage even after high hadron fluences. Besides, as opposed to holes charge carriers, electrons undergo a Lorentz drift which is three times bigger due to their higher mobility. For sensors placed in barrel layers, the depletion electric field is perpendicular to the axial magnetic field of the solenoid and the Lorentz angle is maximum: even if a particle traverses the detector surface perpendicularly, the charge signal is spread over more than one pixel allowing the evaluation of the cluster position by interpolation (barycenter method). Therefore, instead of a binary read-out, the detector front-end and read-out electronics preserve the complete analog signal recorded by all cells of pixel clusters. Within forward detectors, the magnetic and drift field directions are almost parallel and the effect of Lorentz drift is much smaller. Hence, the charge sharing is induced by implementing a turbine-like geometry in which the sensors are tilted by 20° (Appendix A.0.2).

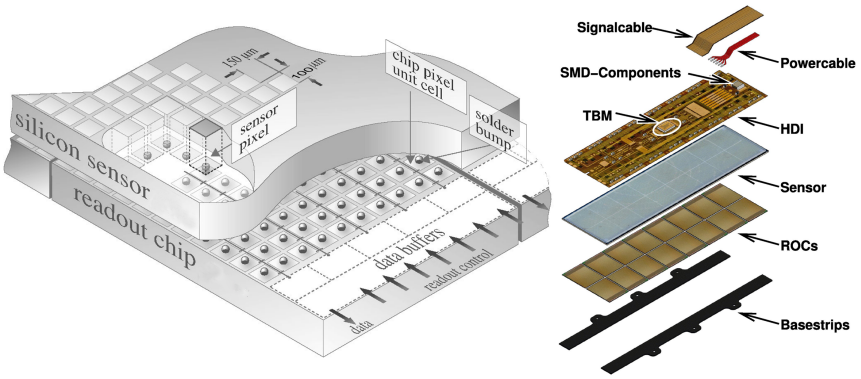

Figure 2.2: Detail of the interface between the matrix of pixel cells and the read-out chip (left); exploded view of a detector as those installed in the barrel layers of the pixel system (right).

Due to the small dimension of the pixels and to the analog read-out that allows charge interpolation, barrel (forward) detectors achieve a spatial resolution in the range of 15-20 μm for both ϕ and z (r) coordinates. However, the depletion depth of pixel sensors reduces with the integrated luminosity because of the radiation damage to the silicon sensor bulk. For the same reason, the bias voltage has to be increased progressively to maintain an adequate level of signal. Both effects imply a reduction of the charge-sharing

[1]The occupancy is defined as the ratio of the number of detector channels above the noise level to the total number of channels.

and a consequent gradual degradation of the sensors spatial resolution [56].

Strip detectors

Unlike the pixel system, less stringent requirements allow the strip tracker to cover a large area by means of cheaper sensors produced with single-sided processing. The p-on-n sensitive elements include strips on the front side as $p+$ implantations into the n type bulk; on the back side, a uniform $n+$ implantation is connected to a positive voltage and provides the bias up to about 500 V [57, 58].

Each single $p+$ strip is covered by an aluminum layer from which it is electrically insulated by means of a silicon oxide and nitride multi-stratum. This integrated capacitor allows for an AC coupling of the signals between the silicon strips and the read-out electronics connected to the aluminum strips. At the same time, the front-end electronics is protected from the high leakage currents that manifest after irradiation of sensors. Every metal strip has two bond pads on each end, which are used to make a wire bond connection to the read-out chip. In case of the daisy-chained sensors, wire bonds connect the pair of sensors that are used in a single detector module (Fig. 2.3).

All barrel layers support rectangular-shaped sensors with their strips running parallel to the z-axis, while forward disks accommodate trapezoidal sensors with radial strips. Hence, the ϕ coordinate is measured by both types of layers. Some radii of the barrel and end-caps systems are instrumented with double-sided modules. These are pairs of detectors that are installed back-to-back at small stereo angle and that provide a rough measurement of a second coordinate: z in the barrel region and r in the forward one. Preferable to a 90° configuration, the small stereo angle readout offers the required resolution with a minimum of particle hit ambiguity.

The strip pitch varies between 80 μm and 200 μm as a function of the radius of the sensor location. The sensors on the innermost and outermost layers have the smaller inter-strip distance, since a better resolution is recommended to properly extrapolate the track parameters from the Silicon Strip Tracker (SST) towards the other sub-detectors of CMS.

Like in the pixel system, the channels of the SST are read out analogically and the positions of strip clusters are evaluated exploiting charge-sharing. A resolution of 23 μm on the measurement of the ϕ-coordinate is achieved in the innermost strip layers consisting of sensors with 80 μm pitch. The corresponding single point resolution of the stereo measurement is around 230 μm. For particles in the range $|\eta| < 2.4$, the Silicon Strip Tracker layout ensures at least 9 measurements, where approximately four of them are two-dimensional measurements on double-sided modules.

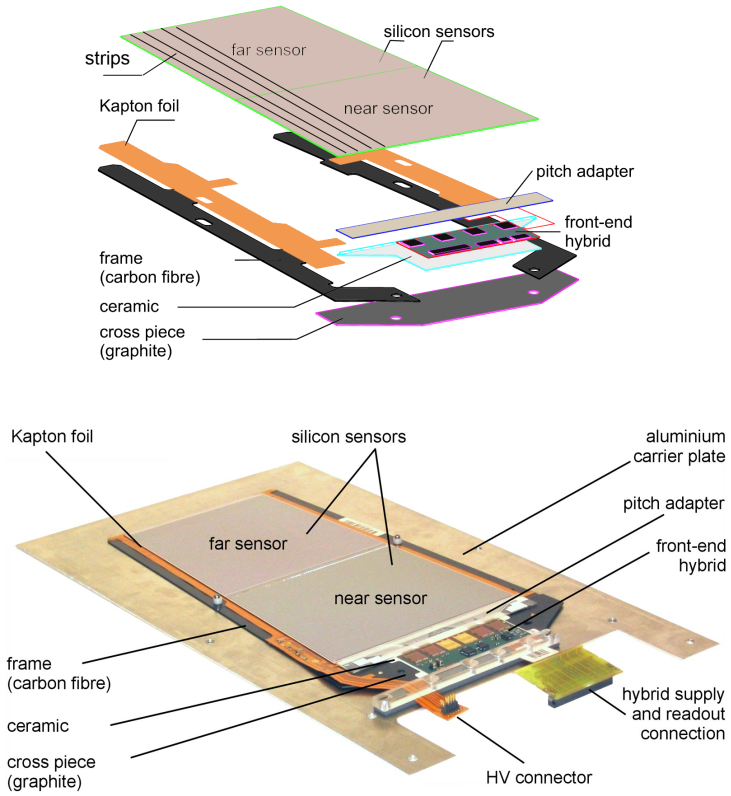

Figure 2.3: Exploded view of daisy-chained strip detector as those used in the forward disks of the Strip Tracking system (upper); picture of a real detector of the same kind (lower). The detector consists of two trapezoidal strip sensors that are wire-bonded to effectively double the strip length. Four read-out chips amplify and process the signal from 4×128 strip channels.

In spite of the small strip pitch, the inner region of radii 20 cm $< r <$ 55 cm contains sensors with a typical strip length of approximately 12 cm, leading to an occupancy of up to 2-3% per strip module and per LHC bunch crossing. To reduce the rate of fired channels and to facilitate the matching of measurements on double-sided modules, the detectors installed in the inner region of the barrel are tilted to compensate for the Lorentz drift (Fig. A.5). Unlike the layers closer to the collision point, in the outer region of radii 55 cm $< r <$ 118 cm the flux of particles is smaller so that both the sensor pitch and the strip length can be increased without compromising the average occupancy of the detectors. Given the large surface that needs to be instrumented, the use of bigger-sized modules is critical to limit the number of read-out channels and to reduce the material budget of the consequent electronics and cooling circuits. Since the electronic noise is a linear func-

tion of the strip length, in order to maintain a good S/N ratio the signal is enhanced using thicker silicon sensors with a deeper depletion region (see Appendix A).

2.2.2 Radiation hardness

Radiation affects both pixel and strip silicon sensors mainly through bulk damage, *i.e.* modifications to the silicon crystal lattice which are caused by non-ionizing energy loss. The consequences are an increase of the leakage current, doping inversion (with a corresponding change in the depletion voltage) and creation of "trapping centers":

- The increased detector leakage current can lead to a dangerous positive feedback of the self heating of the silicon sensor and to the exponential dependence of the leakage current on temperature. This effect, known as *thermal runaway*, is avoided through the cooling of the silicon sensors. The whole tracker volume is indeed operated at -10°C.
- In order to have an efficient charge collection, the creation of trapping centers and the type-inversion are compensated increasing the bias voltage and over-depleting the sensors. The radiation damages to the bulk limit the sensor lifetime because eventually the breakdown voltage is reached.

A second type of radiation damage is called *reverse annealing* and requires to keep the silicon sensors permanently well below 0°C except for short maintenance periods. This effect is caused by the interaction of radiation induced defects in the silicon sensors which can lead to more serious damage and to an even stronger change in depletion voltage with fluence [59] .

A third kind of radiation damage does not directly affect the sensors, but influences the read-out electronics. Ionizing particles can produce charge in the electronic circuitry and change the state of memory cells disturbing, or even stopping, the proper activity of the read-out chips. Usually the chips are recovered by means of a reset command and this kind of radiation damage, usually known as Single Event Upset (SEU), is therefore transient. The read-out chips utilized in the CMS tracker are produced in 0.25 μm CMOS technology, which is less affected by SEUs due to the thin gate oxide.

2.2.3 Tracker read-out

Pixel read-out and control system

The Read-Out Chip (ROC) of pixel detectors is a custom ASIC (Application-Specific Integrated Circuit) that amplifies and buffers the signals from a ma-

trix of 52 × 80 pixel cells [60]. To reduce the amount of data transfered to the DAQ system, this chip also performs zero-suppression using thresholds that can be adjusted differently for each pixel.

The TBM [56] (Token Bit Manager) chips control the read-out of groups of ROCs. A TBM is located directly on each pixel barrel module (Fig.2.2) or in the proximity of a group of pixel forward sensors (see Sec. A.0.2). Each TBM, and the group of ROCs it controls, are connected to a single analog optical link over which the data is sent to the underground electronic room for digitization.

The principal functions of the TBM chip include:

- distribution of clock and Level-1 triggers to the ROCs.
- control and formatting of the ROCs read-out for each incoming L1 trigger.
- multiplexing ROCs analog output into the analog-optical converter, prior to the optical transmission towards the digitizer. The optical links of the pixel system are identical to those used in the silicon strip tracker.

Control commands are sent using a fast line to the TBM, and then forwarded to the ROCs, to refresh the pixel threshold bits and to correct rapidly the effect of SEUs. Since updating the front-end settings may disrupt data taking, small updates are synchronized to orbit gaps in the LHC.

The pixel read-out and control system [61] is completed by other three types of components. The pixel Front-End Drivers (pxFEDs) receive the optical analog data from the front-end electronics and send the digitized information to the CMS-DAQ event builder. The pixel Front-End Controllers (pFECs) send the 40 MHz clock and other fast control signals, including L1A triggers and resets, to the pixel detectors. Slow controls, like the monitoring of temperature and humidity or like the setting of low and high voltages, are conducted by standard non-pixel-specific Front-End Controller (FECs).

The pxFEDs, pFECs and FECs are all VME [62] modules located in the underground electronics room.

Strip read-out and control system

The APV25 [63] is a ASIC designed in 25 μm radiation-hard CMOS technology and used as front-end chip within SST detector modules. It is responsible for amplification, shaping and storing of signals from 128 strips of a silicon sensor. For each read-out channel, the APV includes a pre-amplifier, a 50 ns shaper and a 192-elements-deep pipeline, which samples the shaped

analog signal at the LHC frequency of 40 MHz. In case of a L1A signal, the APV can either pass the signal as sampled at the maximum of the 50 ns pulse (peak mode) or form a weighted sum of three consecutive samples (deconvolution mode). At the price of a slightly worse S/N ratio, the latter effectively reduces the shaping time to 25 ns. Such a smaller shaping time is necessary at the design LHC luminosity in order to associate the signals to the correct bunch crossing.

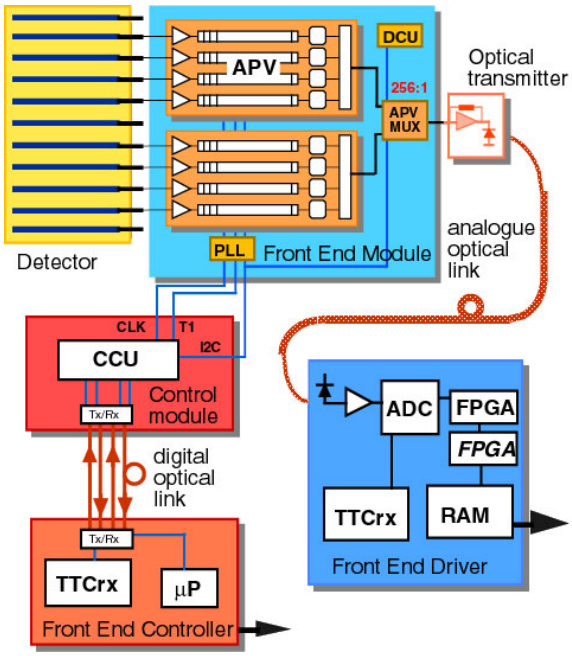

Figure 2.4: Scheme of the read-out and control system implemented for the the Silicon Strip Tracker.

Another custom ASIC, the APVMUX, is used to multiplex the data streams from two APV25 chips onto one single electrical channel. The latter is subsequently converted into an optical signal by the Analogue Opto-Hybrid (AOH) circuit [64, 65], which sits close to each detector module. Analogue optical links are used to transmit the data streams from the strip detectors to the service cavern over a distance of about 100 m. Like the pixel read-out system, the SST design employs optical links, instead of an electrical distribution scheme, since they have minimal impact on the material budget and are immune to electrical interference. Each SST module utilizes from a minimum of 2 data links (single-sided detectors with 512 strips and 4 APVs) to a maximum of 6 links (double-sided modules with 768 strips and 6 APVs on each detector).

The Front End Drivers (FEDs) [66] are the VME modules located in the underground room which receive data from optical fibers, convert the signals

into electrical levels and digitize them every 25 ns with 10 bit Analog-to-Digital Converters (ADC). The digitizer sampling-point for each fibre can be programmed independently in 1 ns steps to optimize the S/N ratio, taking into account the different delays for signals coming from different parts of the SST system. Unlike the pixel system, the zero suppression of the SST data is not performed by the front-end electronics placed on the detector modules, but is conducted by the FEDs in the counting room. This design choice allowed the simplification of APV chips with a consequent reduction in power consumption, cooling system constraints and an overall reduction of the SST material budget. Only the strip clusters surviving the zero suppression are sent from the FEDs to the central DAQ.

Clock and triggers from the global TTC are transmitted to the tracker Front End Controller (FEC) boards [67], which redistribute them, together with other fast control signals, to the strip tracker via digital optical links. FECs are VME modules located in the service cavern as close as possible to the detector in order to reduce trigger latency. The signals from each FEC are distributed to the strip modules by a local network, called *control ring*, of Communication and Control Units (CCUs) [68], which are mounted on the Tracker. A single CCU communicates with up to 16 silicon detector modules sending trigger and clock signals to the Phase-Locked Loop (PLL) [69] chip located on each module; other controls are directly transmitted from the CCU to the APVs.

A Detector Control Unit (DCU) [70] is an ASIC, mounted on each strip detector, that is used to monitor the electronics low voltages, the silicon sensor leakage current and the temperatures of different detector components. The measurements of several DCUs are collected by a control ring and sent, through a digital link, to the FEC. To implement the bi-directional digital optical connection between FECs and control rings, Digital Opto-Hybrid (DOH) circuits [71] are used to convert the optical signals to electrical ones and vice versa.

The Hybrid [72] is the multi-chip module that contains all the front-end components of each silicon strip detector (4 or 6 APVs, one APVMUX chip, one PLL and one DCU) and routes low voltages and data among them.

2.2.4 Material Budget

This section contains the estimate of the Tracker material budget based on the most recent, up to date, simulation of the detector. In its final setup, the silicon tracking system is about twice "thicker" than originally envisaged [2]. Figure 2.5 shows, as a function of pseudorapidity, the fractional radiation length X/X_0 and the fractional nuclear interaction length X/λ_0 seen by particles originating from LHC collisions and traversing the whole

Tracker volume. The amount of crossed material is minimum for particles in the barrel (X/X_0=0.4-0.6 and X/λ_0=0.12-0.2) and maximum for particles moving over the barrel-endcap transition region, where X/X_0 and X/λ_0 peak at 1.8 and 0.55 respectively.

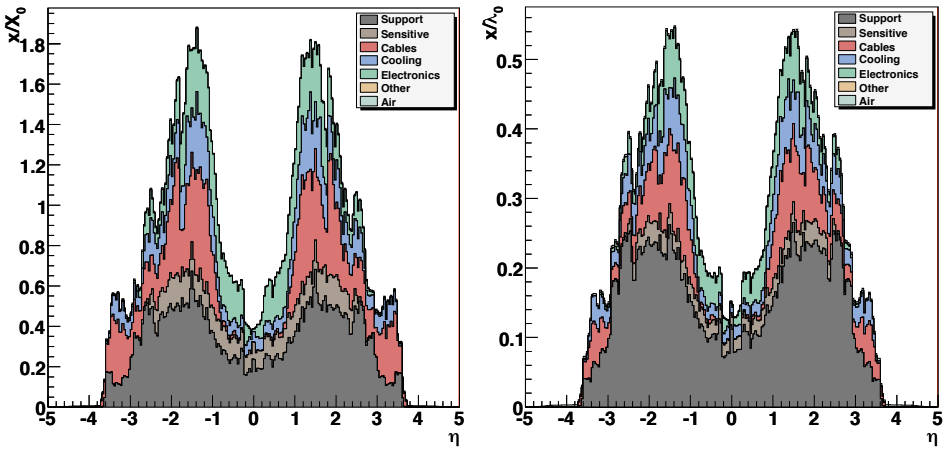

Figure 2.5: Material budget seen by a particle produced in the center of CMS and crossing the whole volume of the Tracker. The material is expressed in terms of fractional radiation lengths (left) and fractional interaction lengths (right) as a function of the particle pseudorapidity.

Unlike the silicon sensors, which contribute to the material budget uniformly in the range $|\eta| < 2.5$, the material associated to the support structures varies by a factor two between the barrel and the end-caps. Also, the material corresponding to cables and cooling is almost negligible in the barrel, but is responsible for about 50% of the total material budget in the transition region. The reason of this inhomogeneity can be understood from Figure 2.6, which shows the distribution of the Tracker material in the R-z plane. In order to avoid dead zones in the cylindric layers of the Tracker, all the services of the barrel are routed through the two gaps between the barrel and the end-caps. As a consequence, a particle with $|\eta|$ between 1.1 and 1.4 cross up to three times the same cables and cooling pipes.

In the barrel-endcap transition region and in the two end-caps, electrons lose between 60% and 85% of their energy by bremsstrahlung radiation before they reach the electromagnetic calorimeter. In the same η range, about 30 to 40% of hadrons interact inelastically before they extend to the outer layers of the tracking system. As Chapter 3 describes in detail, both effects limit the precision with which the momenta of electrons and hadrons can be measured.

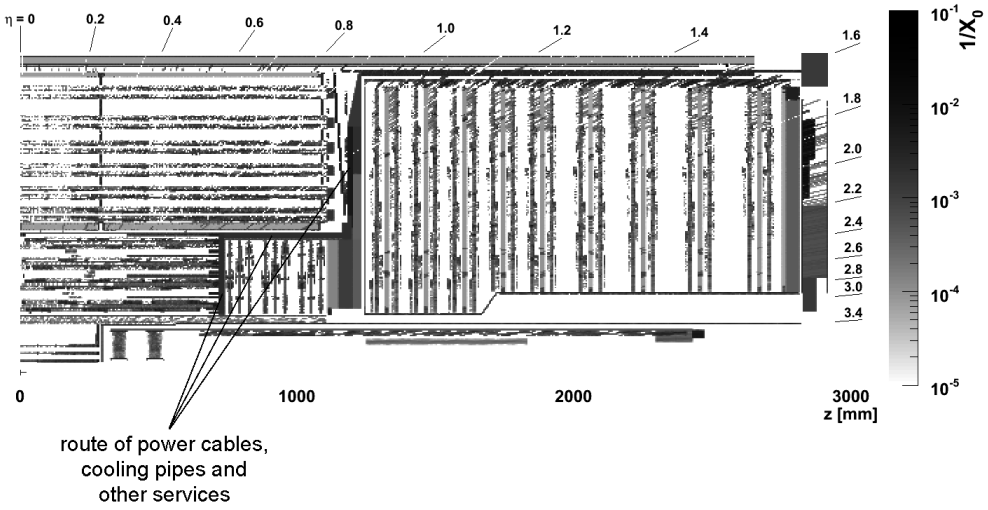

Figure 2.6: Longitudinal R-z view the silicon Tracker and material budget expressed in terms of inverse radiation length. The darker area within the gab between the barrel and the end-cap is due mostly to cables, cooling pipes and other services.

2.3 Test and Qualification of silicon strip modules

This section illustrates the ARC[73] test procedure employed on TIB and TID detector modules. The module test activity program aimed at selecting extremely high quality detectors, with a number of defects on strips and corresponding readout channels not greater than 2%.

The modules which passed the qualication procedure were ready for installation on the final substructures of the TIB and TID. Given the specifics of the mechanical construction of the inner barrel and inner disks, it was difficult and risky to replace the modules once they were integrated onto the final mechanical structures. Therefore, each test procedure was implemented and optimized in order to spot production problems as efficiently as possible.

2.3.1 Module defects definition

The different defects that can affect the several components of a silicon strip detector module are described in the following:

- **Mechanical Defects and Handling Damages** All kinds of defects due to accidents or mistakes in the production of the module fall into

this category. The possible defects include touched or broken wire bonds, scratches, glue excess preventing assembling or bonding or installation.

- **Noisy Channels.** Readout channels that exhibit noise which is consistently higher than the expected average for a standard module.
- **Dead Channels.** Readout channels which do not respond to calibration signals.
- **Open Channels.** The opens are readout channels that are not connected to the corresponding sensor strip due to a missing wire bond between the pitch adapter and the sensor (or the APV25), or an interrupted PA track.
- **Short Circuit Channels.** The shorts are two or more neighbouring readout channels that are shorted together due to defects such as scratches on the silicon sensor surface, unwanted metal contact between strips, problematic wire-bonds or PA tracks.
- **Metal Strip Breaks.** Interruptions in the metallization of the strips, which can cause a module efficiency loss.
- **ASIC Defects** The ASICs on the modules can have electronic defects which usually show as problems in communication, addressing or readout.
- **APV Defects.** Defects which are strictly related to the APV25 behaviour.
- **Silicon Crystal Bias Current Defects.** There is a class of defects which may usually cause a large increase in the current when applying the voltage bias.
- **Pinholes.** Anomalous connections between the p+ implant and the metal strip above the implant; as a consequence the coupling between the implant strip and the APV25 input is no longer purely capacitive, but it has a DC component. Pinhole identication is very important, since undetected pinholes may, in the worst case, prevent the normal functioning of an entire APV25.

2.3.2 Measurement Denitions

The basic measurements and inspections used to determine if a module could be considered as suitable for installation in the CMS Tracker are described in the following list.

- **Optical inspection.** A detailed visual inspection by eye followed by a deeper inspection with a microscope. The goal of this analysis was to spot damages to the silicon sensor, to the hybrid or to the wire bonds.

- **Pedestal and Noise measurement.** The pedestal is dened as the average value of each APV25 channel output when no signal is applied at its input. This measurement gives the baseline prole of each APV25. The noise is calculated as the root mean square deviation of each APV25 channel with respect to its pedestal value. This test is used for the identication of dead channels or intrinsically noisy APV25 channels.

- **Pulse Shape measurement.** The APV25 has internal calibration circuitry to inject charge in each pre-amplier channel. The full pulse shape can be reconstructed varying the injection time with respect to the main 25 ns chip clock. The pulse shape measurement can be analysed to extract the peak time and the pulse height of the APV25 output. They are then used, together with the noise measurement, to identify opens, shorts and PHs.

- **Silicon bias current versus applied voltage (I-V).** This measurement checks the I-V characteristics of the module and verifies that the current limit of 10 μA is not reached up to the maximum test voltage.

- **LED test.** By means of infra-red light, this test produces large charge signals on the implant strips on the far end with respect to the readout hybrid. The test was intended to find possible breaks in the metal strips.

- **Specic Pinhole Test.** To find PHs, infra-red diodes were used. Those diodes generate a variable current in the silicon crystal as a function of light intensity. In case of a PH, the APV exhibits an abnormal behavior which can be detected by comparison with calibration measurements taken at different photo-generated currents.

2.3.3 Test Equipment

The ARC system was adopted for all the single module tests. It is based on a main ARC board, connected to a PC through a ISA (Industry Standard Architecture) interface card and linked to the module under test via a front-end board. In addition, the ARC system include a dedicated HV supply and an infra-red LED pulser. The former provides the low and bias voltages to the module, the latter is used for LED and pinhole measurements.

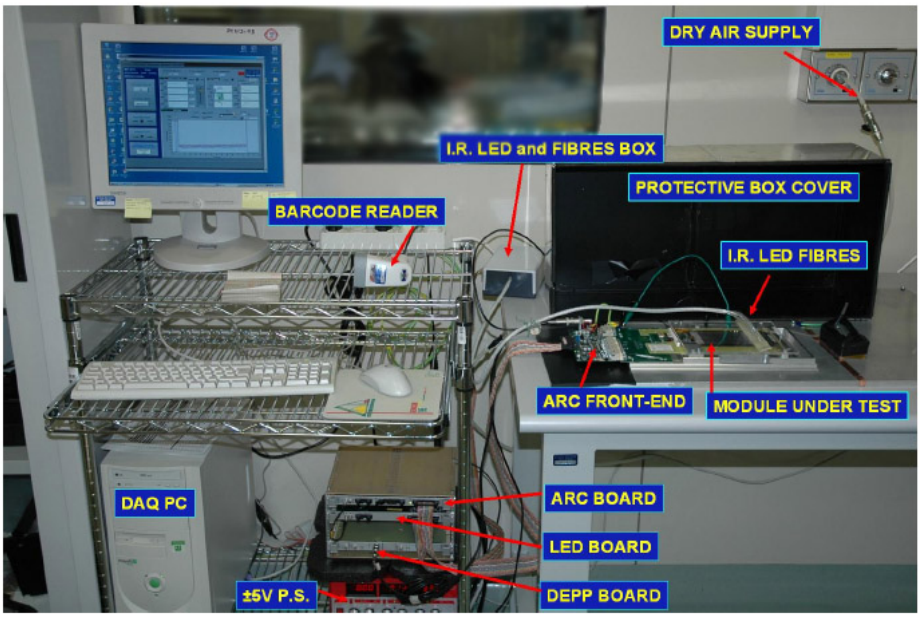

Figure 2.7: ARC setup used for the test of single detector modules, prior to their installation on the substructures of the silicon strip Tracker.

The system was configured to perform automatically two set of measurements, which were indicated as "Fast Test" and the "Full Test" respectively. The Fast Test included nine functionality checks aiming to spot major anomalies which would directly cause the rejection of the module without continuing the following time consuming analysis. The Fast Test verified minimal module functionalities, including components communication, LV levels, pedestal and noise. The Full Test was an extended test of all module functionalities. All the measurements described in section 2.3.2 were performed. The time required to run a Full Test was about 20 minutes.

Figures 2.8 -2.11 show some of the histograms produced by the ARC system during the test of a module. In Fig. 2.8 the noise profile of a six-APV module is shown, which has two channels flagged as noisy and one flagged as open. The cut values for noisy and open channels are superimposed on the profile. The noise profile of a module with a pinhole defective strip is shown in Fig. 2.9: the noise of strip number 388, tagged as a pinhole, is more than 30% lower than the noise of an open strip such as the one visible in Fig. 2.8. In Fig. 2.10 the calibration pulse amplitude is shown for the same defective strip (number 388). As expected, a large variation of the amplitude as a function of the infra-red light intensity is clearly visible. For reference, in Fig. 2.11 the standard calibration amplitude is plotted for the strip next to the pinhole during the same test.

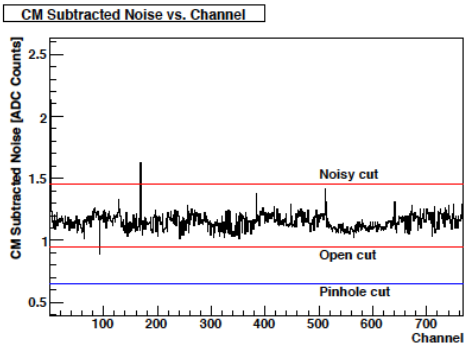

Figure 2.8: Noise profile of a six-APV module (768 strips) which has two channel flagged as noisy and one flagged as open.

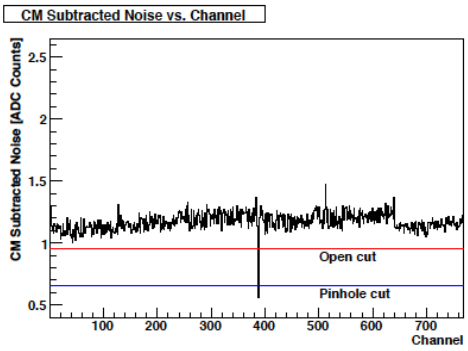

Figure 2.9: Noise profile of a six-APV module (768 strips) which has channel 388 flagged as pinhole.

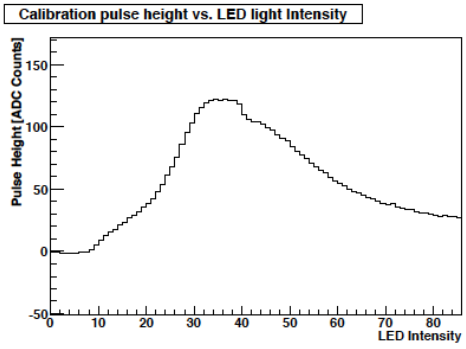

Figure 2.10: Anomalous profile of the calibration pulse amplitude, as a function of the infra-red injected light, for a defective strip containing a pinhole.

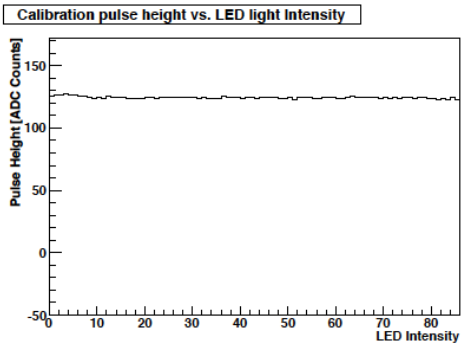

Figure 2.11: Regular profile of the calibration pulse amplitude, as a function of the infra-red injected light, for a strip without defects.

2.3.4 Results

The test procedure carried out in the INFN Laboratories of Pisa and Florence qualified altogether 3 814 modules, of which 3 540 were installed in the TIB and TID sub-detectors. The overall fraction of rejected modules during final production, for all possible reasons, was equal to 3.1%. Less than 1 per mille of the readout strip channels presently installed in the CMS Tracker has been classified as noisy or not working.

Chapter 3

Software for track reconstruction

This chapter describes the software which is used for the reconstruction of charged particle trajectories in the Silicon Tracker of CMS. After an introduction devoted to general principles of *track-finding* and *track-fitting* (Section 3.1), the use of the Kalman Filter formalism for simultaneous identification and fit of particle trajectories [74] is presented (Section 3.2). The implementation of the tracking algorithm used by the CMS collaboration, the Combinatorial Track Finder (CTF), is analyzed in detail (Section 3.3), indicating the personal contributions of the author. Finally, the general performance of the tracking software on different types of simulated events is shown (Sections 3.4 and 3.5).

3.1 General principles of track reconstruction

In HEP experiments, a tracking system generally records the passage of charged particles on several layers of detectors. Given the complete or partial knowledge of the positions of the particles on the sensitive surfaces, the twofold goal of a track reconstruction algorithm is the identification of the trajectories traced by those charged particles which are of interest and the evaluation of their original direction and energy: the former task is commonly denoted as *track-finding*, the latter as *track-fitting*.

Usually the momentum of a charged particle is evaluated measuring its deflection when it passes through a region which is embedded in a known magnetic field. In this case the tracking system works as a *magnetic particle spectrometer*.

3.1.1 Track finding

Given a set of position measurements that are initially provided by the detector, track-finding methods aim to split the original list into smaller, not necessarily disjoint, subsets, each containing only those measurements caused by the same particle. Such a group of correlated positions defines a *track candidate*. A special subset contains all the measurements that cannot be associated to any particle with sufficient certainty. In addition to signals distorted by detector's defects (e.g. dead channels) or by intrinsic interferences (e.g. δ-rays), this subset contains also measurements that are due to noise or are produced by particles that are deliberately excluded from the search, like low momentum ones.

A *track model* is a function in space that describes the path, or *trajectory*, followed by a particle through the detector volume in accord with the *parameters* of the particle. These are a minimal list of values that unambiguously define the state of a particle from a track reconstruction point of view: e.g. its momentum, charge and position on a detecting plane.

The track model determines some constraints among the position measurements of the same track candidate: for example, if a charged particle goes through a volume embedded in a uniform magnetic field and the effect of energy loss is negligible, all the measurements have to be placed along a helicoidal trajectory of constant curvature; similarly, when the energy loss is not negligible, the measurements are arranged along a helix of decreasing curvature. In most of the cases, not-analytical track models are necessary to describe the trajectories of charged particles inside a non-uniform magnetic field. Track-finding methods first use loose criteria to cluster measurements into track candidates in the most efficient way and subsequently they verify the compatibility of the found candidates with the track model. This last check principally aims to reject *ghost* tracks, i.e. spurious combinations of uncorrelated measurements.

Given n points on each of m detecting surfaces, a overly conservative selection of the measurements can provides up to n^m track candidates: this is a huge figure even for a simple event consisting of 20 particles that cross 10 detector layers. Considering that the final rejection test can be a lengthly operation and that often the speed of a track-finding algorithm is as important as its efficiency, some tracking algorithms aim to avoid "combinatorial explosions" using the constraints of the track model already *during* the search of the track candidates. For examples, position measurements can be added to the track candidates in sequence, performing a compatibility check with the track model already at the measurement-level before each point is added to the "growing" (or proto-) track candidate. Balancing the measurement selection between promptness and accuracy, the number of track candidates that are finally compared to the track model can be sensibly reduced and

the overall track-finding process can be considerably speeded up.

Both the selection of track candidates and the one of single measurements have to properly take into account other aspects, in addition to the track model, that are detector-specific: these include the geometric acceptance of the tracking system and the efficiency and resolution of its sensitive components. A common method to reject a spurious track candidate, for instance, is to count the number of detector planes that are crossed by the reconstructed trajectory, but don't contain the corresponding measurements: if that part of the detector is known to be efficient, few missing measurements are enough to indicate that the track candidate is probably a ghost. On the other hand, in the case the trajectory passes through inefficient parts of the detector which are damaged or are not powered, the missing measurements are "expected" and the track candidate shouldn't be rejected. Similarly, the compatibility criteria used to select measurements have to consider not only their positions, but also the associated position uncertainties: e.g. given the estimated crossing point of a trajectory on a detector plane, a measurement that is close and has a very small error can be genuinely less "compatible" to the trajectory than another measurement that is farther away, but has also a larger uncertainty. After the track reconstruction algorithms are developed and tested on a simulated model of the detector, it is of utmost importance to keep monitoring the detector's performance and setup during the real data-taking in order to have a steady control of its actual efficiency and resolution and promptly propagate this information into the reconstruction software.

Methods of track finding

A procedure to classify different track-finding algorithms is to distinguish the *local* methods from the *global* ones.

Local methods. One track candidate is selected at a time. After a set of few measurements is used to initialize the reconstruction procedure, additional points are added in sequence to the growing track candidate, or are rejected, according to their compatibility with the track's parameters. Trajectory parameters are either interpolated or extrapolated to other detector planes where new compatible measurements are looked for. Examples of local methods for track-finding are given in the following:

- **Track-following:** method based on a sequence of alternating selection and extrapolation steps. The parameters of the partially reconstructed track candidate are evaluated on the surface of the last added measurement and then extrapolated to the next sensitive layer of the

tracking system, where an additional position measurement can be chosen and attached to the track candidate. The selection rule can be as simple as one based on a fixed acceptance window or more sophisticated as one based on the actual uncertainty of the extrapolated track parameters. After a measurement is selected, the procedure continues iteratively until the last plane of the tracking system is reached. In order to start, the algorithm needs initial track segments which are used as *seeds*. They are usually, but not always, formed far away from the interaction region, where the trajectories are more separated.

The main advantage of this method is that the extrapolation of track parameters is performed from one detector plane to the close ones only, allowing the use of a simple, and hence fast, track model.

- **Kalman filter**: it is a special refinement of the track-following algorithm. Along with the selection and extrapolation steps, an additional *update* stage is performed after the selection of a news measurement. The track candidate's parameters are combined with the parameters of the selected position measurement in a statistically optimal way. This method naturally combines track-finding and track-fitting and will be described in detail in section 3.2.

- **Track roads:** interpolation is the key of this method. Starting from initial measurements at both ends of a trajectory, the trajectory parameters are interpolated on each intermediate plane of the tracking system. Additional measurements are scrutinized inside a *road* which runs along the intersections of the trajectory with the detector layers. The better the track model which is used, the narrower the road. This method is usually slower than track-following approaches based on extrapolation, but can be the only workable strategy in case of widely spaced detector planes [75].

- **Track segments:** a track candidate is constructed in two steps. Firstly short track segments are built from close measurements, usually inside "natural" subdivision of the tracking system. Secondly, these segments are combined together using the track-following or other track-finding methods. The algorithms have to be adapted to use track elements, instead of single measurements, but the general strategies remain the same.

All the last three track-finding methods described have found application in the context of the CMS reconstruction software. An algorithm based on track roads has been implemented [76] as an alternative strategy for the reconstruction of charged particles using the inner tracking detector, while the track segments approach is used for the reconstruction of trajectories in

the muon system [50]. The default algorithm used by the CMS collaboration for track reconstruction with the silicon tracking system is a *global* version of the Kalman filter track-following strategy, which is described below.

Global methods These methods consider all position measurements simultaneously and identify all track candidates at the same time. In many cases, the list of track candidates is the result of some sort of global minimization that ponders the correlation among all (or most) of the measurements recorded by the tracking system. Algorithm of this kind are usually applicable to a wide range of problems in cluster analysis: the *histogramming* method (add reference) is not used in CMS for track reconstruction, but has found an application for primary vertex finding [77]; a *template matching* method which uses finer and finer approximations was successfully adopted for the trigger of CDF detector [78]; *Hough trasform*, *Minimum Spanning Tree*, *Hopfield networks* and *elastic arms* methods have been used in several different HEP experiments. Because none of these global methods is currently used in CMS for track reconstruction, they are not discussed here. An overview of global methods for track-finding and an accurate list of references can be found for example in [79].

Global Kalman filter In a purely local implementation of the track-following method, after a position measurement has been associated to the "growing" track candidate under scrutiny, the point is removed from the pool of measurements which are available to build additional track candidates. In this way, no account is taken of the fact that a trajectory reconstructed subsequently could match even more closely that particular measurement. While this aspect can be secondary in a setup with few and well separated particles, it cannot be ignored in a crowded high-multiplicity environment like that of LHC collisions.

A symmetric global competition among all track candidates for the entire set of measurements can be implemented in at least two ways:

- All track candidates are built in parallel. The trajectory states of all the proto track-candidates are propagated simultaneously on a specific detecting surface and, after that, the predicted states and the measurements are matched all together by means of a global decision: every prediction is associated with at most one measurement and vice versa. The global decision is usually represented by the minimization of the sum of the state-measurement residuals. This kind of problem is often tackled using artificial neural networks as Hopfield nets [80].
- The trajectories are built one after the other, using a classic track-following method, but the measurements associated to every new can-

didate are *not* removed from the pool of positions that are available to candidates analyzed later. Besides, if a trajectory state is consistent with more than one position measurement on the detector plane under consideration, a new "branch" is open for each compatible point: for any given proto track-candidate which reaches a detector plane, one or several track candidates are propagated afterwards on the next surfaces of the tracking system. Following this strategy, the track-finding procedure returns a final set of candidates which are allowed to share measurements, but are often incompatible with each other: an additional *cleaning* procedure has to be implemented to remove track candidates which are clearly duplicates corresponding to the same original charged particle.

The second approach is the one used in CMS to implement a *global* Kalman Filter (Section 3.3).

3.1.2 Track fitting

Assuming that the clustering of position measurements into track-candidates has already been performed, the aim of track-fitting is to evaluate trajectory parameters and corresponding uncertainties. From the physics point of view, the estimation of parameters is as important as the estimation of their errors: while the momenta and directions of tracks are essential, for example, to determine the invariant mass of a resonance, a good control of track parameters' uncertainties is necessary to estimate how reliable the value of the reconstructed mass is. Besides, from a more operative point of view, the uncertainty on track parameters is also a key element to properly "weight" different tracks in the context of vertex reconstruction.

Before describing concrete techniques for the estimation of trajectory parameters, some general concepts related to track-fitting have to be introduced. In presence of magnetic field, the path of a charged particle through the detector volume is described as a function of the particle parameters, by the *track model*. The latter is usually defined as an application into the *measurement space*, i.e. a n-dimensional space in case of n coordinates measured along the particle trajectory: for example, for m detecting surfaces, each providing 2-dimensional measurements, $n = 2m$ and the track model is:

$$\mathbf{f} : \mathbf{p} \rightarrow f_i(\mathbf{p}), i = 1, \cdots, n \tag{3.1}$$

where $\mathbf{p}$ is the *state vector* of trajectory's parameters and f_i is the i-th coordinate corresponding to the track defined by $\mathbf{p}$. Due to the constraints given by the equations of motion of a charged particle in a magnetic field,

a track is represented by a point on a sub-space of the complete *measurement space*: this sub-space is a *five-dimensional hyperplane* also known as "constraint surface".

Given a state vector $\mathbf{p}$, while $\mathbf{f}$ is a deterministic function, the "measured" coordinates of the trajectory differ from $f_i(\mathbf{p})$ due to stochastic experimental effects. The *measurement vector* $\mathbf{c}$ is defined as

$$\mathbf{c} = \mathbf{f}(\mathbf{p}_{\text{true}}) + \epsilon \tag{3.2}$$

where $\mathbf{p}_{\text{true}}$ is the true unknown value of particle parameters and ϵ is a superimposed term representing contributions from *process noise* (e.g. the random effects of Multiple Scattering and Energy Loss processes) and from *measurement noise* (e.g. the statistical nature of any detector measurement).

After the previous terminology has been introduced, the task of track-fitting can now be reformulated as the determination of an inverse mapping $\mathbf{F}$ from the set of *measurement vectors* $\{\mathbf{c}\}$ onto the set of track parameters $\{\mathbf{p}\}$ without *bias* (eq. 3.4) and with *minimum variance* (eq. 3.5):

$$\mathbf{p}_{\text{estim}} = \mathbf{F}(\mathbf{c}) \tag{3.3}$$

$$< \mathbf{p}_{\text{estim}} > = \mathbf{p}_{\text{true}} \tag{3.4}$$

$$\sigma^2(\mathbf{p}_{\text{estim},i}) \equiv < (\mathbf{p}_{\text{estim},i} - \mathbf{p}_{\text{true},i})^2 > \rightarrow \text{minimum} \tag{3.5}$$

where $<>$ denotes the expectation value of a random variable. In order to define an estimator $\mathbf{F}$ without bias and minimum variance, it is necessary to take into account the effects of both process and measurement noise. The former requires a precise knowledge of the amount of material traversed by the analyzed particles, the latter demands a carefully estimate of detector resolution. Before data-taking, both aspects are evaluated from theoretical considerations and from calibration tests; subsequently, detector resolution and alignment can be monitored directly during the running of the experiment using reconstructed tracks that have sufficient constraints: e.g trajectories corresponding to particles produced from the decay of known resonances or trajectories of particles, like muons, that are independently reconstructed also in other sub-systems of the experimental apparatus. Analogously, the actual value of the material budget can be inferred from the rate of conversion of photons or inelastic interaction of hadrons.

Least Squares method

The *global* track-fitting by Least Squares method (LS) is described in this section. It holds an important role in the field of statistical data analysis

since, under general conditions (see [81]), it is the estimator with minimum variance among the class of linear and unbiased estimators. It is discussed here in order to understand the Kalman filter estimator, which is a recursive application of the LS method.

Usually the track model, as it has been introduced in eq. 3.1, is a *nonlinear* function of particle parameters. While the LS method can be applied directly to linear track models, it works also in more general cases using the first term of the track model's Taylor expansion:

$$\mathbf{f}(\mathbf{p}) = \mathbf{f}(\mathbf{p_0}) + \mathbf{A} \cdot (\mathbf{p} - \mathbf{p_0}) + O((\mathbf{p} - \mathbf{p_0})^2) \tag{3.6}$$

with:

$$\mathbf{A} = \partial \mathbf{f}(\mathbf{p}) / \partial \mathbf{p} \text{ at } \mathbf{p} = \mathbf{p_0}$$

The range of $\mathbf{p}$ values for which the approximate linear model is sufficiently close to the exact track model depends strongly on the choice of the parameters that are used to define the state vector and on the layout of the tracking system [82].

For the aim of track-fitting, the statistical nature of the measurements provided by the detector is contained into the *covariance matrix* $\mathbf{V}$

$$\mathbf{V} = < (\mathbf{c} - < \mathbf{c} >)(\mathbf{c} - < \mathbf{c} >)^T > \tag{3.7}$$

with $\mathbf{c} = \mathbf{f}(\mathbf{p}_{\text{true}}) + \epsilon$ from equation 3.2. In order to estimate the particle's state vector $\mathbf{p}$, the LS method minimizes the function:

$$M(\mathbf{p}) = [\mathbf{f}(\mathbf{p_0}) + \mathbf{A} \cdot (\mathbf{p} - \mathbf{p_0}) - \mathbf{m}]^T \cdot \mathbf{W} \cdot [\mathbf{f}(\mathbf{p_0}) + \mathbf{A} \cdot (\mathbf{p} - \mathbf{p_0}) - \mathbf{m}] \tag{3.8}$$

where $\mathbf{W} \equiv \mathbf{V}^{-1}$ is called *weight matrix* and $\mathbf{m}$ is a "realization" of the random variable $\mathbf{c}$, i.e. it is the specific set of measurements defining a track-candidate. The goal of minimizing M is to determine the *state vector* $\mathbf{p}$ such that the probability to get the *measurement vector* $\mathbf{m}$ is the highest. This can be achieved differentiating M with respect to $\mathbf{p}$ and requiring $\partial M / \partial \mathbf{p} = 0$:

$$\mathbf{p}_{\text{estim}} = \mathbf{p_0} + (\mathbf{A^T W A})^{-1} \cdot \mathbf{A}^T \cdot \mathbf{W} \cdot (\mathbf{m} - \mathbf{f}(\mathbf{p_0})) \tag{3.9}$$

Several proprieties of the LS estimator can be drawn (for details, see [83]):

- If the measurement vector $\mathbf{c} = \mathbf{f}(\mathbf{p}_{\text{true}}) + \epsilon$ is *unbiased*, i.e. $< \epsilon >= 0$, then also the estimated parameter vector $\mathbf{p}_{\text{estim}}$ is determined without *bias*.

- The error matrix of the estimated parameters, i.e. the covariance matrix $\mathbf{C}$ of $\mathbf{p}_{\text{estim}}$, is a simple expression of the weight matrix $\mathbf{W}$ and of the linear model $\mathbf{A}$:

$$\mathbf{C}(\mathbf{p}_{\text{estim}}) = (\mathbf{A}^T\mathbf{W}\mathbf{A})^{-1} \tag{3.10}$$

- Given a linear track model (or a linearized one) and unbiased measurements, the LS method is the *optimal* linear estimator, that is it is the estimator with least variance (Gauss-Markov theorem) among linear methods.

- If the track model is linear and the measurements, in addition to be unbiased, have also gaussian distributed errors, the LS method is not only optimal, but also *efficient*: this means that LS method is an estimator, among linear and non-linear ones, with variance equal to the minimum bound (Cramer-Rao inequality).

- Defined the "readjusted" measurements as $\tilde{\mathbf{c}} := \mathbf{f}(\mathbf{p}_{\text{estim}})$, the corresponding covariance matrix is given by:

$$\mathbf{C}(\tilde{\mathbf{c}}) = \mathbf{A}(\mathbf{A}^T\mathbf{W}\mathbf{A})^{-1}\mathbf{A}^T \tag{3.11}$$

  The *residual* vector $\mathbf{r} := \mathbf{c} - \tilde{\mathbf{c}}$ has a covariance matrix $\mathbf{C}(\mathbf{r})$ equal to:

$$\mathbf{C}(\mathbf{r}) = \mathbf{V}(\mathbf{c}) - \mathbf{C}(\tilde{\mathbf{c}}) \tag{3.12}$$

After the error matrix $\mathbf{V}(\mathbf{c})$ of the measurements has been evaluated, the important advantage of the LS method is that it explicitly provides both the fitted parameters (see eq. 3.9) and the corresponding covariance matrix (see eq. 3.10). If the measurement errors are *uncorrelated*, the covariance matrix takes a very simple form

$$(\mathbf{V})_{ij} = \delta_{ij} \cdot \sigma_j^2 \tag{3.13}$$

where σ_j^2 is the standard deviation of the j-th measurement. Being $\mathbf{V}(\mathbf{c})$ diagonal, the weight matrix $\mathbf{W}$ is trivially calculated .

Nevertheless, despite of all the positive characteristics of the LS method that have been listed, if the correlation between measurements is not negligible and $\mathbf{V}$ is not diagonal [84], the time which is necessary to invert the matrix and to calculate $\mathbf{W}$ and $\mathbf{C}(\mathbf{p}_{\text{estim}})$ scales as n^3, where n is the number of measured coordinates per trajectory. In experiments like CMS, for which hundreds, or even thousands, of trajectories are reconstructed in every event using $O(10)$ sensitive layers, the application of the global LS method can be particularly slow and it is preferable to use a recursive approach as the one described in the following section.

3.2 Kalman Filtering for track finding and fitting

The previously described LS method for track-fitting has two severe limitations in case the effect of Multiple Scattering is not negligible. The first one is due to the fact that the covariance matrix of the measurements is not diagonal and it is cpu-expensive to evaluate and invert it. The second comes from the fact that the LS estimator performs a sort of single *global* fit of the track-candidate's measurements: all the information provided by the detector is used to estimate the track parameters $\mathbf{p}_{\mathbf{r},\mathrm{estim}}$ on a unique specific *reference point*, $\mathbf{r}$. The calculated crossing points of the trajectory on the other detector planes, given by $\mathbf{f}(\mathbf{p}_{\mathbf{r},\mathrm{estim}})$, correspond to an unscattered propagation of the track parameters $\mathbf{p}_{\mathbf{r},\mathrm{estim}}$. While a track model can describe the corrections to the particle trajectory due to the average energy loss, stochastic fluctuations from the average scattering angle and from the average energy reduction are event-specific effects which cannot be included in a general deterministic function.

In the contest of the DELPHI collaboration [85], a novel method for track-fitting, called *progressive fit* [86], [87], was proposed: according to this approach, the trajectory is "followed" from one detector plane to the next one and the track parameters are estimated repeatedly using, as reference point, the position of each measurement encountered along the trajectory. There is not a unique *global* fit of the track parameters at one specific reference point (as it is done by the LS method), but there are as many parameters' estimations as many measurements. After its first proposal, it was shown [74] that the progressive fit is a special case of the Kalman filter, a well known algorithm used for analysis of linear dynamic systems [88]. The method is described here using the same general formalism.

The Kalman filter is a method for estimating the states of dynamic systems, that is evolving stochastic models of some time-varying phenomenon. The notion of time can be extended to a general parameter $\mathbf{t}$ on which the *system state* $\mathbf{p}(t)$ is functionally dependent. Then the application of Kalman filter to track-fitting is straightforward if the track is considered as a *discrete* dynamic system: the system state for a track (the 5-dimensional state vector $\mathbf{p}$ of track parameters) is considered as a function $\mathbf{p}_k$ of the *discrete* variable k, that is the index corresponding to the k-th detecting surface of the tracking system.

The evolution of track parameters between two consecutive detectors is described by a *deterministic* function $\mathbf{f}_k$ plus a *stochastic* term δ_k and is summarized by the following *system equation*:

$$\mathbf{p}_k = \mathbf{f}_k(\mathbf{p}_{k-1}) + \mathbf{P}_k\delta_k, \qquad <\delta_k>= 0, \qquad \mathbf{Q}_k := \mathbf{C}(\delta_k) \tag{3.14}$$

The random variable δ_k describes the *process noise*, which is, in the context

of track-fitting, the stochastic effects of matter on the particle trajectory (e.g. Multiple Scattering and Energy Loss). There is one covariance matrix $\mathbf{Q}_k$ for each index k because, in general, the amount of material is not the same between all pairs of detector's surfaces; the $\mathbf{P}_k$ matrix projects the effect of *process noise* only on the subset of track parameters involved (e.g. MS principally affects the trajectory direction, whereas the position offset is negligible). Similarly to the track model introduced earlier (eq. 3.1), the function $\mathbf{f}_k$ propagates the state of the trajectory between different points inside the volume of the tracking system, but, instead of returning only the trajectory intersection with the detector surfaces, $\mathbf{f}_k$ returns the full 5-dimensional state vector. On the other hand, with respect to the track model, the function $\mathbf{f}_k$ has a smaller scope since it has to be valid only between the two consecutive detector's surfaces k-1 and k.

The *measurement equation* reflects the relation among state vector, experimental observations and measurement uncertainties on the surface k:

$$\mathbf{m}_k = \mathbf{h}_k(\mathbf{p}_k) + \epsilon_k, \qquad < \epsilon_k >= 0, \qquad \mathbf{V}_k = \mathbf{W}_k^{-1} = \mathbf{C}(\epsilon_k) \tag{3.15}$$

The *measurement noise* ϵ_k is assumed to be unbiased and uncorrelated to the *process noise* δ_k and its covariance matrix $\mathbf{V}_k$ describes the detector resolution. The function $\mathbf{h}_k$ maps the track parameters on the same space of the measurements provided by the detector: depending on the track parameterization that is used, $\mathbf{h}_k$ is often a simple projection function (see [89]). While the Kalman filter assumes a linear system, in many cases $\mathbf{f}_k$, and in general also $\mathbf{h}_k$, are nonlinear functions. Hence, both are usually approximated by their Taylor first-order expansions for all k:

$$\mathbf{f}_k(\mathbf{p}_{k-1}) = \mathbf{F}_k \cdot \mathbf{p}_{k-1} + \mathbf{c}_k, \qquad \mathbf{h}_k(\mathbf{p}_k) = \mathbf{H}_k \cdot \mathbf{p}_k + \mathbf{d}_k, \qquad k = 1, \ldots, n \tag{3.16}$$

In the following the constant terms $\mathbf{c}_k$ and $\mathbf{d}_k$ are suppressed for the sake of convenience.

Sorting the n position measurements $\mathbf{m}_1, \ldots, \mathbf{m}_n$ of a track candidate according to the detector's surfaces $1, \ldots, n$, let $\mathbf{p}_{k|j}$ indicate the vector state on surface k-th which is estimated using only the information from measurements $\mathbf{m}_1, \ldots, \mathbf{m}_j$. Given $\mathbf{p}_{k-1|k-1}$ and $\mathbf{m}_k$, the Kalman filter evaluates $\mathbf{p}_{k|k}$ through two steps:

- **prediction:** the vector state $\mathbf{p}_{k|k-1}$ and the corresponding covariance matrix $\mathbf{C}_{k|k-1}$ are evaluated on the base of the first $k-1$ measurements

$$\mathbf{p}_{k|k-1} = \mathbf{F}_k \cdot \mathbf{p}_{k-1|k-1} \tag{3.17}$$

$$\mathbf{C}_{k|k-1} = \mathbf{F}_k \mathbf{C}_{k-1|k-1} \mathbf{F}_k^T + \mathbf{P}_k \mathbf{Q}_k \mathbf{P}_k^T \tag{3.18}$$

The predicted covariance matrix is the sum of two terms: the first represents the simple error propagation applied to $\mathbf{C}_{k-1|k-1}$, while the second is due to the process noise.

- **filtering:** the estimate of the vector state on plane k is calculated combining the predicted state $\mathbf{p}_{k|k-1}$ and the measurement $\mathbf{m}_k$ in a statistically optimal way. Applying the Least Squares method described in 3.1.2, the state vector $\mathbf{p}$ which minimizes the following function

$$\begin{aligned} M(\mathbf{p}) &= (\mathbf{m}_k - \mathbf{H}_k\mathbf{p})^T\mathbf{W}_k(\mathbf{m}_k - \mathbf{H}_k\mathbf{p}) \\ &\quad +(\mathbf{p}_{k|k-1} - \mathbf{p})^T\mathbf{C}_{k|k-1}^{-1}(\mathbf{p}_{k|k-1} - \mathbf{p}) \end{aligned} \tag{3.19}$$

represents the *filtered* state $\mathbf{p}_{k|k}$. After imposing $\partial M/\partial\mathbf{p} = 0$ and some matrix algebra:

$$\mathbf{p}_{k|k} = \mathbf{p}_{k|k-1} + \mathbf{K}_k(\mathbf{m}_k - \mathbf{H}_k\mathbf{p}_{k|k-1}) \tag{3.20}$$

$$\mathbf{K}_k = \mathbf{C}_{k|k-1}\mathbf{H}_k^T(\mathbf{V}_k + \mathbf{H}_k\mathbf{C}_{k|k-1}\mathbf{H}_k^T)^{-1} \tag{3.21}$$

$$\mathbf{C_{k|k}} = (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{C}_{k|k-1} \tag{3.22}$$

The state vector $\mathbf{p}_{k|k}$ results as the predicted $\mathbf{p}_{k|k-1}$ plus a correction term involving the measurement $\mathbf{m}_k$. The correction is rescaled by $\mathbf{K}_k$, which is known as *gain matrix* in the context of Kalman filtering formalism.

For each measurement $\mathbf{m}_k$, there are simple expressions for the residuals of predictions (*predicted residuals*) $\mathbf{r}_{k|k-1}$, the corresponding covariance matrices $\mathbf{R}_{k|k-1}$ and the χ^2:

$$\mathbf{r}_{k|k-1} = \mathbf{m}_k - \mathbf{H}_k\mathbf{p}_{k|k-1} \tag{3.23}$$

$$\mathbf{R}_{k|k-1} = \mathbf{C}(\mathbf{r}_{k|k-1}) = \mathbf{V}_k + \mathbf{H}_k\mathbf{C}_{k|k-1}\mathbf{H}_k^T \tag{3.24}$$

$$\chi^2_{k,P} = \mathbf{r}_{k|k-1}^T \cdot \mathbf{R}_{k|k-1}^{-1} \cdot \mathbf{r}_{k|k-1} \tag{3.25}$$

For process and measurement noises that are both gaussian, $\chi^2_{k,P}$ is χ^2-distributed with $\dim(\mathbf{m}_k)$ degrees of freedom. Under the same assumptions, the residuals are independent and they can be summed up to define the total χ^2 of the trajectory:

$$\chi^2_k = \chi^2_{k-1} + \chi^2_{k,P} \tag{3.26}$$

Residuals can be defined also for the filtered states (*filtered residuals*). The corresponding formulas are:

$$\mathbf{r}_{k|k} = \mathbf{m}_k - \mathbf{H}_k\mathbf{p}_{k|k} \tag{3.27}$$

$$\mathbf{R}_{k|k} = \mathbf{C}(\mathbf{r}_{k|k}) = \mathbf{V}_k - \mathbf{H}_k\mathbf{C}_{k|k}\mathbf{H}_k^T \tag{3.28}$$

$$\chi^2_{k,F} = \mathbf{r}_{k|k}^T \cdot \mathbf{R}_{k|k}^{-1} \cdot \mathbf{r}_{k|k} \tag{3.29}$$

The covariance matrix $\mathbf{R}_{k|k}$ has smaller values with respect to the matrix for predicted residuals (the filtered state vector is "pulled" toward the measurement), nevertheless it can be shown that the χ^2 values of predicted and filtered residuals are exactly the same:

$$\chi^2_{k,P} = \chi^2_{k,F} = \mathbf{r}^T_{k|k} \cdot \mathbf{R}^{-1}_{k|k} \cdot \mathbf{r}_{k|k} \tag{3.30}$$

This formalism encourages the combination of track-fitting with a track-following procedure for track-finding [90, 91]: the predicted state $\mathbf{p}_{k|k-1}$ indicates where to look for new measurements on the k-th detector's plane; the corresponding covariance matrix $\mathbf{C}_{k|k-1}$ can be used to define a sensible search window and $\chi^2_{k,P}$ is an appropriate statistic test to reject or to accept a measurement. In case the measurement $\mathbf{m}_k$ is accepted, the filtered state vector $\mathbf{p}_{k|k}$ is the starting point to iterate the prediction/filtering steps on the next surface of the tracking system.

The Kalman filter technique that has been described has the limitation that only the last state vector $\mathbf{p}_{n|n}$ exploits the full information from all n measurements $\mathbf{m}_1, \ldots, \mathbf{m}_n$. In order to propagate the knowledge provided by all the n measurements also to the other state vectors $\mathbf{p}_k$ (with $1 \leq k < n$), an additional step has to be added to the iterative track-fitting procedure:

- **smoothing:** the *filtered* vector state $\mathbf{p}_{k|k}$ is combined with an extrapolation of the *smoothed* state $\mathbf{p}_{k+1|n}$, corresponding to layer $(k+1)$-th, to get a smoothed state $\mathbf{p}_{k|n}$ also on the layer k-th. On the last layer n, smoothed and filtered states are equivalent and this allows the iterative smoothing procedure to start and to obtain the smoothed states on all the remaining $1, \ldots, n$-1 surfaces. The relevant formulae are:

$$\begin{aligned}
\mathbf{p}_{k|n} &= \mathbf{p}_{k|k} - \mathbf{A}_k \cdot (\mathbf{p}_{k+1|k} - \mathbf{p}_{k+1|n}) && (3.31)\\
\mathbf{C}_{k|n} &= \mathbf{C}_{k|k} - \mathbf{A}_k \cdot (\mathbf{C}_{k+1|k} - \mathbf{C}_{k+1|n}) \cdot \mathbf{A}^T_k && (3.32)\\
\mathbf{A}_k &= \mathbf{C}_{k|k} \cdot \mathbf{F}^T_{k+1} \cdot \mathbf{C}^{-1}_{k+1|k} && (3.33)
\end{aligned}$$

  where the matrix $\mathbf{A}_k$ is commonly indicated as *smoother gain matrix*. The state $\mathbf{p}_{k+1|n}$ contains the full information from all the n measurements of the trajectory and therefore its covariance matrix $\mathbf{C}_{k|n}$ has smaller terms respect to the matrix $\mathbf{C}_{k|k}$ of the state $\mathbf{p}_{k|k}$, which is estimated using only the first k measurements.

Similarly to the filtered results, also the smoothed ones are accompanied by residuals and χ^2:

$$\begin{aligned}
\mathbf{r}_{k|n} &= \mathbf{m}_k - \mathbf{H}_k \mathbf{p}_{k|n} && (3.34)\\
\mathbf{R}_{k|n} &= \mathbf{C}(\mathbf{r}_{k|n}) = \mathbf{V}_k - \mathbf{H}_k \mathbf{C}_{k|n} \mathbf{H}^T_k && (3.35)\\
\chi^2_{k,S} &= \mathbf{r}^T_{k|n} \cdot \mathbf{R}^{-1}_{k|n} \cdot \mathbf{r}_{k|n} && (3.36)
\end{aligned}$$

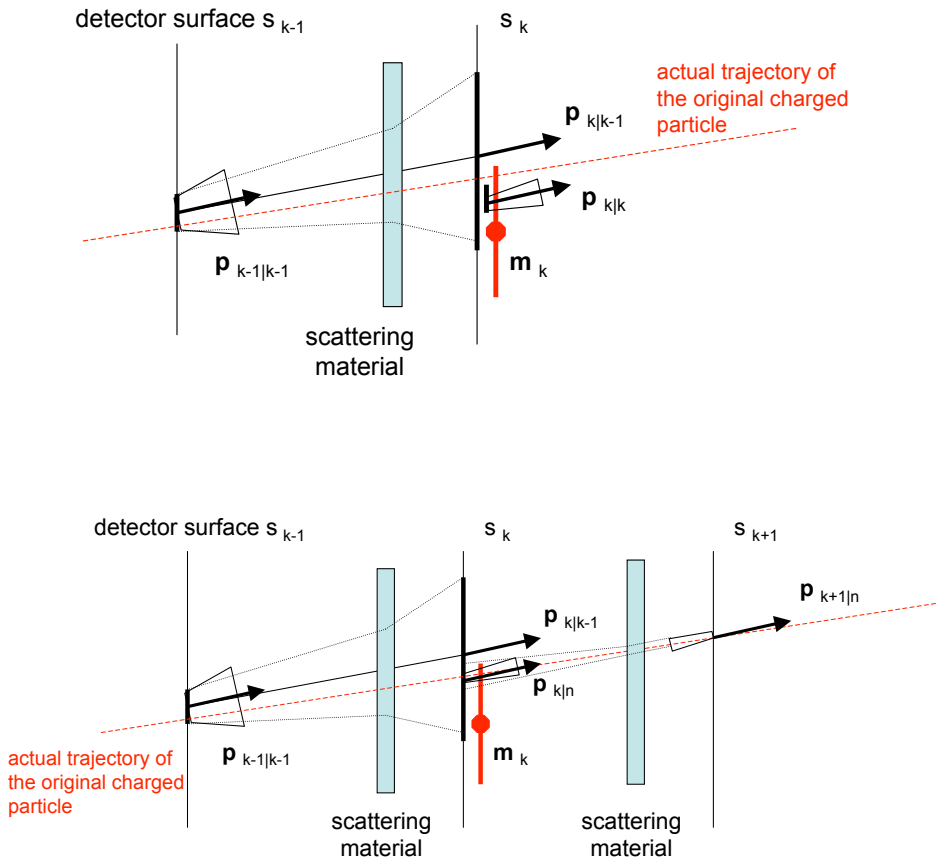

Figure 3.1: *Pictorial representation of the Kalman prediction and filtering steps (top) and of the smoothing step (bottom).*

Being $\mathbf{C}_{k|n}$ smaller than $\mathbf{C}_{k|k}$, equations 3.28 and 3.35 show that the mean squared smoothed-residual is larger than the mean squared filtered-residual. While the filtered state is the result of a weighted sum between the position of the measurement and *one* predicted state, the smoothed state is obtained combining the measurement with *two* predicted states: the extrapolation on the surface k of state $\mathbf{p}_{k-1|k-1}$ and that of state $\mathbf{p}_{k+1|n}$ (Fig. 3.1). Compared to the filtered one, the smoothed result is "pulled" less toward the specific measurement on a sensitive plane (and has consequently larger residuals), but is expected to be closer to the actual state of the original particle.

Under the assumption of gaussian and unbiased errors, the smoothed $\chi^2_{k,S}$ $(k = 1, \ldots, n)$ are still χ^2 distributed. Nevertheless, unlike the filtered $\chi^2_{k,F}$,

they contain the same information from the same set of measurements and are not independent: they cannot be summed up to define a global χ^2_S of the trajectory.

To recapitulate, the main advantages of the Kalman filter/smoother with respect to a global fit are that:

- no large matrices have to be inverted.
- the approximation of the track model doesn't need to be valid over the entire length of the track, but only between adjacent detector's surfaces.
- the estimated track parameters follow closely the real path of the particle, considering event-specific fluctuations in energy loss and scattering.
- track-finding and track-fitting can be performed simultaneously.

On top of that, the Kalman filter/smoother method is based *locally* on the LS principle and therefore it is an *efficient* estimator under the same two general conditions: the track model is linear (or can be approximated by a linear function) and the measurement and process noises are unbiased and gaussian distributed. While the first condition can usually be met by a suitable choice of the track parameters, sometimes the second is violated: measurement noise can contain non gaussian tails due to outlying effects (e.g. δ-rays effects on silicon pixel detector); process noise can be non gaussian if the Multiple Scattering is due to an inhomogeneous scatterer or the energy loss is dominated by bremsstrahlung. In these conditions, an estimator based on LS method is still optimal among linear estimators, but it is not efficient. A more robust nonlinear estimator is the Gaussian-sum filter (GSF) [92], [93]: extending the Kalman filter/smoother method, it can be used when noise effects are not gaussian, but they are described at least by a finite sum of gaussian distributions. In CMS, due to the significant effects of bremsstrahlung within the tracker volume, the GSF approach is used for the track-fitting of electron trajectories.

3.3 Combinatorial Track Finder algorithm

The Combinatorial Track Finder (CTF) is the algorithm used to reconstruct the trajectories of charged particles inside the silicon tracking system of CMS. The bulk of the algorithm and most of the ancillary services were originally developed in the old software framework of the experiment, COBRA [94], using the reconstruction suite ORCA [95]. The cooperation of several physicists spanning about eight years produced the final version of the track reconstruction program that was released in 2005. During the same year, a re-engeneering of the CMS software was decided by the management of the experiment and a new software framework, CMSSW [96], was implemented to replace the previous one. Most of the algorithms for offline reconstruction that were originally implemented in ORCA have been "ported" to the new framework. Also the CTF algorithm, which proved to be successful in ORCA-based physics analyses [50], has been ported to the new framework through an effort which took approximately a year.

The author contributed to improving the CTF algorithm during the last year of development in the COBRA framework and co-directed the porting of the tracking software to the new framework in 2006. Since then the author has steered the group of people aiming to further improve the performance of the track reconstruction software of CMS. In these last two sections of the chapter, the implementation of the CTF algorithm, as of December 2008, is described and several performance plots, produced using the same version of the software, are shown. The next chapter will analyze in detail some of the improvements to the tracking software that have been developed by the author and that are now part of the standard reconstruction program.

3.3.1 Tracking software modules

The track reconstruction sequence is modularized into 5 logical parts:

- **Local reconstruction** consists in clustering into *hits* the strip and pixel signals produced by charged particles on the silicon detectors of the tracking system. The positions of the hits are estimated along with the corresponding uncertainties.
- **Seed generation** provides initial track-candidates for the full track reconstruction. A seed defines initial trajectory parameters and errors.
- **Pattern recognition** module is based on a global Kalman filter and is responsible of finding the track-candidates that correspond to charged particles of interest. Because trajectories are build in parallel and they are allowed to share position measurements, this module is also

responsible of cleaning the track-candidates collection by removing duplicates.

- **Final Track fit** module estimates the parameters of trajectories with ultimate precision by means of Kalman filter and smoother.

While the first reconstruction step provides the necessary input to any global reconstruction algorithm which aims to reconstruct the trajectories of charged particles, it is not a specific component of the CTF reconstruction sequence. The algorithms used to identify pixel and strip hits and to measure their positions are described in literature [97, 56] and [50] and are not illustrated again here. All the other three reconstruction modules have been developed in the context of the CTF algorithm and are described in the following.

3.3.2 Seed generation

The trajectory seeds define the starting trajectory parameters and errors that are used to initiate the iterative track-following procedure employed to identify the track-candidates. The seeds can be obtained externally to the Tracker, using inputs from other subsystems such as the muon chambers or the calorimeters, yet the precision of initial trajectory parameters obtained in this way is generally poor. Besides, exclusively those particles that extend to the other subsystems (mostly muons, electrons and converted photons) can be seeded externally to the Tracker. Another strategy is to construct seeds internally. In this case each seed is composed from a small subset of the position measurements in the Tracker that are supposed to come from the interaction of one single charged particle with the detector. Since 5 parameters (including the trajectory curvature) are needed to start trajectory building, at least 3 hits, or 2 hits and a beam constraint, are necessary to properly define a seed.

Because the seed roughly estimates the initial parameters of the trajectory, the early iterations of the track-finding module employ necessarily large search windows to collect fully the rest of the position measurements that correspond to the particle. While it seems fairly natural to facilitate the pattern recognition constructing the trajectory seeds on the external layers of the tracking system where the detector occupancy is lower, there are three important reasons why the opposite approach is preferred. Firstly, even if the average occupancy of the *strip* layers of the Tracker system decreases faster than $1/r^2$ for detector placed at larger and larger radii, the occupancy of the inner *pixel* subsystem, due to its high density of read-out channels per surface unit, is about one order of magnitude lower than that of the most external strip detectors (Fig. 3.2). Secondly the pixel sensors provide

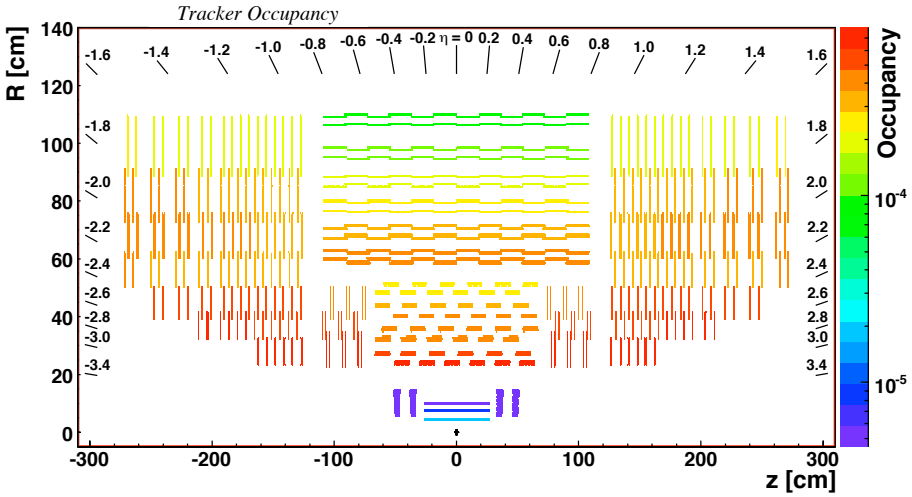

Figure 3.2: *Occupancy of the Silicon Tracker detectors for Minimum Bias events simulated with superimposed pile-up collisions.*

truly 2-dimensional measurements that, given the same lever arm, allow the construction of trajectory seeds that have better defined parameters than equivalent seeds obtained with strip measurements. Finally, the most important reason to construct seeds in the innermost layers of the tracking system is due to the non negligible material budget of the Tracker. Many particles produced in LHC collisions suffer destructive interactions before they reach the outermost layers of the tracking system: while most of the muons cross the whole tracker volume without being absorbed (Fig. 3.3), between 5% and 15% of the pions that are produced in collisions interact inelastically even before they reach the fifth layer of the Tracker (Fig. 3.4); similarly many electrons lose most of their energy due to bremsstrahlung radiation after they cross only a few layers. It is therefore apparent that, given the characteristics of the CMS detector, a track finding algorithm can be highly efficient only by initiating the reconstruction of charged particle trajectories on the innermost layers and moving inside-out.

More than 90% of charged particles produced in LHC collisions, and which are inside the geometrical acceptance of the Tracker, cross pixel sensors on 3 different layers and can be reconstructed starting from trajectory seeds that are obtained from triplets of pixel measurements. Nevertheless, since the barrel layers of the pixel detector are not fully hermetic (Sec. A.0.1), seeds are also produced from pairs of pixel measurements plus a vertex or beam-spot constraint. Finally, additional seeds are constructed using a combination of pixel and strip measurements or uniquely strip measurements in order to extend the track reconstruction efficacy beyond the geometrical

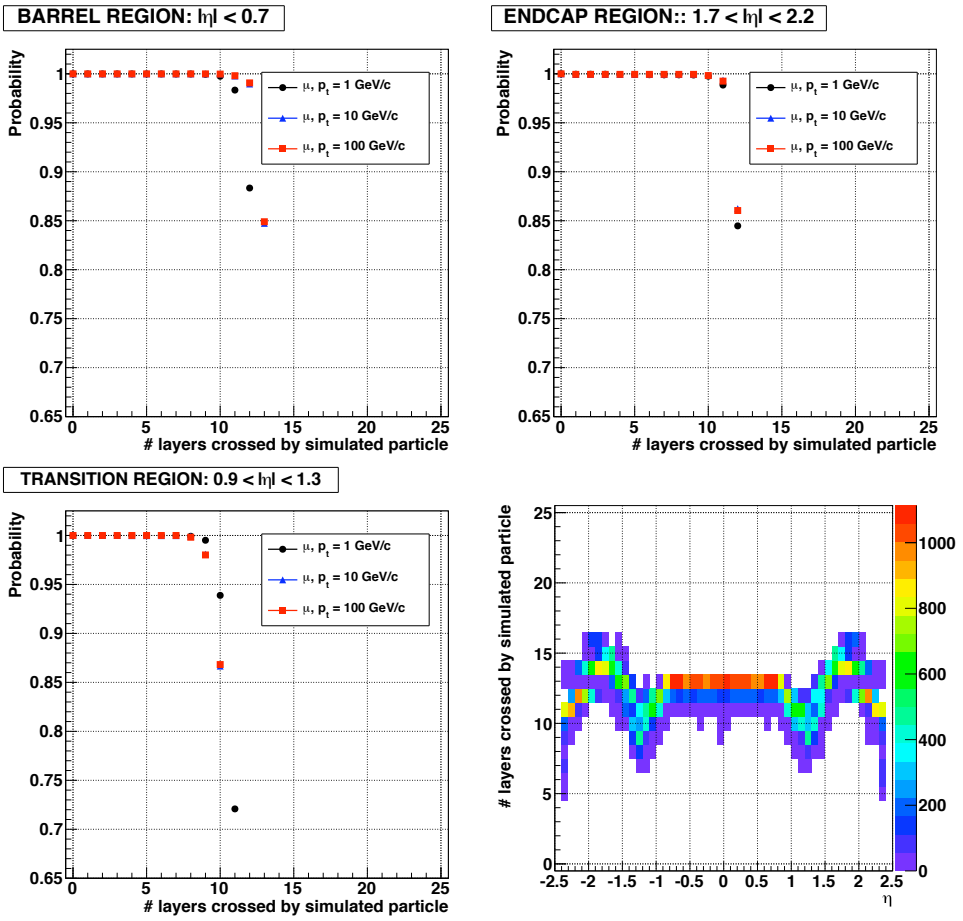

Figure 3.3: *The first three plots show the probability that muons of different momenta are not absorbed by the detector material as a function of the number of Tracker layers that they cross moving outward from the collision point. Probabilities are evaluated for particles moving through the barrel layers (top-left), the endcaps ones (top-right) and the barrel-endcap transition region (bottom-left). The bottom-right plots show the number of crossed layers versus pseudorapidity for muons with transverse momentum equal to 10 GeV/c.*

acceptance of the pixel system (Sec. 4.1). In the implementation of the CTF algorithm used for this analysis, four collections of seeds are used:

- **Pixel Triplets:** seeds produced from triplets of pixel measurements allow the reconstruction of most of the primary charged particles produced in LHC collisions and facilitate a fast pattern recognition on the remaining layers of the silicon Tracker due to a pretty precise estimate of the trajectory seed parameters. A loose beam-spot constraint is em-

ployed to remove from the seeds collection those combination of pixel triplets that are not compatible with particles produced in the luminous region and that correspond presumably to ghosts. However, the beam-spot constraint is not used in the estimation of the trajectory seed parameters and errors.

- **Pixel and Strip pairs with vertex constraint:** using the previous collection of pixel triplets, the vertices corresponding to the primary collision and to the pile-up events are reconstructed. Pairs of tracker measurements are combined with the positions of these vertices to define new trajectory seeds and assure a ∼99% seeding efficiency also in the barrel region of the pixel tracker, despite its non complete hermeticity. While the vertex constraint is initially used to define the seed parameters, its position is not used during the final fit of the track in order to remove any bias. In addition to all the pixel measurements, the strip measurements that are used for the construction of this collection of seeds are hits on the two innermost rings of the three innermost TEC layers.

- **Pixel and Strip pairs with beam-spot constraint:** the set of measurements that are used to produce this collection of seeds is the same employed previously, but this time, instead of using a vertex constraint, a looser beam-spot compatibility is used in the definition of seed parameters. This seeds collection is designed mainly to recover trajectories corresponding to particles produced from the decay of long-living hadrons, like b-mesons: while these particles are produced within 1-2 mm from the beam line, they are not compatible with the primary vertex position. The bias from the beam-spot constraint is removed during the final fit of the tracks.

- **Strip-only pairs with beam-spot constraint:** these seeds are constructed using only strip measurements on the two innermost TIB layers and on the disks of TID. The goal of this collection is twofold: assuring the reconstruction of trajectories that originate so far from the interaction point (e.g. K_s decay products) that they cross only one, or even zero, pixel layers; allowing the reconstruction of tracks at $\eta \sim 0$ where the pixel detector is intrinsically inefficient. The beam-spot constraint is very loose and its bias is completely removed during the final fit of the tracks.

Recently, the software to produce a fifth collection of seeds from measurements on the two double-sided layers of TOB and a combination of TEC rings has been implemented in order to extend the track reconstruction efficacy at higher radii. These seeds allow the reconstruction of secondary particles from the inelastic nuclear interactions of primary hadrons with the

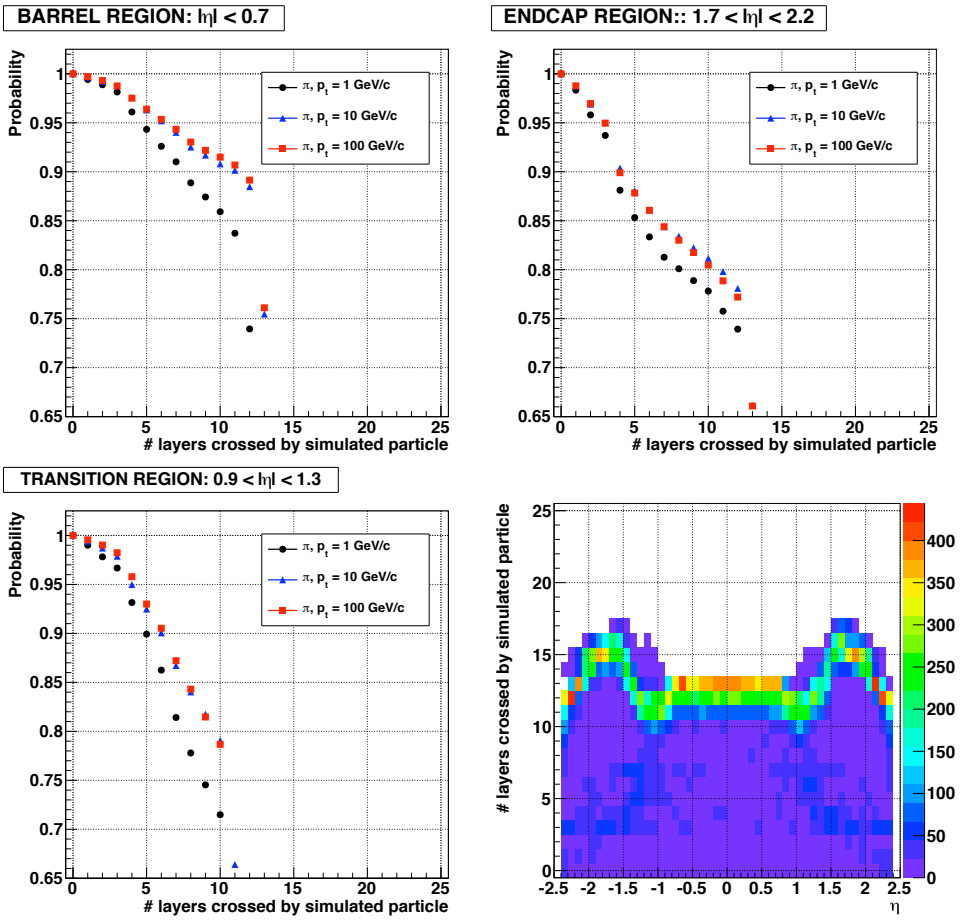

Figure 3.4: *The first three plots show the probability that pions of different momenta are not absorbed by the detector material as a function of the number of Tracker layers that they cross moving outward from the collision point. Probabilities are evaluated for particles moving through the barrel layers (top-left), the endcaps ones (top-right) and the barrel-endcap transition region (bottom-left). The bottom-right plots show the number of crossed layers versus pseudorapidity for pions with transverse momentum equal to 10 GeV/c.*

Tracker material; similarly this collection of seeds can be used for the identification of photons converting up to 50 cm far from the interaction point and the reconstruction of particles that, according to some physics models beyond the standard model, are expected to have an extremely long life time. However, the use of these additional seeds does not appreciably affect the efficiency of reconstructing primary tracks from LHC collisions and it was not used in this analysis.

3.3.3 Pattern recognition

The pattern recognition module of the CTF algorithm is based on the Kalman filter method. The filter proceeds iteratively from the seed layer, starting from a coarse estimate of the track parameters provided by the trajectory seed, and includes the information of the successive detection layers one by one. On each layer, i.e., with every new measurement, the track parameters are known with a better precision, up to the last point, where they include the full tracker information. Each iteration of the Kalman filter is implemented in separate steps that are executed in the following order:

1. **Navigation:** given the track-candidate's updated state on a specific layer, a dedicated navigation service determines which layers of the tracking system, among those adjacent, may be crossed by the extrapolated trajectory. In determining the compatibility of a layer surface with the path of the particle, the uncertainty on the trajectory's parameters is considered. The navigation service is configurable and can be used to build the trajectory along and opposite to the momentum's direction. During the very first iteration of the trajectory building process, the input of the navigation service is the trajectory state defined by the seed.

2. **Search for compatible detectors:** after the track-candidate state is propagated to the surface of a compatible layer, the list of the layer's detectors that are compatible with the trajectory is determined. A detector is considered incompatible with the trajectory when the position of the trajectory's intersection with the detector surface is n standard deviations outside the detector's bounds. The default value of n is 3, but this parameter is configurable and it can be changed if track reconstruction efficiency is preferred to the algorithm's speed or vice versa. The propagation of the trajectory parameters (and the corresponding errors) on the sensor surface involve mathematically intensive operations and routines that are generally time-consuming [89]. Hence the code responsible for searching for the compatible detectors has been optimized as much as possible to limit the number of layer's sensors that are queried and additionally to assure an efficiency in finding the relevant sensors which is higher than 99%.

3. **Search for compatible measurements:** for each compatible detector, the direction of the trajectory on the sensor surface is used to calculate more accurately the drift of the ionization charge carriers inside the silicon bulk and to improve the estimation of the position of the pixel or strip clusters in the detector. A reconstructed clusters is considered compatible with the trajectory when the χ^2 of the predicted residual, as calculated in Eq. 3.25, is smaller than a configurable

parameter (that by default is equal to 30). In order to reduce the processing time, the residuals and the χ^2 are estimated only for those strip clusters that have a reasonable chance of being compatible with the trajectory. At the moment, there is no corresponding optimization responsible for finding compatible pixel clusters and therefore all the clusters in a pixel detector are analyzed.

4. **State update:** if a measurement is compatible with the trajectory, the track-candidate is extended by this additional point and the trajectory's parameters are updated to the filtered state, which is defined by Eq. 3.20. When several measurements on the new layer are compatible with the predicted trajectory, several new track-candidates are created, one per compatible hit. Finally, one additional track-candidate is created, in which no measured position is added, to account for the possibility that the charged particle did not leave any hit on that particular layer. This fake hit is called an "invalid hit".

All resulting track-candidates are then grown in turn to the next compatible layer(s) and the procedure is repeated until either the outermost layer of the tracker is reached or a "stopping condition" is satisfied. To avoid an exponential increase of the number of candidates, only a limited number of these are retained at each step, based on their normalized χ^2 and the number of valid and invalid hits. Besides, the building of a trajectory can be stopped as soon as the uncertainty on its parameters is below a given threshold or a minimum number of measurements has been added: this kind of stopping conditions are typically used in the context of HLT, where the required accuracy on track parameters is often reached when 5 or 6 position measurements are added to the track-candidate, and the continuation of the pattern recognition to 12-13 hits would be a waste of CPU time.

When a trajectory is propagated to a given layer, the uncertainty of the predicted state has a direct effect on pattern recognition, since it determines the compatibility between the trajectory and nearby hits. The number of compatible hits found on a layer determines the increase of the number of trajectories to be propagated, as the initial trajectories are multiplied by the number of hits found. For particles moving through the barrel region of the Tracker, a large fraction of the seeds are composed of hits in the first three pixel layers (Sec. 3.3.3), and the parameters of these seeds are propagated from the third pixel layer to the first layer of TIB. The uncertainty on the position of the predicted state is significant at this stage (Fig. 3.5), due to the large gap between the two layers (approximately 13 cm) and to the short lever arm (about 6 cm) that is used to determine the trajectory's curvature by means of only three pixel measurements. The average uncertainty on the transverse r-ϕ coordinate is about 800 μm, twice the value of that on the longitudinal z coordinate. The ill defined trajectory's momentum from

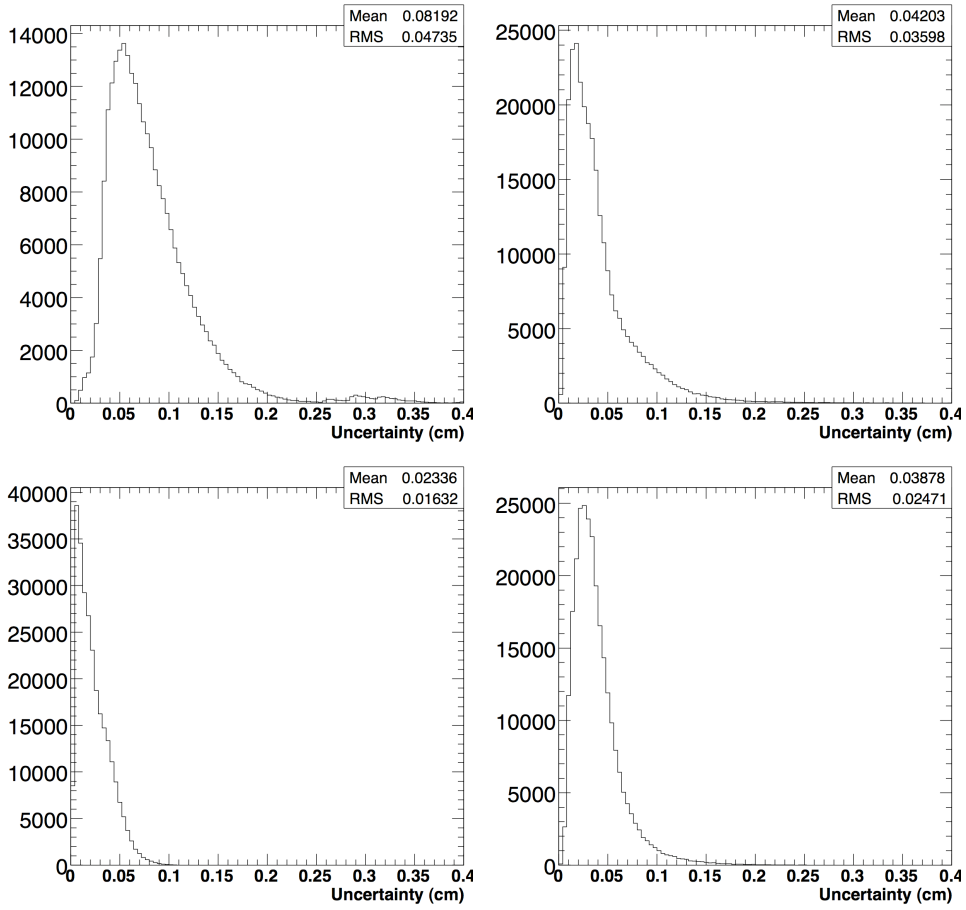

Figure 3.5: *Distributions of the uncertainties on the X (left column) and Y (right column) position coordinates of the trajectory's predicted state during the pattern recognition step of track reconstruction. Coordinates refer the local reference frames of the detectors. Upper row shows histograms for states propagated from the third pixel barrel layer onto the first TIB layer; histograms for states propagated from TIB layer 1 onto TIB layer 2 are on the bottom row.*

this early stage makes it difficult to extrapolate the trajectory within the transverse plane. When a measurement on the first TIB layer is included into the track-candidate, the uncertainty on the trajectory parameters becomes much smaller. In particular, the estimate of the curvature benefits significantly from the larger lever arm that now extends approximately over 20 cm. The position of the extrapolated trajectory on the second layer of TIB is estimated with an average precision of 230 μm in r-ϕ and 380 μm in z (Fig. 3.5, bottom). Compared with the first propagation, which was influenced principally by the uncertainty on the trajectory curvature, start-

ing from the second layer of TIB the accuracy of the propagated state is dominated by the precision of the last position measurement added to the track-candidate. For double-sided strip detectors, the z coordinate of a cluster is determined with a precision that is about 10 time worse compared to the r-ϕ coordinate. Besides, on single-sided detectors, the z coordinate is not measured at all and it is constrained only by the length of the sensor's strips (approximately 10 or 20 cm). For this reason, excluded the first extrapolation, the uncertainty on the z coordinate of the predicted state is larger than that on r-ϕ for all the layers of the Strip Tracker (Fig. 3.6).

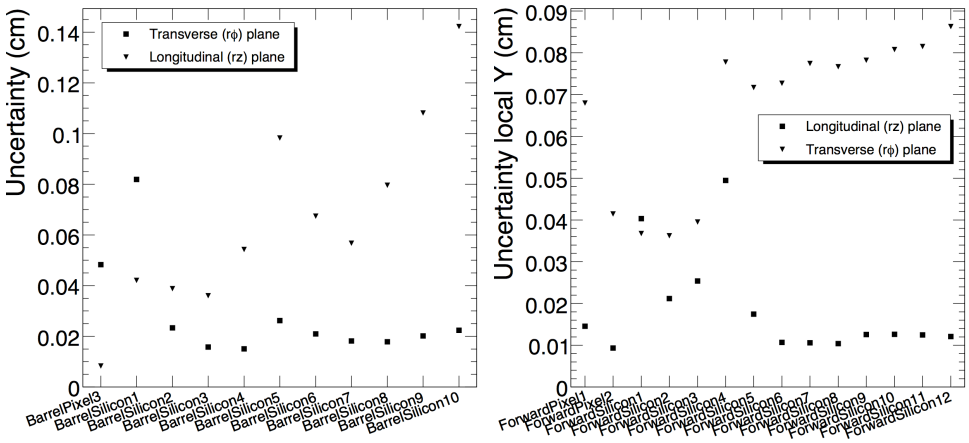

Figure 3.6: *Mean uncertainties of the predicted state in the barrel (left) and forward tracker (right) in the transverse (r-ϕ) and longitudinal (r-z) planes.*

The correlation between the uncertainty on the extrapolated trajectory state and the number of compatible measurements affects the distributions shown in Fig. 3.7: while more than half of the trajectories propagated from the third layer of the pixel barrel prove to be compatible with a spurious hit, this fraction is reduced approximately to one third for trajectories propagated from TIB layer 1 to TIB layer 2.

The specific geometry of the CMS Tracker, which consists of cylindrical barrel layers completed by disks in the two endcaps, facilitates the navigation between barrel layers during the pattern recognition stage of track reconstruction. However, the navigation is more complex in the barrel-endcap transition region, where a single layer can be connected to several others. For example, for a trajectory state that is defined on the Forward Pixel disk 2 and is pointing at high η, all three TID disks and the first three TEC ones could be compatible with the trajectory. When these layers are queried for compatible hits, each of these returns at least the invalid hit. As the propagation distance to some of these layers can be quite large, the uncertainties are comparatively large, and the probability of finding also

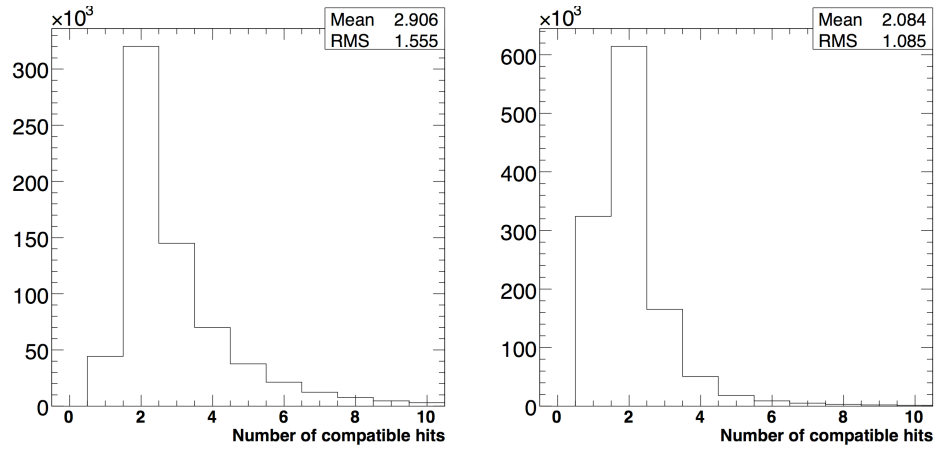

Figure 3.7: *Number of compatible hits (including invalid hit(s)) found per trajectory, when leaving the pixel barrel layer 3 (left) and TIB layer 1 (right). Simulated b-jets events with transverse momenta between 120 and 170 GeV/c are used, with low luminosity pile-up included*

spurious hit increases. Once the trajectory building reaches the endcap, the navigation is easy again and many of the trajectories containing spurious hits are quickly discarded, usually because they fail the condition on the minimum trajectory's χ^2 or that on the number of invalid hits.

Ambiguity resolution

Ambiguities in track-finding arise when the same track is reconstructed twice starting from different seeds or when a given seed develops into more than one final track-candidate. The ambiguity among *mutually exclusive* track-candidates must be resolved in order to avoid that the same charged particle is counted twice, or even multiple times.

The pattern recognition module of the CTF algorithms aims to identify track duplicates through an ancillary service, called trajectory cleaner, that determines the fraction of *shared* hits between two track-candidates:

$$f_{shared} = \frac{N^{hits}_{shared}}{min(N^{hits}_1, N^{hits}_2)} \tag{3.37}$$

where N^{hits}_1 (N^{hits}_2) is the number of hits in the first (second) track-candidate. If this fraction exceeds the cut value $c = 50\%$, the trajectory cleaner removes the track with the least number of hits from the candidates collection; if both tracks have the same number of hits, that with the highest χ^2 value is discarded. The procedure is repeated iteratively until all the pairs of track-candidates in the collection share less than half of their hits.

The probability that a track-candidate is a ghost produced by accidentally aligned hits decreases very quickly with the number of measurements on the trajectory and therefore the ambiguity resolution method is biased toward long trajectories. However, such a simple selection criterion, which is successful in most of the cases, fails in some special conditions, like during the reconstruction of high energy hadrons and electrons. As it is shown in section 3.4.1, sometimes the hits produced by secondary particles (e.g. electrons from conversions or the products of inelastic nuclear interactions) are incorrectly attached by the trajectory builder to the track-candidate corresponding to the primary particle. In this case, it is the shorter track-candidate, containing only the measurements produced by the primary particle, that should be preferred to the longer candidates that include also spurious measurements. According to a more conservative approach, neither of the candidates should be discarded and the ambiguity resolution should be post-poned to a later stage of the reconstruction sequence where additional information becomes available. The author has suggested to use a different configuration of the trajectory cleaner for the tracking of electrons and now the cut value c isn't fixed to 50% anymore, but is a configurable parameter of the tracking algorithm.

3.3.4 Final Track Fit

For each trajectory, the building stage results in a collection of hits and in an estimate of the track parameters. However, the full information is only available at the last hit of the trajectory and the estimate can be biased by constraints applied during the seeding stage. Therefore the trajectory is refitted using the standard Kalman filter and smoother approach described at the end of section 3.2.

The Kalman filter is initialized at the location of the innermost hit with an estimate obtained during seeding. The corresponding covariance matrix is scaled by a large factor in order to avoid any bias. The fit then proceeds in an iterative way through the full list of hits. For each valid hit the position estimate is re-evaluated again using the current values of the track parameters. This first filter is complemented with the smoothing stage: a second filter is initialized with the result of the first one (except for the covariance matrix, which is scaled with a large factor) and is run backward toward the beam line.

This filtering and smoothing procedure yields optimal estimates of the parameters at the surface associated with each hit and, specifically, at the first and the last hit of the trajectory. Estimates on other surfaces, e.g., at the impact point, are then derived by extrapolation from the closest hit.

Rejection of spurious measurements

The Final Track Fit module is also in charge of removing spurious measurements, or *outliers*, that could be included into the track-candidates during pattern recognition. These measurements can be correlated with the otherwise well-defined track, e.g. clusters from δ-rays, or uncorrelated, like hits due to nearby tracks or to electronic noise. Spurious measurements have generally larger residuals than the genuine signal hits produced by the passage of the charged particle through the detector. During the pattern recognition, it is appropriate to discard spurious hits determing the closeness of a measurement to the trajectory by means of the χ^2 of the predicted residuals (eq. 3.29). However, the χ^2 of the *smoothed* residual (eq. 3.36) is a more powerful decision criterion once the full information from the tracker is available to estimate the trajectory parameters on each sensor surface [74].

Once the Kalman smoother is executed, the Final Track Fit module identifies the track's measurement with the largest smoothed residual and removes the hit from the trajectory if its residual is above a configurable threshold C. The remaining measurements are re-fit with the usual mechanism and the trajectory hit with the next higher residual is found. The procedure is repeated iteratively until all the trajectory measurements have their residuals below the bound C. Every time a measurement is removed from the trajectory, it is replaced by an invalid hit: in this context, the invalid measurement is a sort of placeholder informing the Kalman smoother that, despite the lack of a genuine measurement on that specific sensor, the effects of material (like energy loss and multiple scattering) have still to be accounted during the fit.

The default value (20) for the parameter C has been chosen as the best tradeoff between rejecting the maximum fraction of outliers and removing the minimum number of genuine measurements. About 20% of outliers are pulled out from tracks reconstructed in high-density di-jets events, compared to less than 0.2% of good measurements that are removed [98].

While the above described method aims to identify spurious measurements using exclusively geometrical information, a more sophisticated approach is being developed. It exploits also additional attributes of the reconstructed hits, like its cluster charge and cluster width.

3.4 CTF track reconstruction performance

In this section the performance of the CTF algorithm is evaluated on different categories of simulated samples. Firstly, the efficiency and quality of the pattern recognition is analyzed on simple events containing back-to-back pairs of muons, pions or electrons; secondly, the precision of the

track-fitting is examined on the same samples through the resolution and pull distributions of the five reconstructed track parameters. Finally, the efficacy of track-finding and track-fitting is re-evaluated on different types of multi-track events that simulate LHC collisions.

A service known as *Track Associator By Hits* has been implemented for evaluating efficiency, resolutions and the other quantities related to the performance of the track reconstruction algorithm. The associator matches a simulated track to the corresponding reconstructed one if more than 50% of their measurements are connected. The association between the hits of the simulated particles and the clusters in the silicon detectors is possible because the simulation software traces the particles responsible of the signal on each channel of the Tracker. Strips and pixels that are fired by electronic noise are known as well.

3.4.1 Efficiency and fake rate

Both *global* and *algorithmic* efficiencies are analyzed. The former represents the fraction of charged particles, among all the simulated ones, that have a partner within the collection of reconstructed tracks; the latter considers only the simulated particles that leave at least three measurements on three different layers of the Tracker and, therefore, are track-able by the reconstruction algorithm. While the first kind of efficiency is affected also by properties of the detector, like its geometrical acceptance or its material budget, the algorithmic efficiency is exclusively related to the ability of the tracking algorithm to perform the track-finding task.

The effectiveness of the pattern recognition module is quantified also estimating its hit-finding efficiency and the purity of the measurements on-track: the former is the fraction of hits produced by the simulated particle that are included in the corresponding reconstructed track; the latter is the fraction of not-spurious measurements among all the hits of a track.

The fake rate is defined as the fraction of reconstructed tracks that don't have an associated particle among the simulated ones. This quantity represents the probability that a track provided by the reconstruction algorithm is either a combination of uncorrelated measurements or a genuine trajectory that is reconstructed including a large number of spurious hits. The fake rate is evaluated both requesting that the hit purity is greater than 50% and greater than 75%.

Muons

Muons are the charged particles that are best reconstructed in the Tracker. They mainly interact with the silicon detector through ionization and their

energy loss by bremsstrahlung is generally negligible, except when muons are produced with an initial energy higher than about 100 GeV. Therefore these particles usually cross the whole volume of the tracking system, producing detectable hits on *all* the sensitive layers of the apparatus that they intersect. Finally, muon trajectories are scattered exclusively by multiple-scattering, whose effects are straightforward to be included inside the Kalman filter formalism. For isolated muons with a transverse momentum between 1

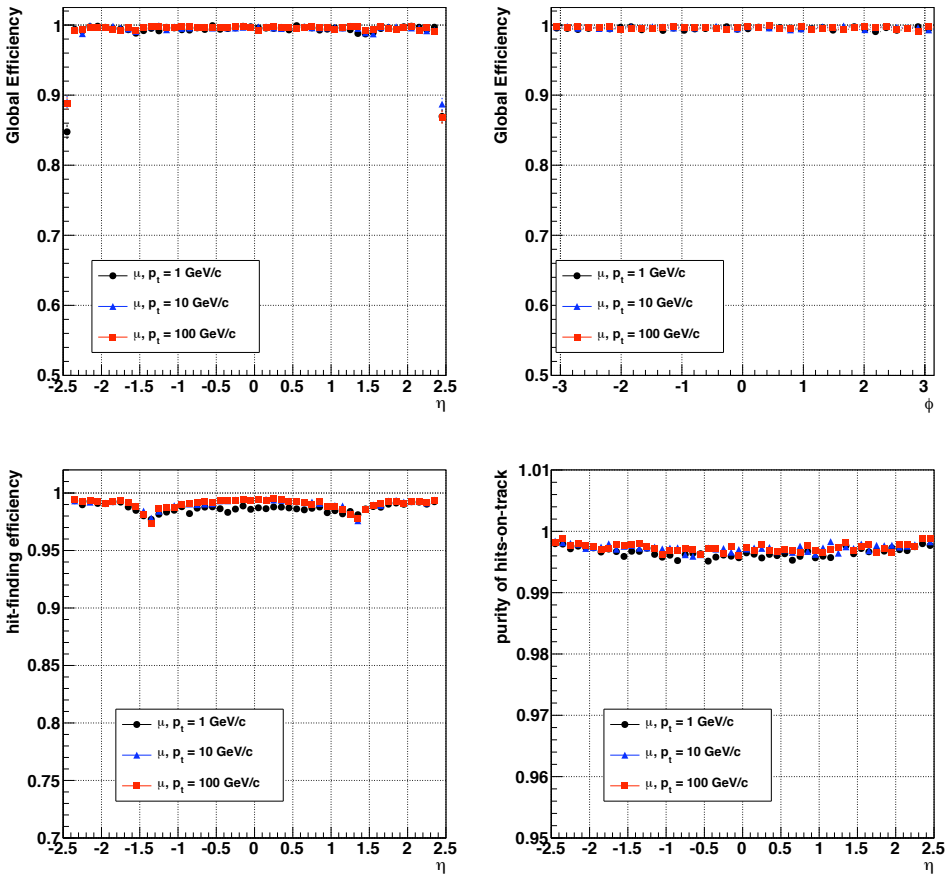

Figure 3.8: *Global track reconstruction efficiency as a function of η (top-left) and φ (top-right) for muons of transverse momenta of 1, 10 and* 100 *GeV/c. Hit-finding efficiency and purity of the measurements on-track (bottom).*

GeV/c and 100 GeV/c, the global tracking efficiency is higher than 99% in the full η-range of the Tracker acceptance and the efficiency does not depends on the ϕ direction of the particles (Fig. 3.8, top). The average hit-finding efficiency (Fig. 3.8, bottom-left) is higher than 99% both in the barrel and in the endcaps; it is still above 97% in the barrel-transition region

where the layer navigation is more complicated (Sec. 3.3.3) and the material budget of the Tracker is more significant. The trajectory contamination due to spurious hits produced by electrical noise or δ-rays is lower than per-mille (Fig. 3.8, bottom-right).

Pions

Charged pion particles undergo multiple scattering and energy loss by ionization while they cross the tracker volume, like muons do. However, pions, and all hadrons in general, are also subject to elastic and inelastic *nuclear* interactions. The former effect introduces tails in the distribution of the scattering angle that is expected assuming the multiple scattering as the only contributor to the stochastic deflection of the particles. In the current implementation of the CTF algorithm, the description of the process noise (eq. 3.2) doesn't account for elastic nuclear interactions and this brings to a lower hit finding efficiency during the reconstruction of hadrons compared to muons (Fig. 3.9, left).

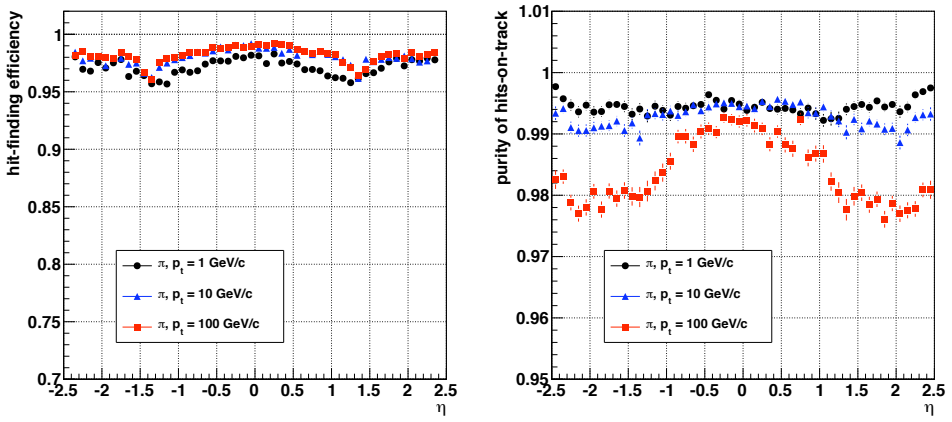

Figure 3.9: *Hit-finding efficiency (left) and hit-purity (right) for pions of transverse momenta of 1, 10 and* 100 *GeV/c.*

A possible effect of an inelastic nuclear interaction between a hadron and the detector's material is the transformation of the incoming particle into two or more secondary particles. Depending on the location within the Tracker volume where this process occurs, the number of measurements produced by the primary particle on the sensitive layers of the detector can be pretty small, deteriorating the precision with which the trajectory parameters can be estimated (Sec.3.4.2). Besides, when the number of crossed layers is less than three, the particle trajectory cannot be reconstructed at all by the CTF

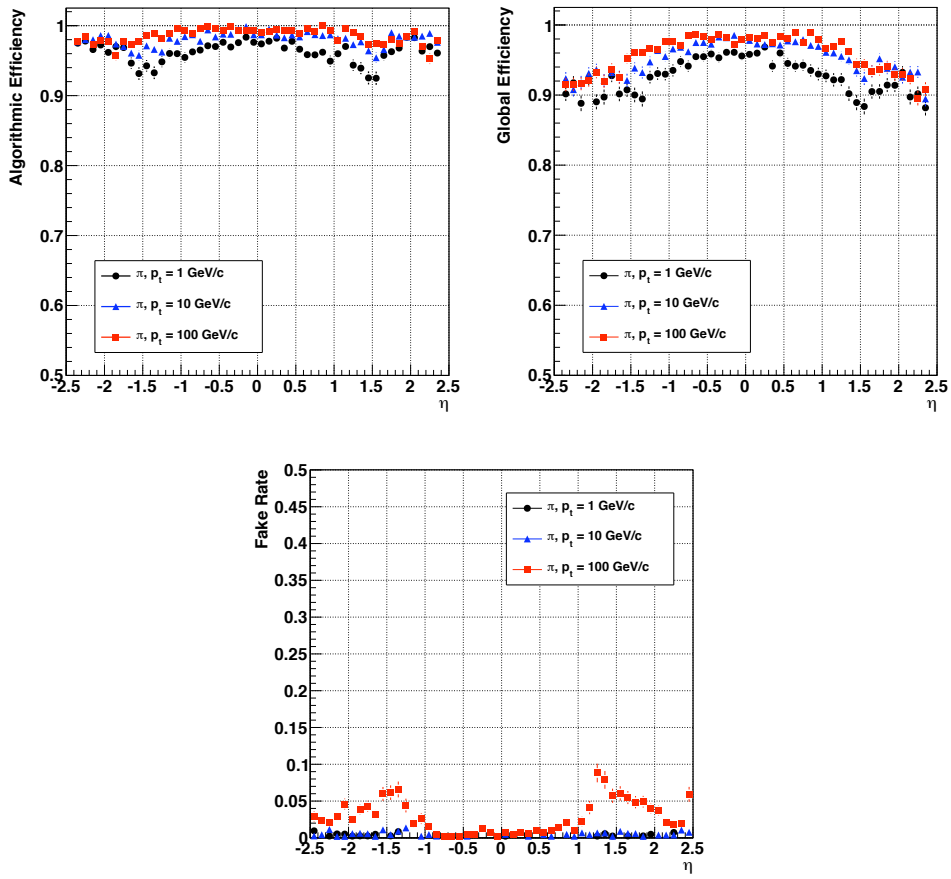

Figure 3.10: *Algorithmic tracking efficiency (top-left), global tracking efficiency (top-right) and fake rate (bottom) as a function of* η *for pions of transverse momenta of 1, 10 and* 100 *GeV/c. Results refer to reconstructed tracks that are requested to have a hit-purity greater than 50%.*

algorithm. Inelastic nuclear interactions are the main source of tracking inefficiency for hadrons, in particular in those regions of the Tracker where the material budget is more significant and the probability that particles undergo nuclear interactions is higher (Fig. 3.10, top-right). Therefore, elastic nuclear interactions and the consequent hit-finding inefficiency, not only reduce the global efficiency of the pattern recognition software, but they also lower the *algorithmic* efficiency: depending on the η-bin, up to 9% of the simulated particles are not reconstructed (Fig. 3.10, top-left), despite the fact that they leave measurements on three or more layers of the tracking system. This effect is more significant for low energy hadrons due to their higher cross section for nuclear interactions [99].

The secondary particles produced in inelastic processes affect also the hit-purity of the reconstructed trajectories. The products of the nuclear interaction are usually emitted from the nuclear vertex with directions roughly tangent to the trajectory of the incoming particle and therefore it is not uncommon that the trajectory builder combines the hits of the primary particle with those of a secondary particle. The degradation in the purity

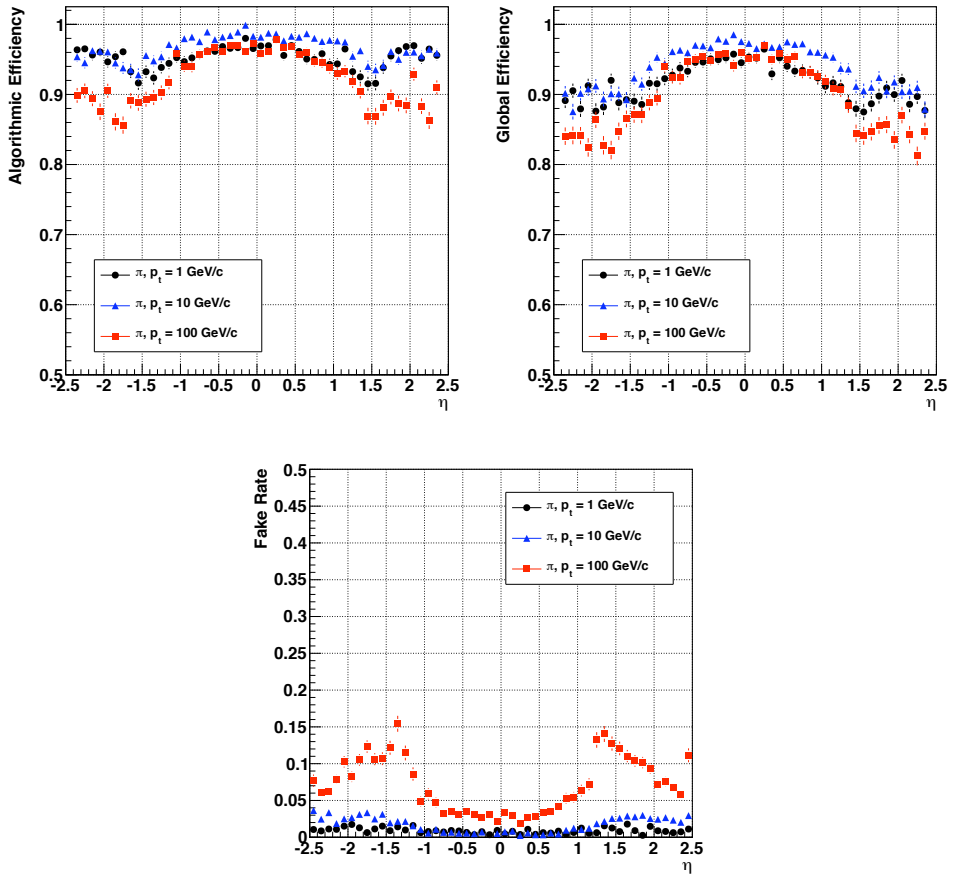

Figure 3.11: *Algorithmic tracking efficiency (top-left), global tracking efficiency (top-right) and fake rate (bottom) as a function of* η *for pions of transverse momenta of 1, 10 and* 100 *GeV/c. Results refer to reconstructed tracks that are requested to have a hit-purity greater than 75%.*

of the measurements on track (Fig. 3.9, right) is more accentuated for high energy hadrons than for low energy ones: indeed the probability that the trajectory builder merges together separate tracks increases with the energy of the primary hadrons because of the larger number of secondary particles that are produced in the interaction and because of the smaller difference between the curvature of the primary particle and that of the secondary

particles.

The merging of hadron trajectories affects significantly the efficiency and fake rate distributions depending on which cut on the hit-purity is used in the definition of correctly-reconstructed tracks: Fig. 3.10 and Fig. 3.11 shows efficiency and fake rate distributions that are obtained requesting respectively a hit-purity of 50% and 75% for the reconstructed tracks. The degradation in efficiency and the increase in fake rate are correlated, as expected, and the variation is more significant for the sample with the highest energy particles. In general, the merging of separate trajectories during reconstruction is more common in the transition and endcap regions of the tracker, due to the higher material budget. In the barrel-endcap transition region, the probability that the reconstruction algorithm returns merged trajectories is even higher because of the large extrapolation length during iterations of the pattern recognition module in which the trajectory builder navigates from barrel to endcap layers. While the fake rate is generally lower than 1% for pions with a transverse momentum of 1 or 10 GeV/c, the probability that a 100 GeV/c pion is incorrectly reconstructed peaks to about 10% for particles produced with $|\eta|$ around 1.5, regardless of a looser cut on the hit-purity (Fit. 3.10, bottom).

Electrons

Electrons and positrons are subject to multiple scattering and energy loss by ionization, like the muon leptons, and they are not exposed to the nuclear interactions that affect significantly the reconstruction of charged hadrons. Due to the remarkable material budget of the tracking system, which ranges between 0.4 and 1.8 radiation lengths depending on the particle direction (Sec. 2.2.4), most of the electrons produced in LHC collisions loose a large fraction of their energy by means of bremsstrahlung radiation before they reach the outer layers of the Silicon Tracker. In some cases, the radiated photons have enough energy to convert into electron-positron pairs and initiate authentic electromagnetic showers inside the Tracker volume.

Although the mechanism responsible for the bremsstrahlung radiation is completely different from that of nuclear interaction for hadrons, it influences the performance of the track reconstruction algorithm in a similar way. The radiation of soft photons usually doesn't prevent the reconstruction of electrons, however it introduces tails in the distribution of the electron's energy loss that are not taken into account by the CTF algorithm in its default configuration. The principal consequence of this aspect is a decrease of the hit-finding efficiency, analogous to the reduction caused by elastic nuclear interactions for hadrons. The radiation of high energy photons has an impact on the reconstruction of the electrons which is similar to that of the inelastic nuclear interactions on the reconstruction of charged hadrons:

firstly, when an electron looses most of its energy before reaching the outer layer of the tracker, there is a reduction in the number of measurements, recorded by the detector, that are really useful to determine the trajectory's parameters at the collision point; secondly, when a radiated photon converts or creates an electromagnetic shower, the trajectory builder can fail to properly recognize the hits of the secondary particles and may try to merge the trajectory of the primary electron with other uncorrelated track segments. The latter effect further deteriorates the precision with which the

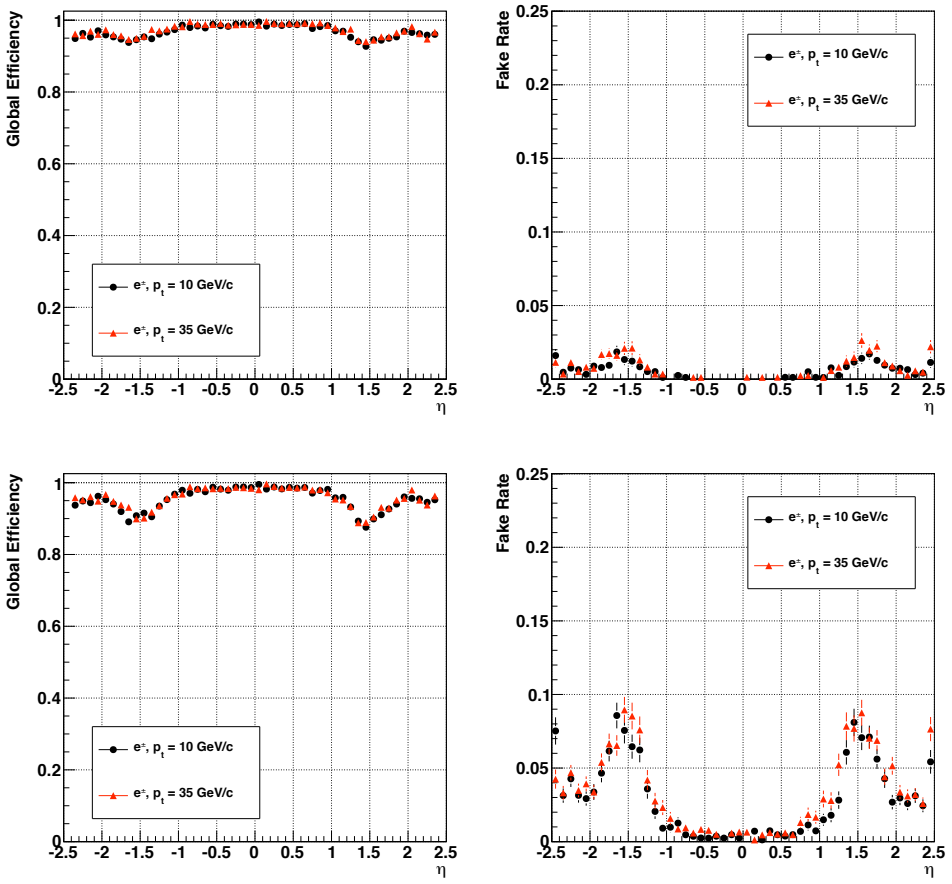

Figure 3.12: *Global tracking efficiency (left) and fake rate (right) for electrons and positrons of transverse momenta of 10 and 35 GeV/c. The histograms are produced with two different requirements on the hit-purity of the reconstructed tracks: 50% purity (top) and 75% purity (bottom).*

trajectories corresponding to early-radiating electrons are reconstructed and it seems to be the principal source of charge misidentification for electrons.

In its default configuration, the CTF algorithm efficiency for reconstructing

electron trajectories that have a hit-purity greater than 50% is above 99% in the barrel region of the Tracker. The efficiency is generally greater than 95% in the full η range of the geometrical acceptance (Fig. 3.12, top-left). If the hit-purity of the reconstructed trajectories is required to be greater than 75%, the global tracking efficiency decreases to about 90% in the transition region, but remains almost unchanged in the barrel and the endcap regions (Fig. 3.12, bottom-left). For both cuts on the hit-purity, the fake rate is correlated to the tracking inefficiency and peaks in the barrel-endcap transition region where the amount of material traversed by the electrons is the largest and the probability of producing electromagnetic showers is maximum (Fig. 3.12, right-column).

Better performance in term of efficiency and fake rate can be obtained using an ad-hoc configuration of the CTF trajectory builder that is optimized for electron reconstruction. In particular, the service used to extrapolate the trajectory parameters between the layers of the Tracker can activate the inclusion of the process noise term due to the bremsstrahlung radiation. Such a configuration of the trajectory builder implies a much slower execution of the track reconstruction sequence that is not affordable for the reconstruction of all the tracks in a generic LHC collision event. Nevertheless it is used for the reconstruction of specific electron candidates within small regions of interest of the Tracker that are identified by the electromagnetic calorimeter [100].

3.4.2 Track reconstruction resolution

This section analyzes and quantifies the level of precision that is achieved in estimating the parameters of the reconstructed trajectories. In the context of the reconstruction software of CMS, the five parameters used to describe a track are: d_0, z_0, φ, $\cot\vartheta$ and the transverse momentum p_t. The track parameters are defined at the point of closest approach of the track to the beam axis (this point is called the impact point); d_0 and z_0 hence measure the coordinate of the impact point in the transverse and longitudinal plane ($d_0 = y_0 * \cos\varphi - x_0 * \sin\varphi$, where x_0 and y_0 are the transverse coordinates of the impact point), φ is the azimuthal angle of the momentum vector of the track at the impact point and ϑ is the polar angle.

For each of the five track parameters, the resolution is estimated as a function of the pseudorapidity of the simulated charged particle. In every η-bin, the resolution is expressed both as the σ of a gaussian fit of the residuals' distribution[1] and as the RMS of the same distribution. The pulls of the five reconstructed parameters are analyzed to verify whether the uncer-

[1]The residual is the difference between the parameter of the reconstructed track and the corresponding value of the simulated particle's parameter.

tainty that is provided for each parameter by the final track fit is a sensible measure of the error of the estimator or not.

Resolutions and pull plots are presented for the same eight samples of isolated muons, pions and electrons that have been used in the previous section to study the efficiency and the fake rate of the reconstruction algorithm.

Muons

Muons are the charged particles whose trajectories are estimated with the best precision in the silicon tracking system of CMS. Since they don't undergo inelastic nuclear interactions and the amount of energy that radiate by bremsstrahlung is negligible (for transverse momenta up to about 100 GeV/c), muons that are produced in LHC collisions cross the whole Tracker volume and therefore can be reconstructed with the largest lever arm available.

Figure 3.13 shows the resolution of the five track parameters for samples of isolated muons with p_{T} of 1, 10 and 100 GeV/c. The resolutions of the transverse and longitudinal impact parameters d_0 and z_0 are shown in the first two plots of Figure 3.13. At high momentum, the d_0 resolution is fairly constant and is dominated by the resolution of the first hit in the pixel detector. At lower momenta, the d_0 resolution is progressively degraded by multiple scattering, until the latter becomes dominant. The z_0 resolution of high momentum tracks is also dominated by the resolution of the first pixel hit, with multiple scattering dominating at low momentum. The improvement of the z_0 resolution up to a pseudorapidity of $|\eta| = 0.4$ can be attributed to the beneficial effect of charge interpolation in the position estimate of pixel clusters (Sec. 2.2.1): in the barrel, as the angle with which the tracks cross the pixel layers increases the clusters become wider, distributing the signal on more than one pixel and improving the position resolution. Similar considerations about the relation between track momentum and effect of multiple scattering and between track pseudorapidity and pixel hit precision are valid also for the plots of the φ and $\cot\vartheta$ parameters, which are shown in the second and third histogram of the same figure. The resolution of the transverse momentum is shown in the last plot of Figure 3.13. At high momentum (100 GeV/c), the resolution is around 1-2% up to a pseudo-rapidity of $|\eta| = 1.6$; for higher values of $|\eta|$ the lever arm used for the measurement is reduced. The degradation around $|\eta| = 1.0$ and beyond is due to the gap between the barrel and the end-cap disks (Fig. 2.1) and to the lower hit resolution of the last hits of the track measured in the TEC ring 7 with respect to the hit resolution in the TOB layers 5 and 6 (Tab. A.8 and A.10). At a transverse momentum of 100 GeV/c, the material in the tracker accounts for between 20 and 30% of the transverse momentum resolution; at lower

momenta, the resolution is dominated by multiple scattering and its distribution reflects the amount of material traversed by the track. In all the five reported resolution plots, the difference between the sigma of the gaussian fit and the RMS of the residuals distributions is generally very small for the 10 and 100 GeV/c samples, indicating that the distributions of the track parameter errors are properly described by single gaussian functions and have few entries populating the tails. The worse discrepancy between sigma and RMS is about 40% and regards the resolution of d_0 for the 1 GeV/c sample: for tracks with $\eta = 0$, the sigma of the fit is equal to 70 μm compared to a RMS of about 100 μm.

The precision with which the CMS implementation of the Kalman filter and smoother estimates the *uncertainty* on the reconstructed track parameters is analyzed through the pull distributions of all the five trajectory parameters (Fig. 3.13). The pull is the distribution of the parameter residual divided by the uncertainty on the reconstructed parameter, as it is provided by the final track fit. If the estimator is unbiased and the parameter uncertainty is evaluated correctly, the pull has a gaussian distribution with mean equal to 0 and sigma equal to 1. All the gaussian functions, obtained from the fit of the pull distributions of the five track parameters, have standard deviations within the range [0.85,1.10], indicating that the parameter errors are generally correctly estimated, but there is still room for improvement. Like for the tails in the resolution distribution, also for the pulls the worse results are those corresponding to the low energy sample. This could be an indication of some limits in the current description of the multiple scattering effects.

Another quantity that can be used to estimate the quality of the track fit is the distribution of the normalized-χ^2 of the reconstructed track . Unlike the pull distribution, the χ^2 includes also contributions from the off-diagonal terms of the error covariance matrix (eq. 3.28 and 3.29). The bottom-right plot of Fig. 3.13 represents the average normalized-χ^2 as a function of the pseudorapidity of the reconstructed tracks in the three muon samples. The values corresponding to 1 GeV/c tracks produced in the barrel region show the worse discrepancy (of about 20%) with the theoretical curve (i.e. a flat distribution equal to 1). The results for the 10 and 100 GeV/c samples are compatible with the expected values within 10%.

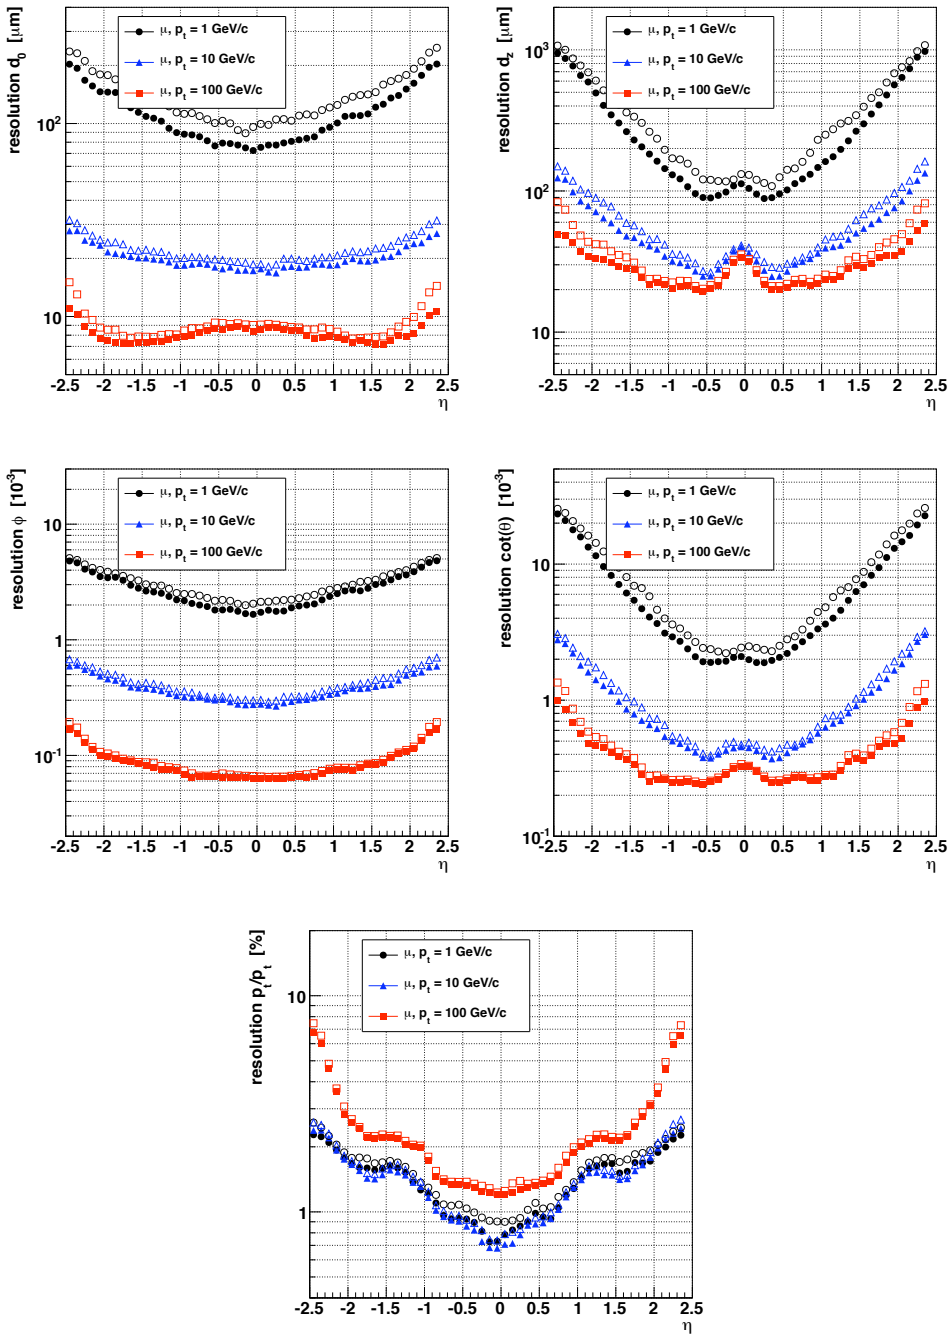

Figure 3.13: *Resolution, as a function of pseudorapidity, of the five track parameters for isolated muons with transverse momenta of 1, 10 and 100 GeV/c. From top to bottom and left to right: transverse momentum, φ, $\cot\vartheta$, transverse and longitudinal impact parameter. For each bin in η, the full symbol corresponds to the sigma of a gaussian fit of the residuals distribution and the empty symbol represents the RMS of the same distribution.*

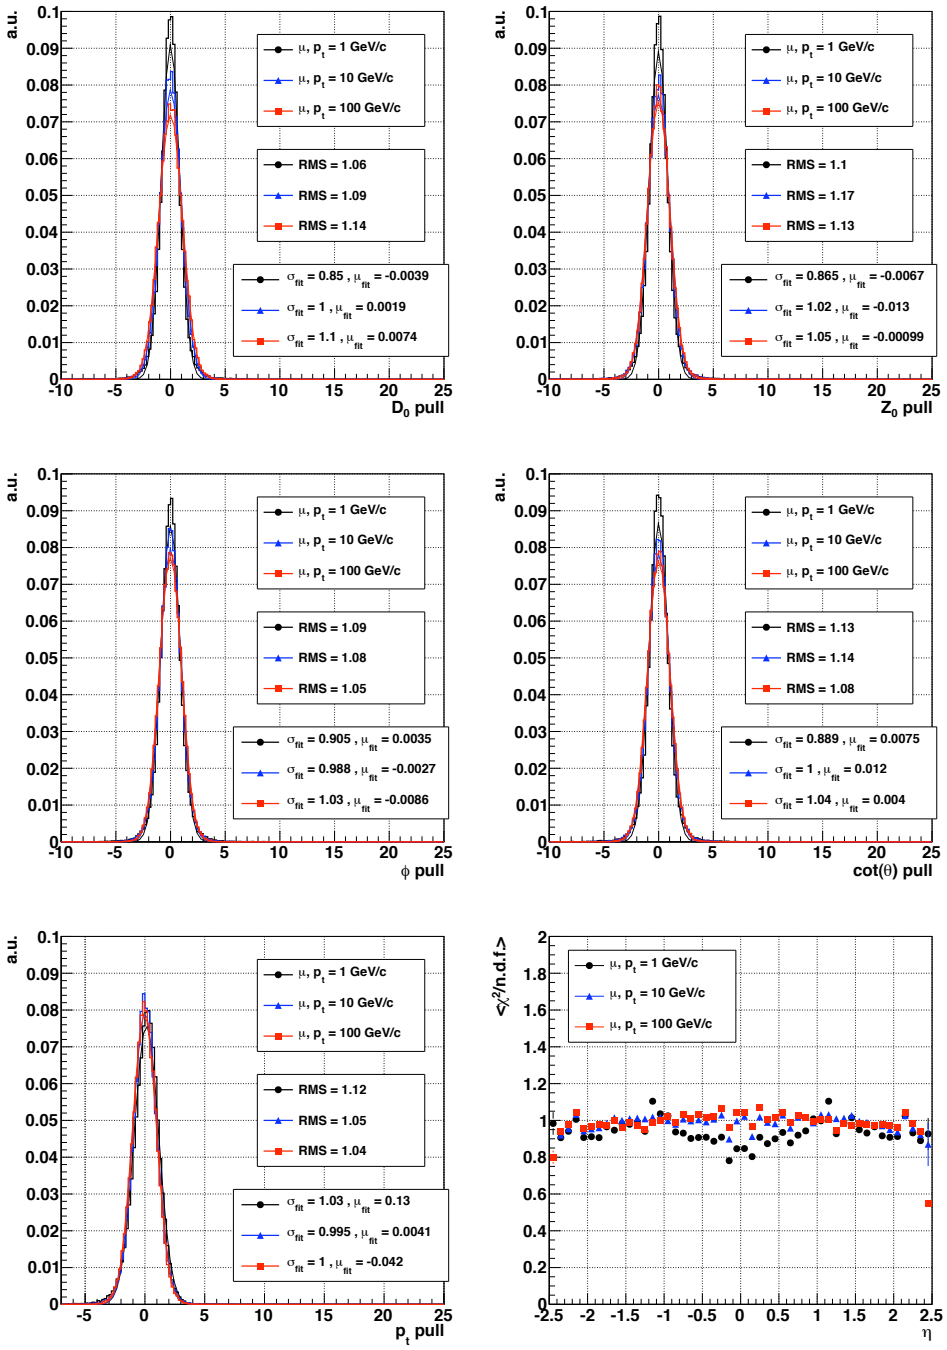

Figure 3.14: *The first five plots correspond to the pull distributions of the five track parameters for isolated muons with transverse momenta of 1, 10 and 100 GeV/c. From top to bottom and left to right: transverse momentum, φ, $\cot\vartheta$, transverse and longitudinal impact parameter. Both the σ and mean of a single-gaussian fit and the RMS of the pull distribution are printed. The sixth plot is the average normalized-χ^2 of reconstructed tracks as a function of pseudorapidity.*

Pions

Charged pions that don't undergo nuclear interactions behave similarly to muons being subject to the same multiple scattering effects and to the same mechanism of energy loss by ionization. The trajectories of this subset of pions are reconstructed by the CTF algorithm with a precision which is very close to that achieved for muons. The five plots in Fig. 3.15 show the resolutions of the five track parameters. As expected, the results obtained with a gaussian fit of the distributions of the residuals are very close to those showed for muon trajectories in Fig. 3.13. However, the resolutions obtained using the RMS show a different pattern with respect to the same resolutions evaluated for the muon samples. The RMS of transverse impact parameter, φ direction and transverse momentum residuals are up to a factor 2 wider than the corresponding results from a gaussian fit. These tracks whose parameters are measured less precisely are mainly pions that interact inelastically; since these particles don't reach the outer layers of the tracking system, their trajectories can only be measured using a shorter lever arm. Considering only the charged pions that succeed to cross the whole Tracker volume without inelastic interactions, the difference between RMS and gaussian fit resolutions is significantly reduced (Fig. 3.17). While the length of the lever arm affects directly the estimate of the transverse momentum, it influences also the measurement of those quantities, like d_0 and φ, that are obtained propagating the trajectory state within the *transverse* plane from the innermost sensitive surface of the tracking system to the impact point.

The pull distributions of the five trajectory parameters and the average normalized-χ^2 are shown in Fig. 3.16. Like for the muons samples, the sigmas of the gaussian fit of the pulls are close to the expected value 1, ranging between 0.86 and 1.14. However, the RMS of the same pull distributions are generally wider, extending between 1.18 and 1.65. The uncertainty on the estimates of the reconstructed trajectory parameters is similarly underestimated due to the effects of nuclear elastic interactions that are not incorporated in the current version of the trajectory propagator (see also analysis of pion hit-finding inefficiency in Sec. 3.4.1).

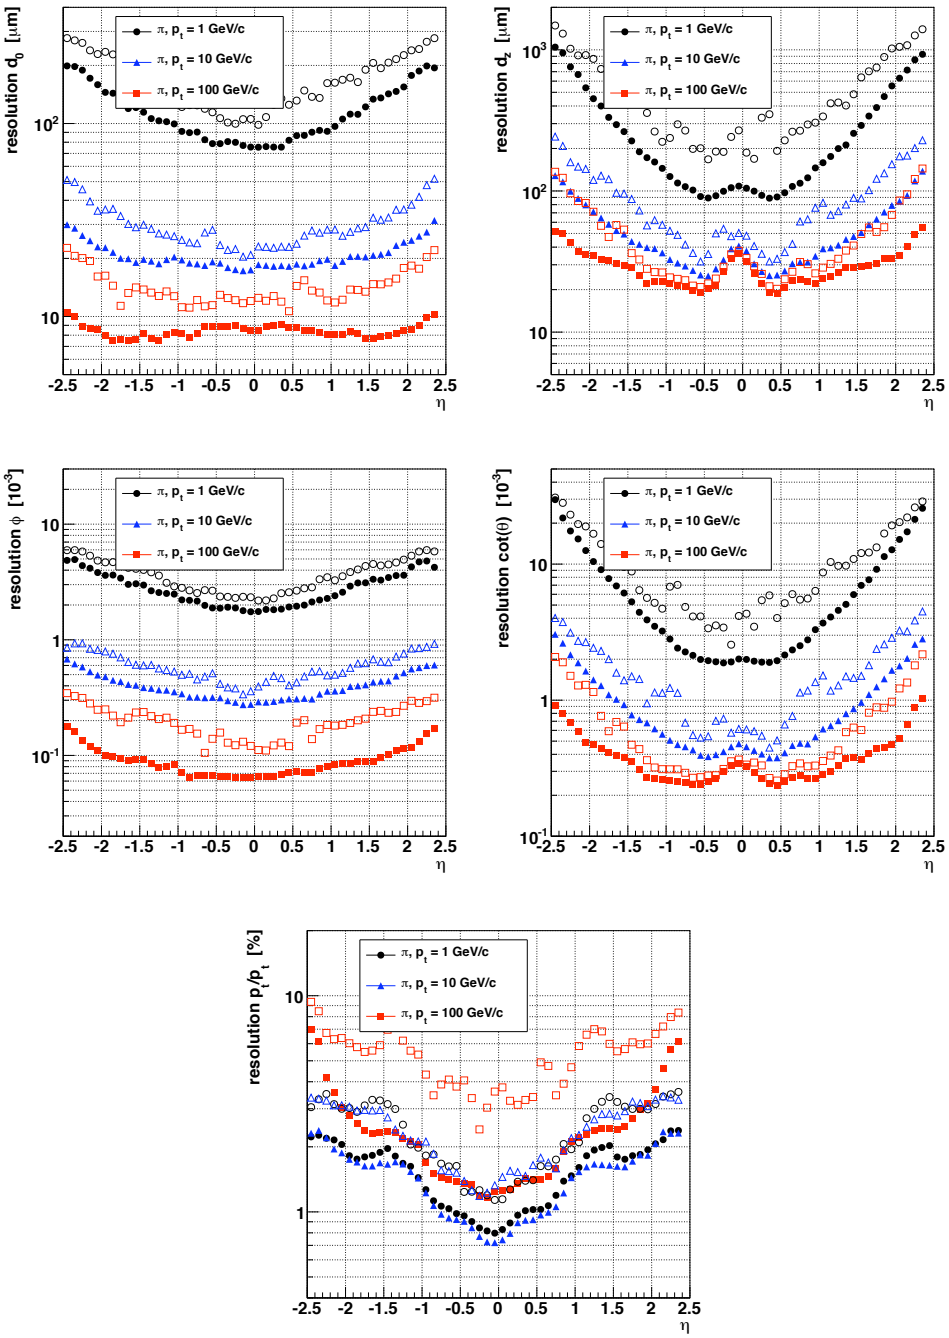

Figure 3.15: *Resolution, as a function of pseudorapidity, of the five track parameters for isolated pions with transverse momenta of 1, 10 and 100 GeV/c. From top to bottom and left to right: transverse momentum, φ, $\cot\vartheta$, transverse and longitudinal impact parameter. For each bin in η, the full symbol corresponds to the sigma of a gaussian fit of the residuals distribution and the empty symbol represents the RMS of the same distribution.*

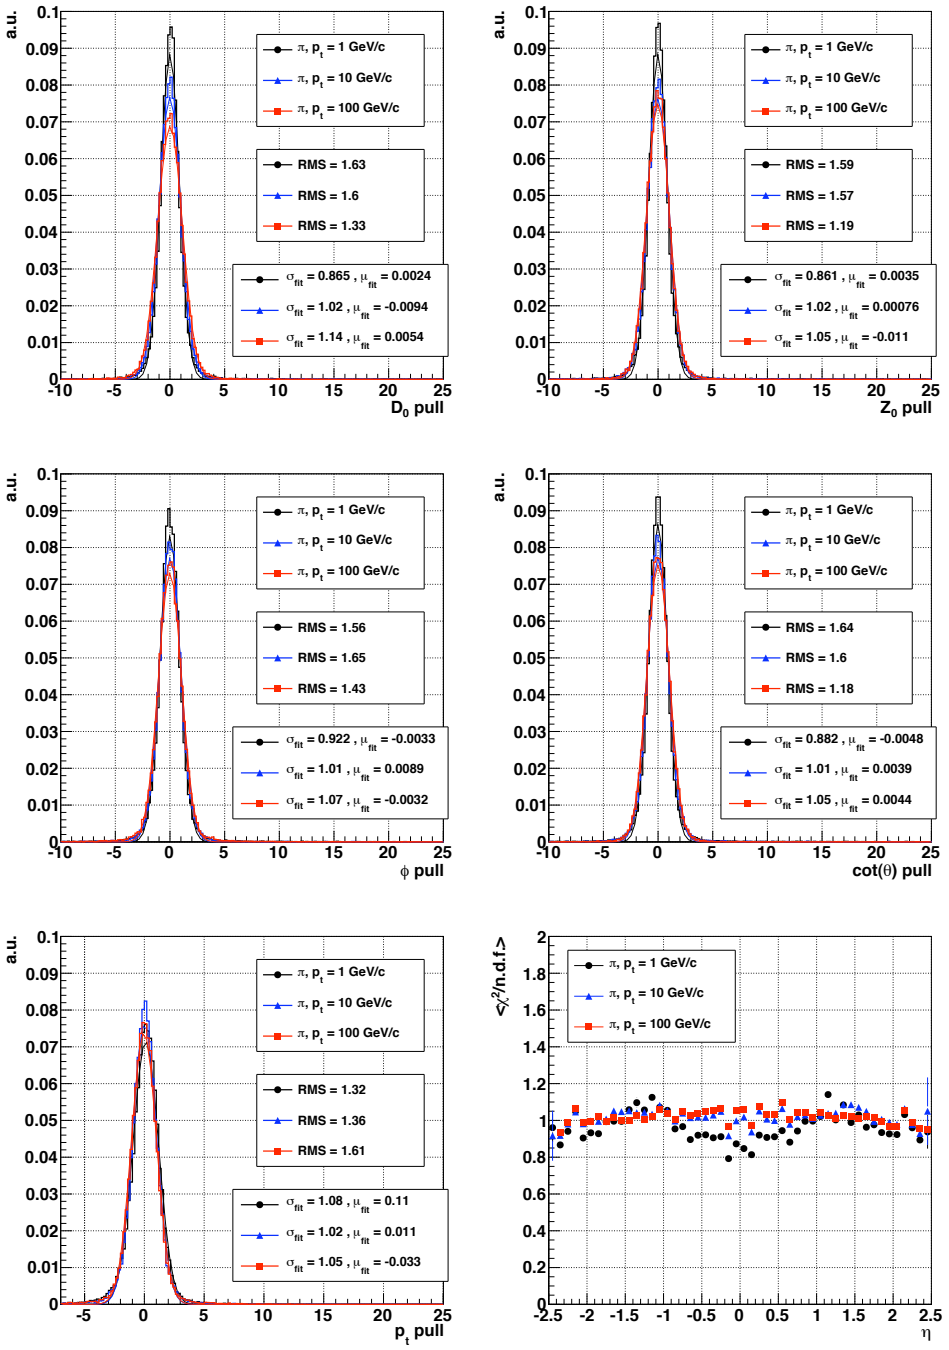

Figure 3.16: *The first five plots correspond to the pull distributions of the five track parameters for isolated pions with transverse momenta of 1, 10 and 100 GeV/c. From top to bottom and left to right: transverse momentum, φ, $\cot\vartheta$, transverse and longitudinal impact parameter. Both the σ and mean of a single-gaussian fit and the RMS of the pull distribution are printed. The sixth plot is the average normalized-χ^2 of reconstructed tracks as a function of pseudorapidity.*

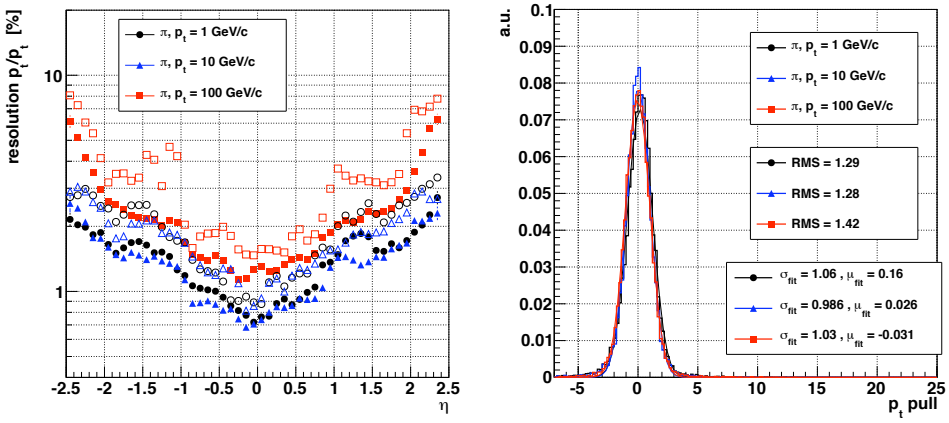

Figure 3.17: *On the left: resolution, as a function of pseudorapidity, of the reconstructed transverse pions for isolated pions with transverse momenta of 1, 10 and 100 GeV/c; for each bin in η, the full symbol corresponds to the sigma of a gaussian fit of the residuals distribution and the empty symbol represents the RMS of the same distribution. On the right: pull distribution of the transverse momentum for the same simulated samples of pions; both the σ and mean of a single-gaussian fit and the RMS of the pull distribution are printed. The two plots have been filled only for those simulated hadrons that reached the outer layers of the tracker without undergoing inelastic nuclear interactions.*

Electrons

The resolutions of the five trajectory parameters are shown in Fig. 3.18 for the two samples of isolated electrons with transverse momenta of 10 and 35 GeV/c. Like for pions, the resolution of the three quantities measured in the transverse plane (d_0, φ and transverse momentum) are significantly different whether they are evaluated with the RMS of the residuals or with a gaussian fit. Due to the effect of bremsstrahlung radiation, some electrons don't reach the outer layers of the tracking system and, like for pions undergoing inelastic nuclear interactions, can be reconstructed only using a shorter lever arm. Besides, since the process noise due to bremsstrahlung is not implemented in the default version of Kalman filter and smoother used by the CTF algorithm, the estimates of the trajectory parameters that are reconstructed in the transverse plane have systematically underestimated uncertainties. Pull distributions of the five track parameters are shown in Fig. 3.19: the trajectory parameters estimated in the r-z plane have pull distributions centered around zero and with standard deviations differing from 1 by less than 15%; the pull distributions of d_0, φ and transverse momentum

measurement have RMS ranging between 3 and beyond 6; finally, the measurement of the transverse momentum is systematically underestimated, as shown by the asymmetric shape of the corresponding pull.

The asymmetric and not-gaussian process noise due to bremsstrahlung can be included in an extended version of the Kalman filter method which is known as Gaussian Sum Filter [93]. This different filter has been implemented in CMS [101] and is used for fitting the trajectories of the particles that are reconstructed in the specific electron reconstruction sequence driven by energy deposits in the electromagnetic calorimeter [100].

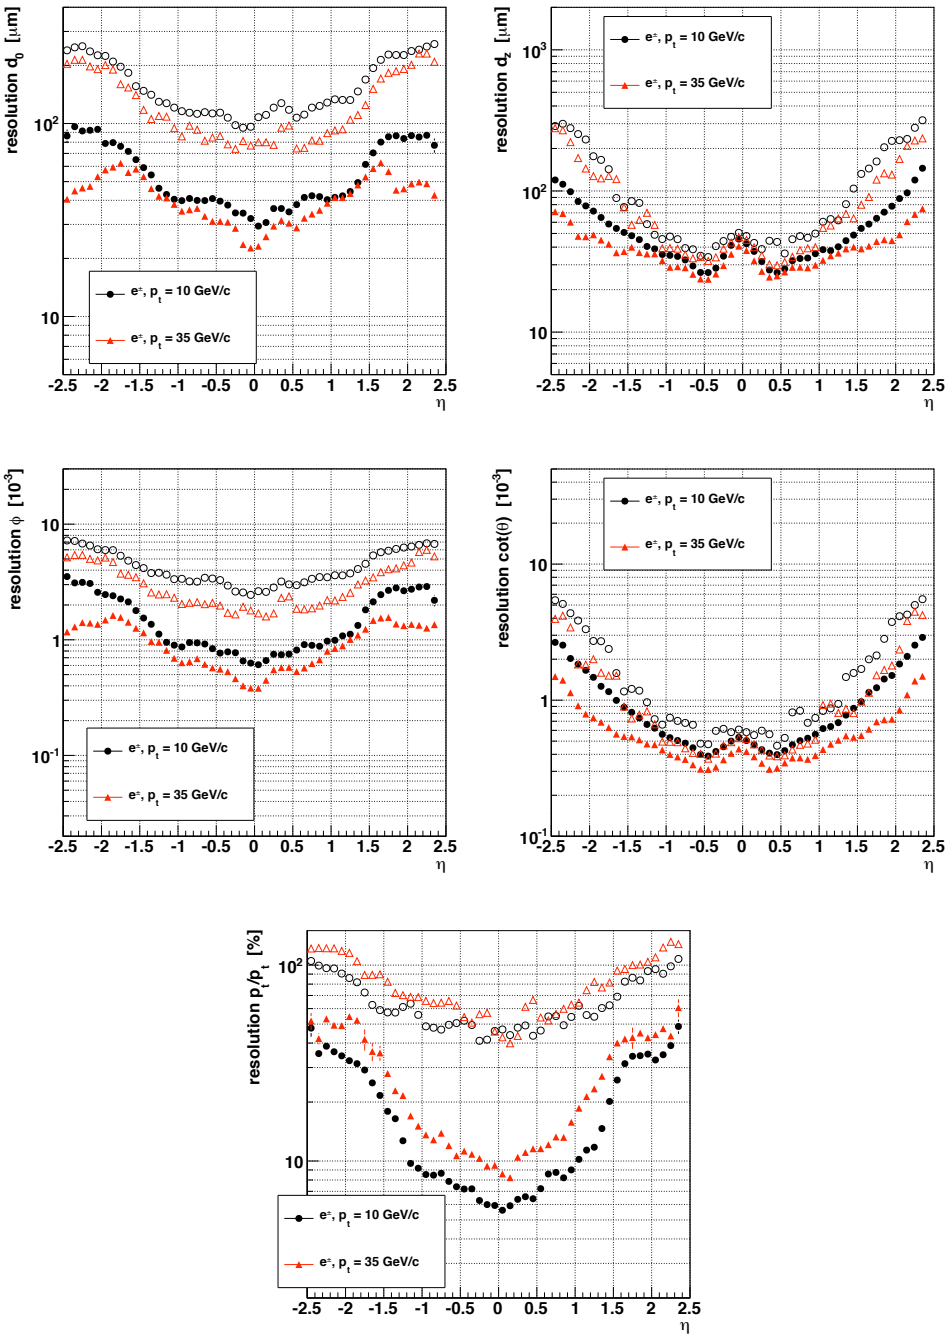

Figure 3.18: *Resolution, as a function of pseudorapidity, of the five track parameters for isolated electrons with transverse momenta of 10 and 35 GeV/c. From top to bottom and left to right: transverse momentum, φ, $\cot\vartheta$, transverse and longitudinal impact parameter. For each bin in η, the full symbol corresponds to the sigma of a gaussian fit of the residuals distribution and the empty symbol represents the RMS of the same distribution.*

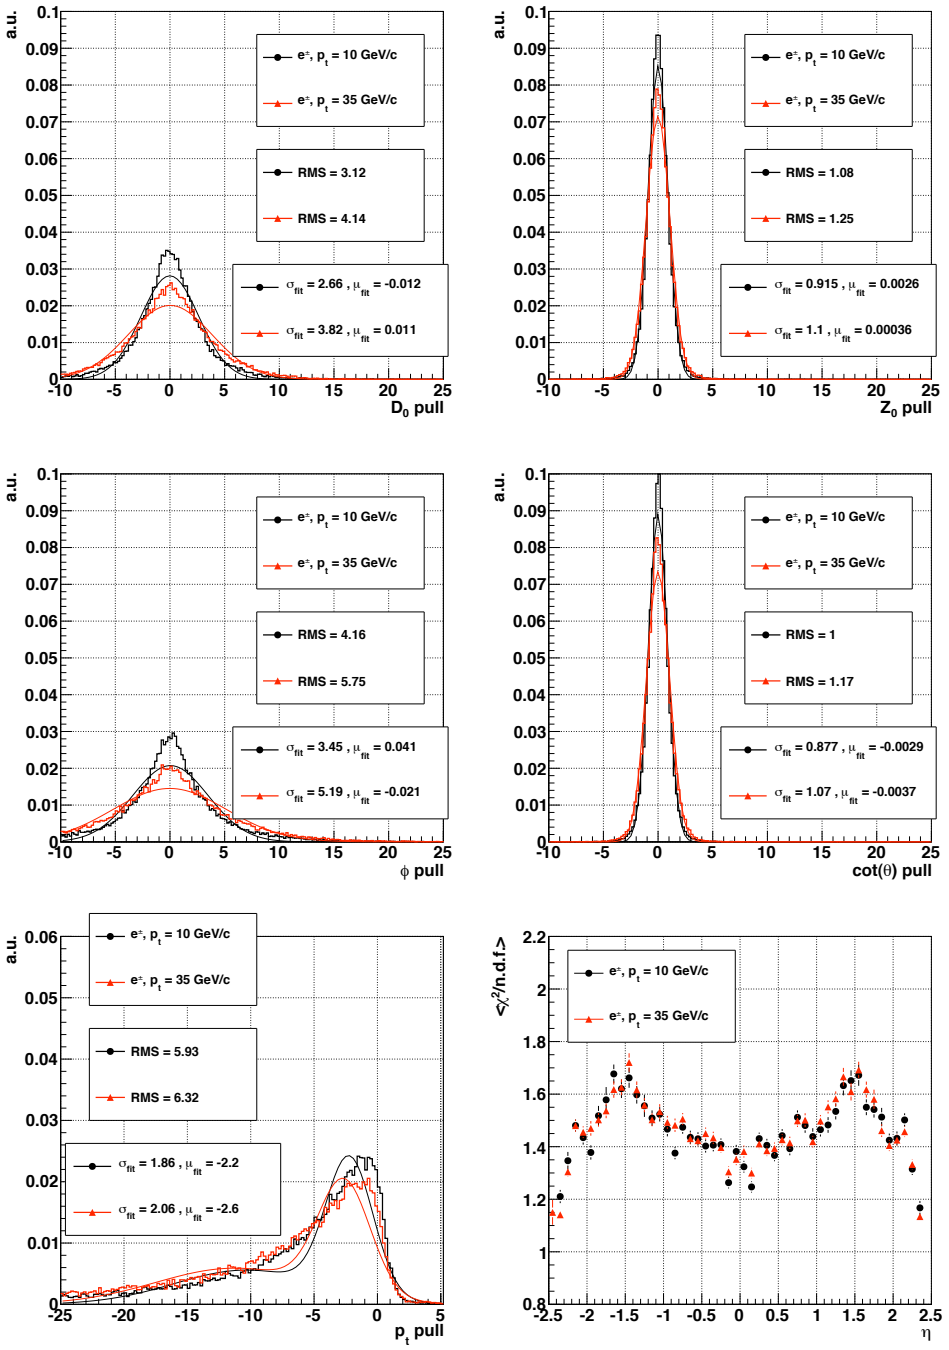

Figure 3.19: *The first five plots correspond to the pull distributions of the five track parameters for isolated electrons with transverse momenta of 10 and 35 GeV/c. From top to bottom and left to right: transverse momentum, φ, $\cot\vartheta$, transverse and longitudinal impact parameter. Both the σ and mean (of the narrower gaussian function) of a double-gaussian fit and the RMS of the pull distribution are printed. The sixth plot is the average normalized-χ^2 of reconstructed tracks as a function of pseudorapidity.*

3.5 Tracking performance on simulated LHC collisions

This section analyzes, using samples of simulated events, the performance of the CTF tracking software in reconstructing the trajectories of *not-isolated* charged particles that are produced in LHC collisions. Compared to the results shown in Sec. 3.4 for isolated particles, the tracking performance that are shown in this section are affected also by a feature that is specific to multi-track LHC collision events: the simultaneous presence, in the Tracker's detectors, of many measurements due to pile-up collisions, low energy spiraling particles (loopers) and hits left from the tens or hundreds of primary particles usually produced in those LHC events that are selected by the Trigger system. In the following some examples of the difficulties encountered by the tracking algorithm during the reconstruction of these events are described:

- Particles are produced within a very collimated jet and their position hits are closer than the typical uncertainty of the extrapolated trajectory position in the detectors (Fig. 3.6). In this case, the trajectory builder is unable to assign the measurements to the corresponding trajectories without ambiguity: e.g. hits corresponding to two distinct charged particles may be mixed into two reconstructed tracks that don't describe accurately any of the actual trajectories of the two particles.

- The trajectories of nearby particles are enough detached on the outer layers of the Tracker and they are correctly identified by the pattern recognition module. Nevertheless, their measurements on the innermost layers are so close to each other that the reconstruction algorithm does not assign correctly the hits to the corresponding trajectories. Distinct particles can be so near that their ionization signals are merged into a single cluster on the innermost silicon detector. In this scenario, even if the individual trajectories are reconstructed and their transverse momentum is generally well-measured, the resolution of the impact parameter is degraded.

- Due to the several hits produced by loopers and particles originating from pile-up collisions, it is possible that completely uncorrelated measurements are arranged as the legitimate trajectory of a charged particle and therefore are reconstructed as a ghost track.

Figure 3.20 shows, as a function of pseudorapidity, the global tracking efficiency of the CTF algorithm in reconstructing primary charged particles produced in LHC collisions. These simulated events contain di-jets with

transverse momenta within intervals of 15-20, 80-120, 120-170 and 3000-3500 GeV/c and multiple-jets produced from $t\bar{t}$ pairs. The efficiency for

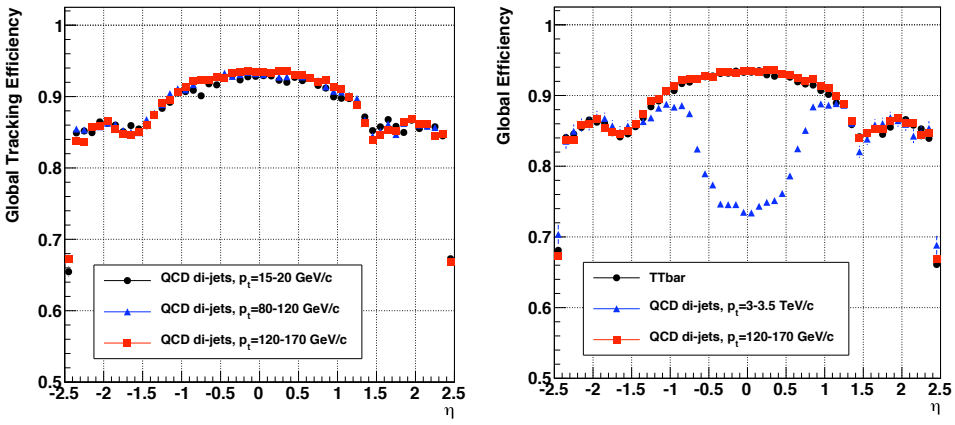

Figure 3.20: *Global tracking efficiency evaluated on simulated LHC collisions containing di-jets with transverse momenta within intervals of 15-20, 80-120 and 120-170 GeV/c (left); the plot on the right contains results also for di-jets of p_t within 3-3.5 TeV/c and multi-jets from $t\bar{t}$ pairs. The efficiency is estimated as a function of pseudorapidity. The events are produced in full simulation including the underlaying event, but without superimposing the pile-up collisions.*

the three samples with jets of lower transverse momenta (Fig. 3.20, left) is a sort of average of the reconstruction efficiencies for isolated pions which were shown in Fig. 3.11. This result is compatible with the fact that the majority of charged particles produced in LHC collisions consists of pions, followed by other hadrons whose trajectories are identified with similar efficacy. The higher Tracker occupancy inside the jets doesn't seem to compromise the pattern recognition in the range of energies of these three di-jets samples. However, the tracking efficiency is significantly degraded within jets of higher energy and that are more collimated: for events with di-jets of more than 3 TeV/c of transverse momentum, the global tracking efficiency decreases from approximately 94% to about 74% for tracks produced in the barrel[2] (Fig. 3.20, right).

The tracking efficiency as a function of the particle transverse momentum is approximately constant in the p_t range between 5 and 30 GeV/c and then starts decreasing (Fig. 3.21, left) because lower energy particles have a larger nuclear cross section. For transverse momenta below 0.8 GeV/c the efficiency

[2]Given the 7 TeV beams of the LHC, energy conservation and geometrical arguments requires that di-jets of 3.5 TeV/c transverse momentum can be produced only within $|\eta| < 1.3$.

decreases quickly from about 80% to zero (Fig. 3.21, right) due to several causes:

- The pions-proton nuclear cross section increases rapidly for pions of energies below 0.7 GeV/c.
- The track selection criteria are much tighter for low momentum trajectories since they are the main source of fake tracks.
- In order to match the constraints on track reconstruction's timing that are set by the available computing resources, the version of the CTF algorithm used in this study reconstructs only trajectory seeds from hit-triplets (pairs) that have a transverse momentum higher than 0.3 GeV/c (0.6 GeV/c).
- In evaluating the process noise, the trajectory parameters propagator assumes that the mass of the particle is that of a pion, i.e. the most common particle among those originating from LHC collisions. While this assumption is legitimate for relativistic particles, it breaks at low energies.

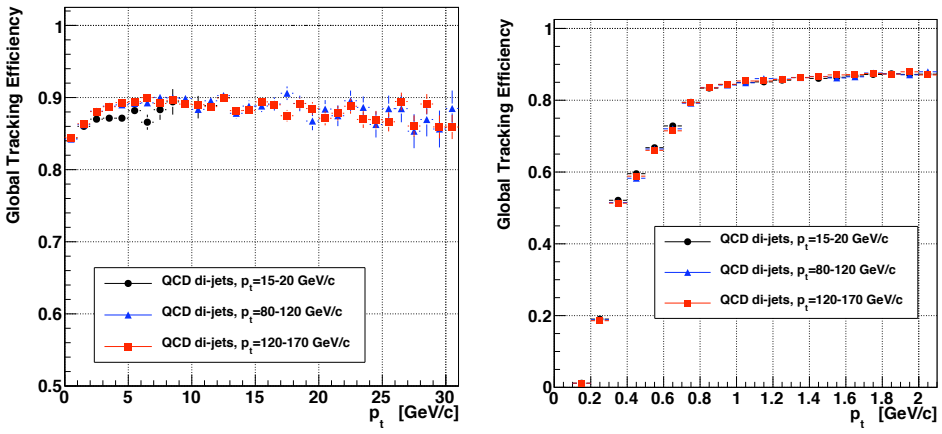

Figure 3.21: *Global tracking efficiency evaluated on simulated LHC collisions containing di-jets with transverse momenta within intervals of 15-20, 80-120 and 120-170 GeV/c; The efficiency is estimated as a function of the transverse momenta of the simulated charged particles (top). The events are produced in full simulation including the underlaying event, but without superimposing the pile-up collisions.*

For the analyzed multi-track samples, the fake rate (Fig. 3.22, top) has the same functional dependency on η as the fake rate evaluated on isolated pions samples (Sec. 3.4.1). The fake rate generally doesn't vary as a function

of the trajectory transverse momentum, except for values below 1 GeV/c: the smaller the energy of an initial trajectory seed, the larger is the search window that is used during the early iterations of the pattern recognition to collect additional trajectory measurements and higher is the probability to accept uncorrelated hits and to produce a ghost track. The hit-finding

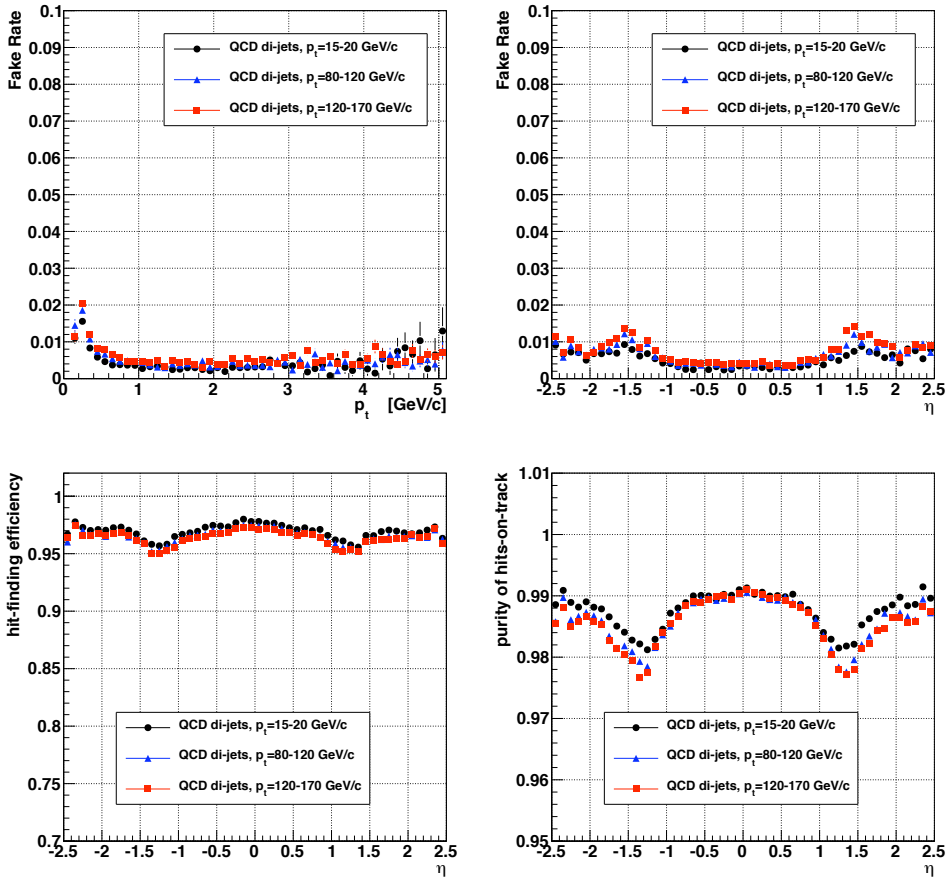

Figure 3.22: *Tracking fake rate as a function of the transverse momentum of the reconstructed trajectory (top-left) and as a function of its pseudorapidity (top-right). Hit-finding efficiency (bottom-left) and hit-purity (bottom-right) as a function of the pseudorapidity of the simulated particle that is reconstructed. All the four sets of histograms are evaluated on simulated LHC collisions containing di-jets with transverse momenta within intervals of 15-20, 80-120 and 120-170 GeV/c. The events are produced in full simulation including the underlaying event, but without superimposing the pile-up collisions.*

efficiency and hit-purity (Fig. 3.22, bottom) are similar to those measured for the isolates pion samples and worsen as the energy of the jets increases.

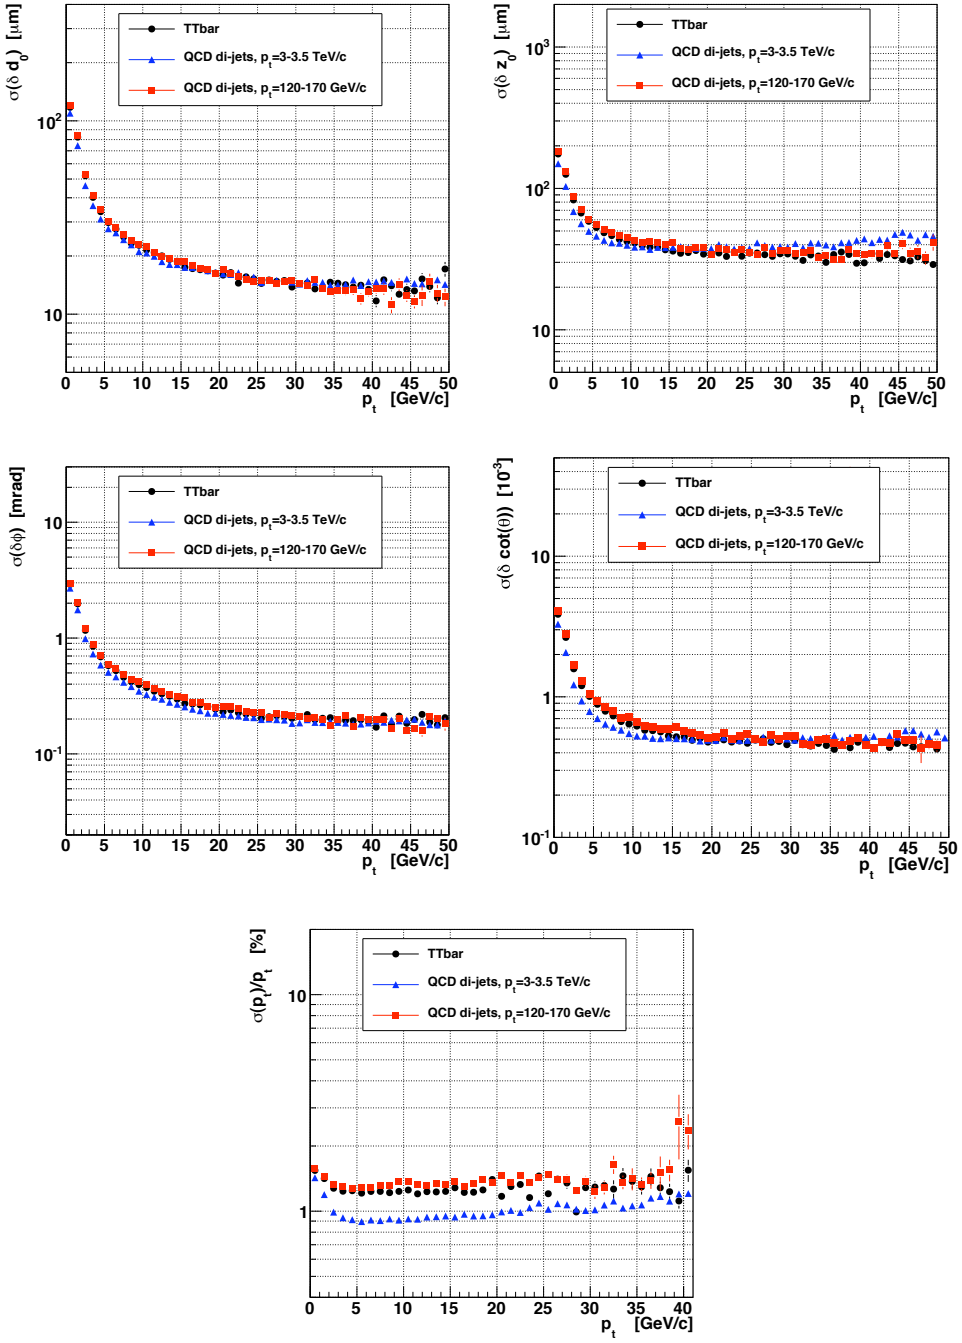

Figure 3.23: *Resolution, as a function of the transverse momentum, of the five track parameters for charged particles produced within 3 cm from the beam-line and ± 20 cm around the center of CMS in simulated LHC collisions containing di-jets with transverse momenta within 120-170 and 3000-3500 GeV/c and multiple-jets from $t\bar{t}$ pairs. From top to bottom and left to right: transverse momentum, φ, $\cot\vartheta$, transverse and longitudinal impact parameter. For each bin in η, the full symbol corresponds to the sigma of a gaussian fit of the residuals distribution. The events are produced in full simulation including the underlaying event, but without superimposing the pile-up collisions.*

The di-jets and $t\bar{t}$ events analyzed in this section contain charged particles that are spread over a wide range of energies and, therefore, these simulated samples allow the study of the relation between the resolution of the track parameters and the particle transverse momentum. In Figure 3.23 the resolutions of the five trajectory parameters are reported as a function of the particle transverse momentum. Due to the increasing influence of multiple scattering, all the five plots show a significant degradation of the resolution for values of transverse momentum below 10-20 GeV/c. For values *above* the same range, the resolutions of the four parameters measuring d_0, d_z, φ and $\cot\vartheta$ reach a plateau since at that energies the resolutions are dominated by the precision with which the trajectory position is measured on the innermost sensitive layer of the Tracker. The precision in measuring the transverse momentum is maximum for transverse momenta around 5 GeV/c (Fig. 3.23, bottom) and then degrades progressively for higher momenta, while the trajectories approximate closer and closer a straight line.

The transverse momentum resolution estimated on the 3-3.5 TeV/c di-jets sample is superior compared to the resolutions measured on the two other kind of events: the explanation is that the jets of the first sample, and the corresponding charged particles, are produced exclusively in the barrel region of the Tracker where the largest lever arm can be used to estimate the trajectory curvature.

Chapter 4

Improvements to the CTF algorithm

This chapter describes the major contributions by the author to improve each of the four modules that compose the CTF tracking algorithm (see Sec. 3.3). Section 4.1 describes the advantages of the current trajectory seed generator, which uses both pixel and strip measurements as input, in comparison to the original seed generator that was based exclusively on pixel hits. Also, a seed generator that doesn't employ pixel measurements at all is described and some possible applications are discussed. Section 4.2 is devoted to an extension of the trajectory builder that allows the reconstruction of trajectories also within those regions of the Tracker where nearby detectors overlap. The use of this new feature is discussed in the context of detector alignment and measurement of silicon sensor resolution. Finally section 4.3 shows how the use of a more sophisticated trajectory propagator, within the final track fitter module, removes a bias in the estimation of the reconstructed tracks' momenta due to the in-hogeneity of the magnetic field in the forward region of the Tracker.

4.1 Pixel-Strip and Pixel-Less trajectory seed generators

In the context of producing the trajectory seeds that initiate the track reconstruction sequence (see Sec. 3.3.2), the use of the pixel detector offers many advantages over the strip sub-system. The pixel sensors provide truly 2-dimensional measurements with similar resolutions on both position coordinates and allow a more precise estimate of the seed direction and impact parameter. Also, the pixel detectors are placed on the innermost sensitive layers of the tracking system and they enable the production of trajectory

seeds also for those charged particles that don't reach the outer layers of the Tracker due to inelastic nuclear interactions or hard bremsstrahlung radiation. However, specific barrel and forward strip detectors can counterbalance some weaknesses of a trajectory seed generator that is exclusively based on pixel measurements:

- Barrel pixel sensors don't overlap in z within the same ladder (Sec. A.0.1) and, because the ladders of the different layers are not staggered, the pixel barrel system is not sensitive to particles in correspondence of *all* the 2.2 mm gaps along z between pairs of adjacent sensors. In the context of tracking efficiency for primary particles, the most relevant gap is is the one at $\eta = 0$, since particles originating from LHC collisions are produced around z=0.

- Due to the geometrical acceptance of the forward pixel disks, particles produced with $|\eta|$ greater than approximately 2.1 often leave less than two hits (the minimum number of measurements that is necessary to produce a trajectory seed) on the pixel detector, even if these particles are detected by several layers of strip sensors and their trajectories could be reconstructed. This inefficiency is amplified for those collisions that are displaced along the z-axis from the center of CMS.

- Charged particles that are produced by the decay of long-living primary particles (e.g. K_s) or by inelastic interactions with the detector (e.g. $e^{\pm}$ pairs originated from the conversion of primary photons) can cross less than two layers of the pixel system when their production vertices are detached from the beam line by several centimeters.

The sketch in Fig. 4.1 shows how trajectories, which cannot be seeded using only the pixel detector, are reconstructed by means of trajectory seeds that use a combination of both pixel and strip hits (in the forward region of pseudorapidity) or solely strip measurements from the two innermost TIB layers (for trajectories with pseudorapidity close to zero). The same two layers of the Tracker Inner Barrel (see Sec. A.0.3) can be used to produce trajectory seeds also for charged particles that have a production vertex detached from the primary LHC collision up to about 25 cm.

In implementing a seed generator which is able to deal also with strip measurements, as much as possible of the original code developed for pixel seeds generation has been re-used. In the following the more relevant extensions that had to be implemented are described.

The mini-framework that finds the pairs and triplets of hits employed to define a seed has been extended in order to use the actual uncertainty on the position of the strip measurements. Indeed, in its first implementation, the

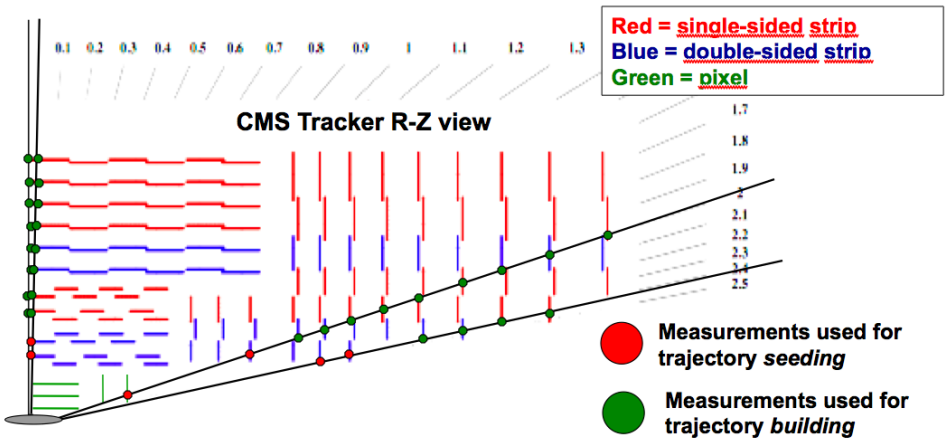

Figure 4.1: *Sketch showing how some charged particles produced at values of $|\eta|$ greater than 2.1 or close to zero cannot have a trajectory seed based exclusively on pixel measurements, but their trajectories are still reconstructed using the strip sub-system of the Tracker.*

pairs and triplets generator was using a rough parameterization of the pixel hit position uncertainty: this parameterization allows the seed generator module to be executed faster and it is a legitimate approximation for pixel measurements. However, the uncertainty on the less-precisely-measured coordinate of a *strip* hit changes by up to an order of magnitude depending whether the measurement is produced by a double-sided strip sensor or by a single-sided one. For measurements on strip layers, the seeding code now uses the uncertainty of each specific hit as it is evaluated by the local reconstruction module (see 3.3.1).

The original seeding code was also assuming that all the detectors of the seeding layers are rectangularly shaped. This premise is correct for the pixel layers, but it fails for the double-sided detectors of the barrel strip layers, which consist of pairs of tilted and coupled rectangular sensors, and for all the detectors of the forward strip layers, which are trapezoids or pairs of tilted and superimposed trapezoids (see Chapter 2). The code has been consequently generalized to cope with all these different topologies of sensors.

After testing the first implementation of the seed generator based on both pixel and strip measurements, it was apparent that many of the produced trajectory seeds were redundant, implying a waste of processing time. The original version of the pixel-based seed generator module used *all* the measurements of the specified seeding layers to produce the trajectory seeds. However, as it is shown in Fig. 4.2, some combinations of measurements on forward strip layers create trajectory seeds corresponding to particles that

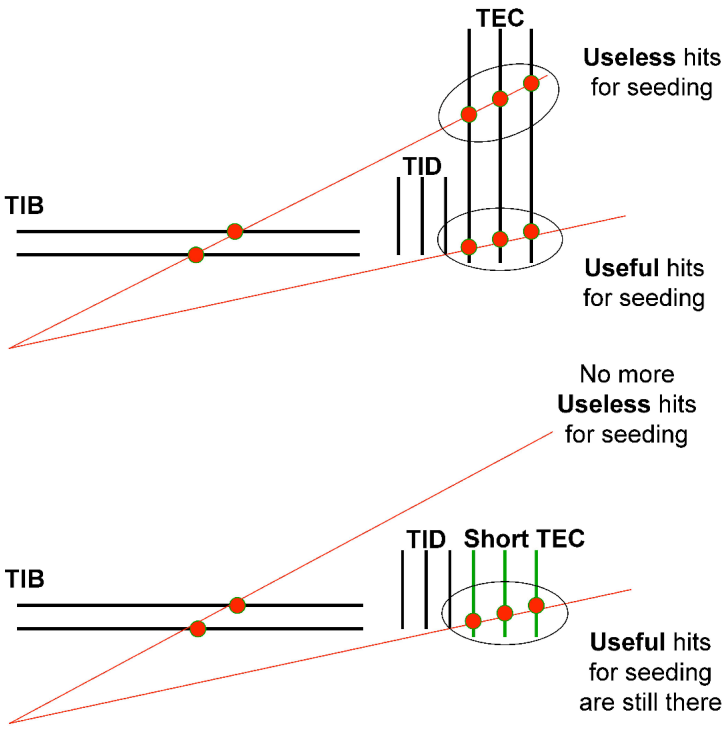

Figure 4.2: *Pictorical sketch of the silicon seed generation using standard DetLayers (up) and appositely implemented short DetLayers (bottom).*

have already been identified by other layers of the inner TIB and Pixel Barrel systems. Despite some level of redundancy is useful to maintain a high seed finding efficiency in case of detector failures, the use of the measurements on the outer rings of TEC layers is too time consuming compared to the benefits to seeding efficiency that it yields. The mechanism used by the seed generator to read the measurements of the seeding layers has been modified in order to offer higher flexibility: for each forward strip layer of TID and TEC, the range of disk rings that are used to create trajectory seeds can be specified in the configuration files of the seeding module; also, the seeder can be instructed to use only "matched" strip hits, single-component hits or a combination of both. In the current configuration of the CTF algorithm, only the first ring of each of the three innermost TEC disks (on both sides of the Tracker) is used to extend the geometrical reach of the reconstruction software up to pseudorapidity of about 2.5-2.6.

Figure 4.3 compares the tracking efficiency that is achievable employing a seed generator based exclusively on pixel measurements and the tracking efficiency of the current version of the software that utilizes a combination of both pixel and strip hits. Using the new seeder, the 2% inefficiency around $\eta = 0$ is completely recovered like most of the inefficiency in the range of

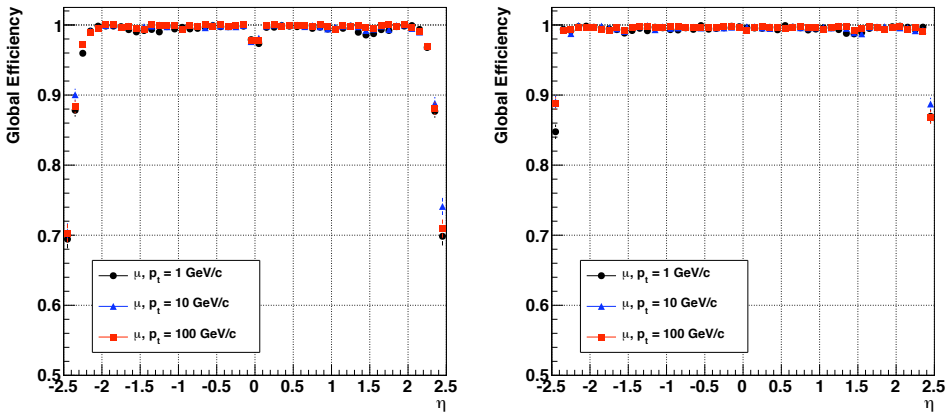

Figure 4.3: *Global tracking efficiency using a seed generator based exclusively on pixel measurements (left) and using the current configuration of the tracking software that employs both pixel and strip measurements in the construction of trajectory seeds (right). The results have been evaluated on simulated samples of isolated muons of transverse momenta of 1,10 and 100 GeV/c.*

$|\eta|$ between 2.1 and 2.5. As an example of the benefits provided to the physics reach of CMS by the increased tracking efficiency, Figure 4.6 shows the invariant mass of the four muons originating from the $H \rightarrow ZZ \rightarrow 4\mu$ decay chain using different configurations of the seed generator module. Given that in producing this plot a toy event selection based on simulated-reconstructed track association was used, the current seed generators allows the reconstruction of more than 1% additional Higgs events compared to same reconstruction sequence initiated by the pixel-only seed generator.

4.1.1 Track reconstruction without the pixel detector

After the seed generator software was extended to use also strip measurements, it was straightforward to define an additional configuration of the seeding layers that doesn't rely on pixel hits at all. Inserting such a configuration of the seed generation module at the beginning of the CTF reconstruction sequence, it is possible to reconstruct charged particles within CMS even in absence of any measurement provided by the pixel sub-system.

The main motivation to implement a "pixel-less" tracking algorithm was the fact that CMS would have not had the pixel detector installed during the LHC pilot run that was originally scheduled for the second half of 2007 [50]. This pilot run never took place due to the delays in the accelerator construction, however the initial pixel-less tracking configuration has been used

as a template to develop the tracking sequences used to reconstruct cosmic rays during different stages of the Tracker detector commissioning: the analysis of cosmic ray data with the Strip Tracker at the Tracker Integration Facility in 2007 [102, 103, 104] and subsequently the global runs of the full silicon Tracker together with the rest of the CMS detector during the Summer of 2008 (see Chap. 5). Even though the CMS detector has the pixel system installed, a track reconstruction sequence that doesn't rely on the pixel measurements is still valuable, since it allows an unbiased measurement from *data* of pixel detector performance (e.g. efficiency, resolution and alignment) using tracks reconstructed in the strip sub-system. Finally, a tracking sequence that doesn't need pixel measurements to reconstruct the charged particles coming from LHC collisions offers an important backup solution in case of hardware problems in the pixel system. In what follows, the performance of the pixel-less reconstruction sequence is analyzed using samples of isolated muons and $H \rightarrow ZZ \rightarrow 4\mu$ events produced in LHC collisions.

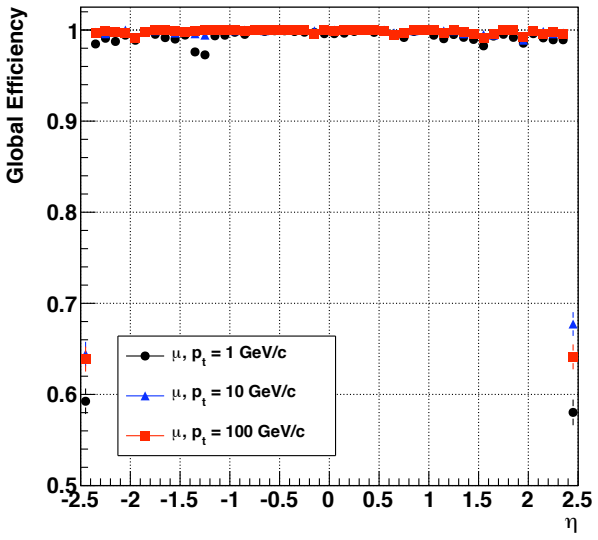

Figure 4.4: *Global tracking efficiency obtained using a seed generator based exclusively on strip measurements. The results have been evaluated on simulated samples of isolated muons of transverse momenta of 1,10 and 100 GeV/c.*

The global tracking efficiency for the configuration of the CTF algorithm that doesn't use pixel measurements within the seeding module is shown in Figure 4.4. The efficiency is higher than 99% in the full range of pseudorapidity, indicating that the pixel detector is not crucial for the hermeticity of the Tracker system.

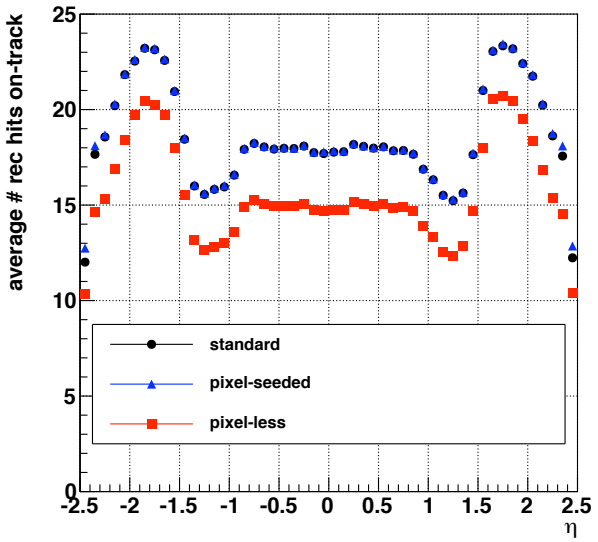

Figure 4.5: *Average number of position measurements included in the reconstructed trajectories as a function of pseudorapidity. The distributions are obtained using different configurations of the seeder module that initiate the track reconstruction sequence: standard, pixel-less and seeded only with pixel hits. The simulated samples were including isolated muons with a transverse momentum of 10 GeV/c.*

Figure 4.5 compares the average number of measurements that are included in the reconstructed trajectories that are obtained using the standard tracking sequence, the one seeded only by pixel measurements (but using the rest of the Tracker during the pattern recognition) or the pixel-less reconstruction sequence. In the last case, the trajectories have between 2 and 3 fewer measurements than in the other two tracking configurations. This is fully compatible with the geometry of the pixel system that has three layers in the barrel region of the Tracker and two disks in the endcaps. The difference in the number of hits in the forward region is not exactly equal to 2 since the forward pixel detector can provide more than two measurements for a trajectory that goes through the region of overlap between two adjacent sensors.

The resolution of the five track parameters as a function of pseudorapidity are shown in Figure 4.7. The track-fitting performance is compared with that obtained using the standard reconstruction sequence on the same three samples of isolated muons of different transverse momenta. Without the use of pixel measurements, the resolutions of the four parameters that depends heavily on the innermost position point added to the track (φ, $\cot\vartheta$, d_0 and

d_z) deteriorate by one order of magnitude or more. The main sources of the resolution decrease include: the much longer extrapolation distance between the innermost measurement of the trajectories and the impact point where the track parameters are evaluated; the multiple scattering induced by the pixel detector material between the innermost measurement and the impact point.

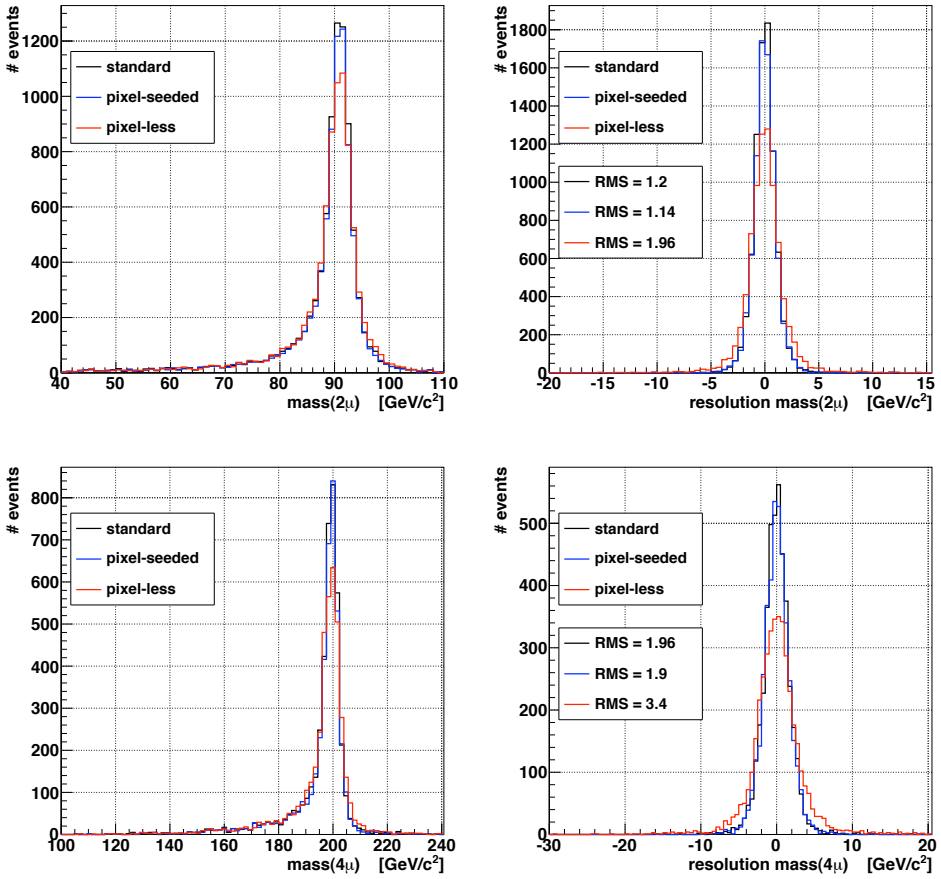

Figure 4.6: *Reconstructed mass (left) of the Higgs boson obtained from the 4-vectors of the four muons in the decay chain $H \rightarrow ZZ \rightarrow 4\mu$. The right plots reports the resolution of the mass measurement. Both histograms are obtained for three different configurations of the track reconstruction sequence: standard, pixel-less and seeded only with pixel hits.*

The resolution of the transverse momentum measurement worsen only slightly for the 1 and 10 GeV/c samples and it is a factor two poorer for the 100 GeV/c sample. The decline of the transverse momentum measurement is more significant in the last sample, since in that case the resolution is dominated by the length of the lever arm that is used to estimated the trajectory cur-

vature and the 6-7 extra centimeters provided by the pixel system are not negligible.

Due to the degradation of track parameter measurements for trajectories reconstructed without pixel hits, the resolution in evaluating the mass of resonances is consequently reduced. Figure 4.6 shows that the resolution in reconstructing the mass of 200 GeV/c^2 Higgs boson is spoiled by about a factor of two.

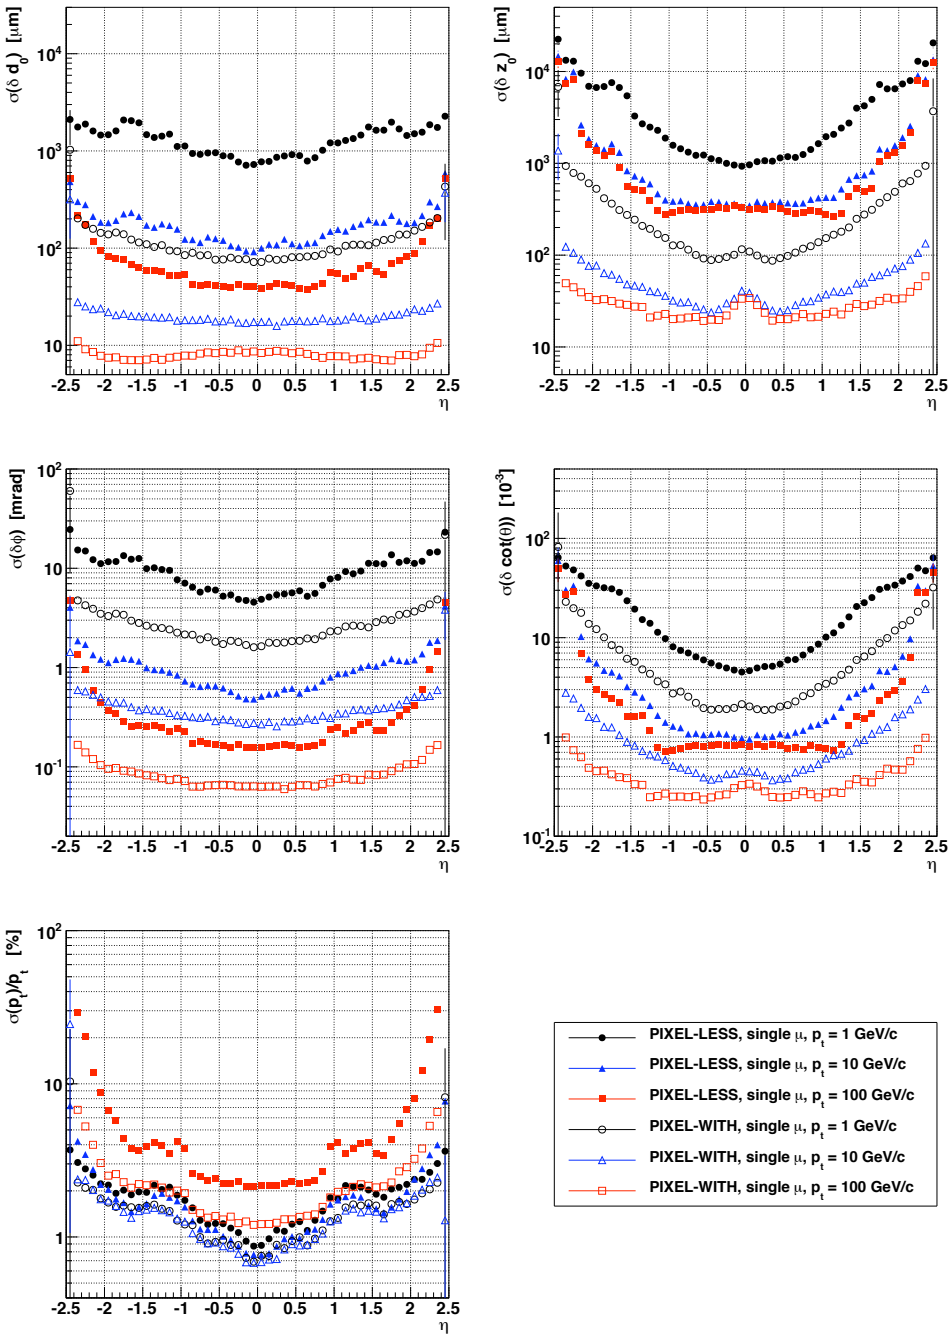

Figure 4.7: *Resolution, as a function of pseudorapidity, of the five track parameters for isolated muons with transverse momenta of 1, 10 and 100 GeV/c. From top to bottom and left to right: transverse momentum, φ, $\cot\vartheta$, transverse and longitudinal impact parameter. For each bin in η, resolutions correspond to the sigma of a gaussian fit of the residuals distribution. The full symbols refers to tracks reconstructed using a pixel-less tracking sequence, while empty symbols refers to tracks reconstructed with the standard configuration.*

4.2 Trajectory building within overlapping detectors

In its original implementation, the CTF module in charge of pattern recognition (see Sec. 3.3.3) was able to attach to the reconstructed trajectories no more than *one* position measurement per Tracker layer. While most of the times the charged particles produced in LHC collisions intersect only one sensor when they cross a layer of the Tracker system, sometimes their trajectories traverse those regions where nearby silicon modules overlap: in this situation, a charged particle can leave two or even more hits on the same layer of the tracking system.

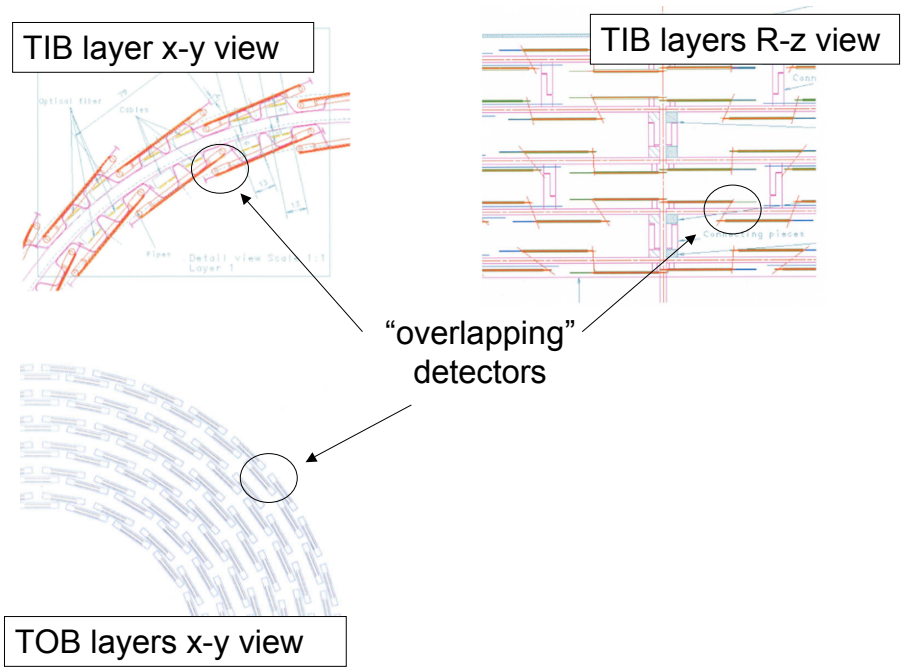

Figure 4.8: *Details of the partial overlap between nearby silicon modules located in different sub-systems of the Tracker: ϕ-overlap (top-left) between sensors on the same side of a TIB half-shell and z-overlap (top-right) between modules on opposite sides of the same half-shell; ϕ-overlap (bottom-left) between modules located on adjacent rods of a TOB layer.*

Figure 4.8 shows different types of overlaps for sensors within TIB and TOB layers. The fraction of a sensor's surface that is superimposed to that of an adjacent detector varies between 3-5% and depends on the Tracker subcomponent where the sensors are located (see App.A). The "effective" over-

lap varies also as a function of the trajectory crossing angle: while the geometry of the Tracker layers has been optimized to maximize the amount of sensors' overlap for trajectories that emerge radially from LHC collisions, the same system is less hermetic for cosmic rays that cross the Tracker vertically.

The reason for efficiently identifying *additional* hits produced by a charged particle on a given layer is not due to any expected improvements of the track-fitting performance: indeed, the measurements that are added to the track within the same detecting layer neither increase significantly the lever arm used to estimate the trajectories curvature nor reduce the extrapolation length between the innermost measurement and the impact point. However, the particle's hits located on adjacent overlapping sensors are valuable in several other contexts:

- **Alignment:** Due to the stochastic effects of the material and the extrapolation length of 3-15 cm between consecutive layers of the Tracker, the uncertainty of the *unbiased*[1] trajectory position over a sensor is greater than 100 μm (Fig.3.6). Tracker modules can be aligned to a comparable or higher level of precision only through a statistical approach that makes use of the residuals distribution of many tracks crossing the same sensor. However, if a trajectory includes both measurements of two overlapping sensors, the trajectory state on one detector can be used to evaluate the trajectory position on the other detector with an uncertainty which is very small due to an extrapolation length of only tens of millimeters between the sensors. Provided the same amount of collected data, these coupled measurements allow a higher level of inter-alignment among sensors *within* the same tracker layer.

- **Resolution measurement:** The same considerations about the uncertainty of the unbiased trajectory position discussed in the context of alignment apply also to the evaluation of sensor's resolution through the study of trajectory-measurement residuals.

- **Particle identification:** Recovering additional measurements produced by a given charged particle can be important for all those cases in which the *total number* of hits on the track makes a difference. An example is the particle identification that can be performed with the silicon Tracker measuring the dE/dx of the particles when they cross the silicon sensors [105, 106]. A larger number of hits correctly included in the trajectory is beneficial, since the dE/dx is estimated averaging the energy loss on all the crossed sensors.

[1]The position of the trajectory on a sensor, when it is evaluated without using the measurement of the sensor itself, is called unbiased

The new extended version of the trajectory builder follows the same iterative logic of the original builder, as it is described in Section 3.3.3, and re-uses most of the ancillary services like the trajectory parameters propagator and the modules implementing all the mathematical transformations requested by the Kalman filter formalism. However, the new trajectory builder is significantly different in the way it gathers the measurements compatible with a trajectory and includes them into the trajectory itself.

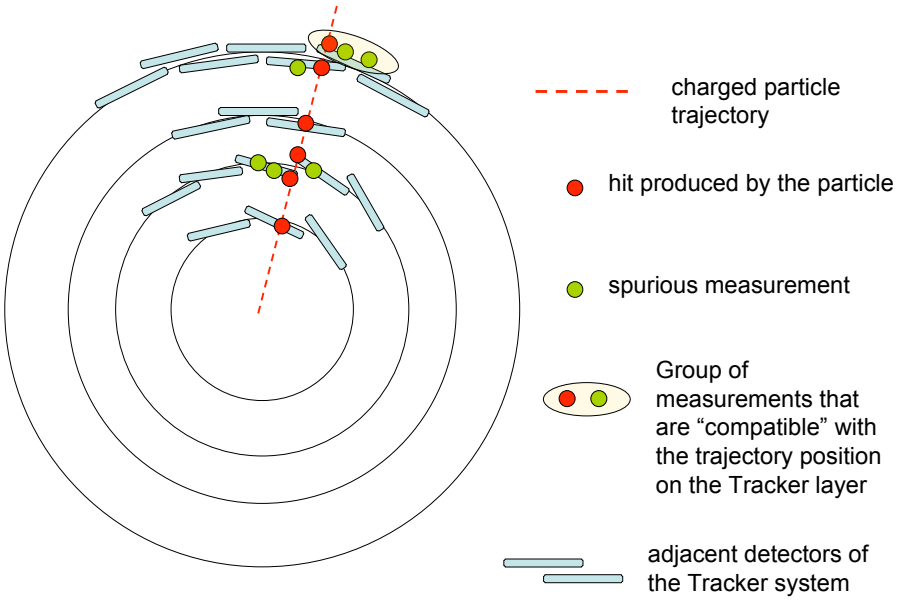

Figure 4.9: *Sketch of the transverse section of barrel layers crossed by a particle that intersects more than one sensor per layer. The picture represents both the hits actually left by the particle (red) and the "spurious" measurements that are compatible with the trajectory (green).*

Figures 4.9 and 4.10 show respectively a sketch of a particle crossing more than one sensor within the same Tracker layers and an outline of the way the measurements compatible with the trajectory are grouped by the two builders. After the trajectory parameters are propagated on a layer, the standard builder determines the sensors that are compatible with the trajectory and, within these sensors, all the compatible measurements. These measurements are provided as a *single group* to the software module in charge of implementing the Kalman filter (update) step. Despite the possibility to create new track-candidates for all (or most of the) compatible measurements (combinatorial Kalman filter, see Sec. 3.1.1), only *one* hit is included into a trajectory. The new builder, instead, organizes the measurements compatible to the trajectory in groups of "exclusive" hits. All

the measurements on the same silicon sensor form a group of exclusive hits since a particle can cross the sensor surface only once (omitting loopers) and an unique measurement is the one truly produced by the particle's ionization; conversely measurements on different sensors within the same layer are gathered in different groups of exclusive hits. The new trajectory builder constructs trajectory segments from pair or triplets of measurements belonging to two or three separate groups of measurements within the same layers. Subsequently, the Kalman filter updates the track-candidate adding trajectory segments instead of a single measurement. When only a group of exclusive measurements is found in a layer, only "segments consisting of single points" are constructed and, effectively, the same procedure of the original builder is implemented.

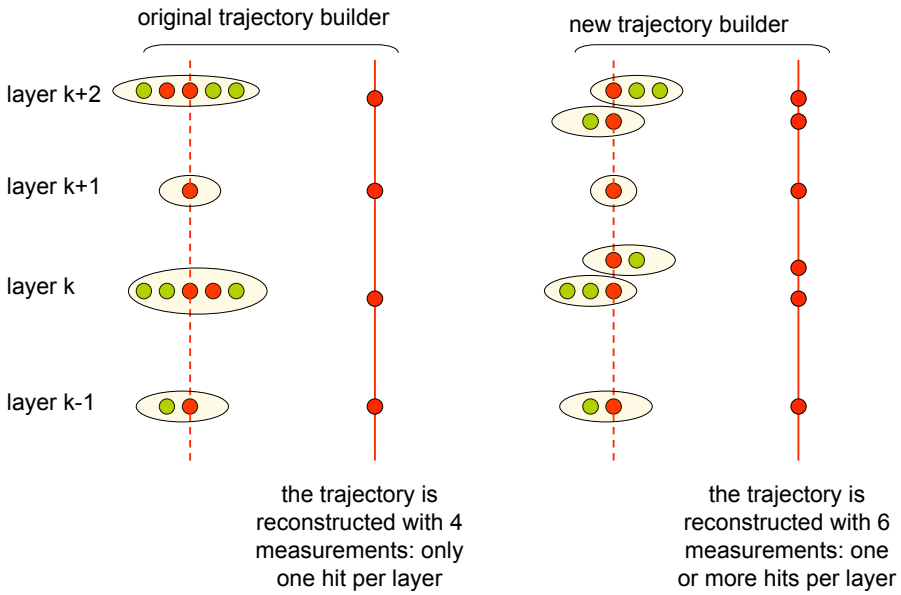

Figure 4.10: *Sketch representing the key difference between the original (left) and the new (right) implementation of the CTF trajectory builder. In the former case the measurements compatible with a trajectory are collected in a single group for each layer and only one measurement from this group is attached to the trajectory; in the latter, compatible measurements are arranged in two or more groups of "exclusive" hits for each layer and one measurement from each group can be attached to the reconstructed trajectory.*

According to the original design of the tracking software, the patter recognition starts from the layer of the outermost hits of the trajectory seed and proceeds looking for compatible measurements on the next layers. However, this approach doesn't allow the identification of hits on overlapping sensors that are located in the seeding layers. Hence, the new trajectory builder im-

plements also the option to re-build the trajectory outside-in within the inner region of the seeding layers. The rebuilding is done after all the trajectory hits have been identified in the outer Tracker layers during the inside-out building.

The plots in Figure 4.11 show the hit-finding efficiency and the average number of measurements on-track as a function of pseudorapidity evaluated on a sample of isolated muons for three versions of the trajectory builder: the standard one, the version recovering hits on overlapping sensors within the pattern recognition layers and the final version of the trajectory builder that is able to collect hits on overlapping sensors located both within the outer layers and within the seeding layers. The last configuration is the default for the current implementation of the CTF tracking sequence and it was used to produce the reconstruction performance plots analyzed in chapter 3. Depending on the pseudorapidity region, the hit-finding efficiency of the new builder is between 8 and 15% higher than the efficiency of the original builder. After the hits are recovered also in the seeding layers, the efficiency is above 99% in most of the η range of the Tracker and it dips to about 98% only in the barrel-endcap transition region. The average number of extra hits that are recovered is about 2 in the barrel region of the tracker and it is larger than 3 for tracks with a psedurapidity of about 1.8.

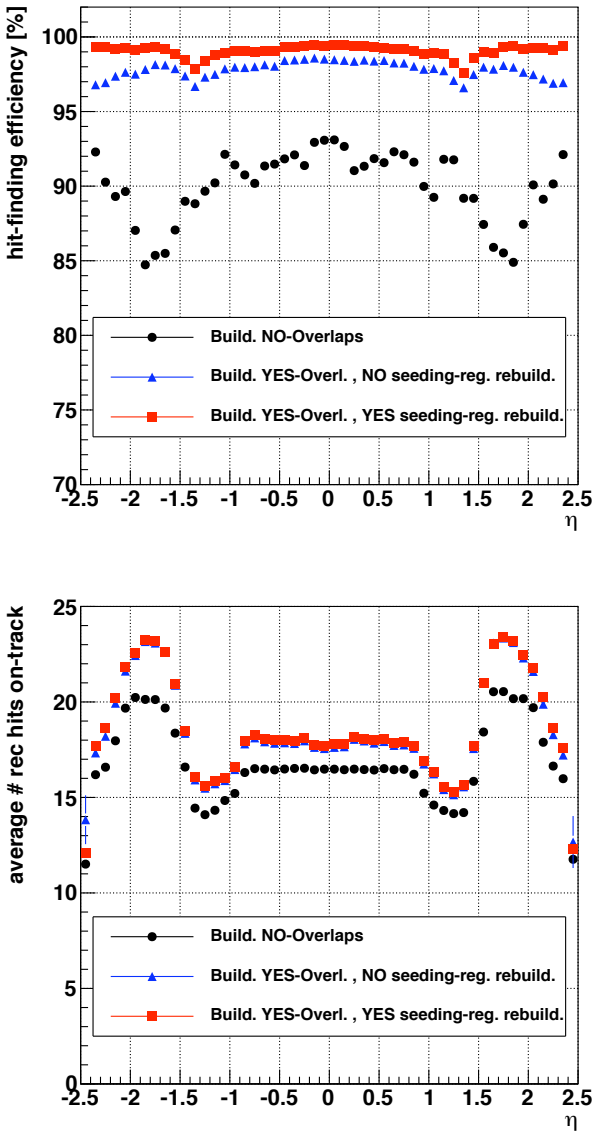

Figure 4.11: *Hit-finding efficiency (top) and average number of hits on-track as a function of η (bottom) evaluated on a sample of isolated muons using different versions of the CTF trajectory builder: original builder which is able to collect only one measurement per layer (black); trajectory builder that collects hits also within sensors' overlaps during patter recognition (blue); trajectory builder that collects hits also on overlapping sensors located within layers used for trajectory seed generation.*

4.3 Trajectory fitter and Runge-Kutta propagator

In Section 3.2 it was shown that an essential element of the Kalman filter estimator is the transport of the trajectory parameters, and of the associated covariance matrix, from one detector surface to another. In the presence of an uniform magnetic field, the Lorentz force constrains the charged particles to follow helical trajectories and therefore the parameters can usually be transported using analytical formulas. In particular, if the magnetic field is of the form $\mathbf{B} = (0, 0, B_z)$, the projection of a trajectory on the transverse x-y plane is a circle and the longitudinal component of the particle momentum is unchanged by the Lorentz force. The intersections between the trajectory and the detector can be calculated exactly as long as the latter consists of cylindric layers (planes) that are parallel (perpendicular) to the z axis.

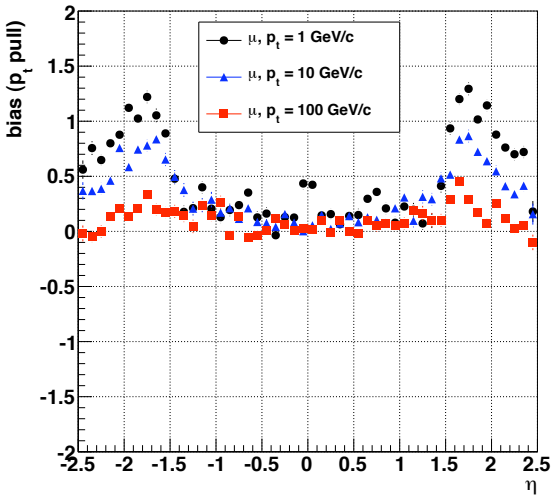

Figure 4.12: *Pull distributions of the reconstructed transverse momentum are evaluated separately for different values of pseudorapidity and using three samples of isolated muons. For each bin in η, the histogram shows the mean of the corresponding pull distribution. All the trajectories have been reconstructed using an analytical track propagator.*

In CMS, the magnetic field is in first approximation uniform and parallel to z within the Tracker volume (see Fig. 1.6). Therefore the implementation of the original trajectory propagator assumes that, at least in the region between two consecutive Tracker layers, the charged particles move exactly along helical trajectories and that the axes of the helices are parallel to z. All the detector modules, except for pixel forward modules, are arranged on cylindric layers or planes, respectively parallel and perpendicular to z (App.A). As a consequence, in the hypothesis of helical trajectories,

the propagator can calculate the intersections between trajectories and most Tracker components using *analytical* formula. However, the actual magnetic field is not strictly homogeneous and the assumption of exact helical trajectories introduces a systematic shift in the position of the propagated trajectory states. When these imprecise states are used by the Kalman filter (see Eq. 3.17 and 3.20), the estimates of the transverse momenta of the reconstructed tracks have a bias up to 2% (Fig.4.12).

The systematic error due to the $\mathbf{B}(\mathbf{x}) = (0, 0, B_z)$ assumption can be removed using a trajectory propagator that solves the equations of motion also for an inhomogeneous magnetic field. Iterative *numerical* methods can provide approximate solutions that match basically any level of precision employing a large enough number of iterations. This section describes the implementation of a new trajectory propagator based on an adaptive version of the Runge-Kutta method [107]. The performance of the new propagator is compared with that of the analytical propagator analyzing the bias in the transverse momenta of the reconstructed charged particles. Also the bias of the reconstructed masses of the Z and the Higgs bosons is examined.

4.3.1 Runge-Kutta algorithm

In numerical analysis, the Runge-Kutta methods are a family of iterative procedures used to find approximate solutions of ordinary differential equations. The fourth-order Runge-Kutta method (RK4) is the member of the family that is most commonly used since it usually provides the best ratio between the accuracy of the solution and the computation time. Given an unknown function $y(t)$ of the single parameter t, a differential equation and the initial condition are expressed as:

$$y\prime = f(t, y), \qquad y(t_0) = y_0 \tag{4.1}$$

The RK4 method provides an approximate solution of $y(t)$ through the iterative application of the following equations:

$$y_{n+1} = y_n + h \cdot (k_1 + 2k_2 + 2k_3 + k_4)/6 \tag{4.2}$$
$$t_{n+1} = t_n + h \tag{4.3}$$

where y_{n+1} is the approximation of $y(t_{n+1})$ and

$$k_1 = f(t_n, y_n) \tag{4.4}$$
$$k_2 = f(t_n + h/2, y_n + k_1 \cdot h/2) \tag{4.5}$$
$$k_3 = f(t_n + h/2, y_n + k_2 \cdot h/2) \tag{4.6}$$
$$k_4 = f(t_n + h, y_n + k_3 \cdot h) \tag{4.7}$$

The RK4 method is called a fourth-order method because the discrepancy between the exact solution $y(t_n)$ and the approximation y_n is of the order of h^4. Given the same step size h, higher-order Runge-Kutta methods can provide solutions that are more precise than that of RK4. However, these methods involve the calculation of a higher number of terms k_i and therefore take longer to execute.

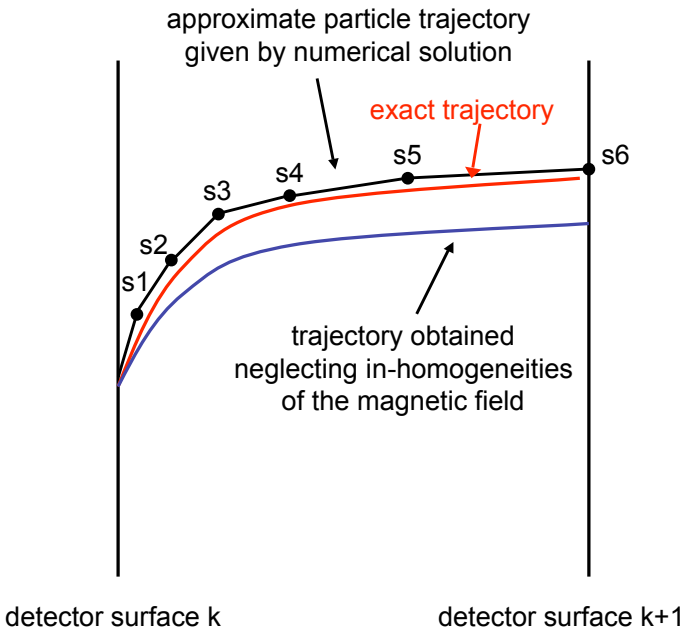

Figure 4.13: *Sketch showing the evolution of one trajectory parameter for a charged particle moving inside an inhomogeneous magnetic field between two consecutive detectors. An approximate solution of the equations of motion can be obtained using iterative numerical procedures. The step of each iteration can be tuned according to the expected level of "fluctuation" of the exact solution.*

Since the precision of the numerical solution increases for a smaller step size h, the required level of accuracy can be achieved with a sufficiently large number of iterations. However, the efficacy of the numerical approximation depends also on how much the actual function $y(n)$ "fluctuates" in a given interval of t. For a function that is "flat" enough, the accuracy of the numerical solution can be adequate also using relatively large steps. The sketch in Figure 4.13 shows the evolution of one of the five trajectory parameters for a charged particle traversing the volume between two consecutive detector layers. A large number of iterations is used to determine the numerical solution in the region where the trajectory parameter varies more, while larger

steps are used in the region where the exact function is flatter.

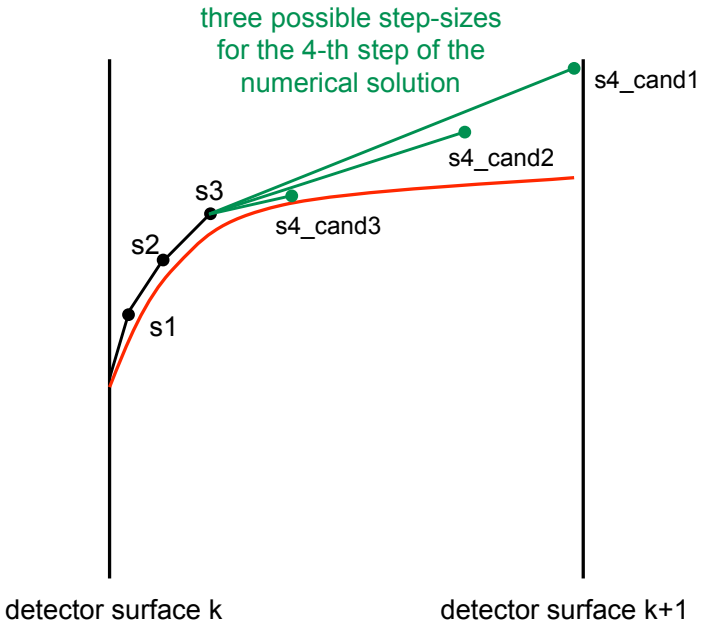

Figure 4.14: *Sketch showing the "adaptive" functioning of the Runge-Kutta method implemented for track propagation. For each step of the iterative procedure, the algorithm aims to determine the longest step size that assures an approximate solution within the requested level of precision.*

The Runge-Kutta method that has been implemented for the transport of the trajectory parameters is not the RK4, but it is the fourth-order *adaptive* method of Cash and Karp [108]. This is a method of the Runge-Kutta family that uses the evaluation of six k_i terms to calculate simultaneously fourth- and fifth-order accurate solutions. For each step of the numerical integration, the difference between these solutions is used as an estimate of the error of the fourth-order solution. The error is then employed by the algorithm to determine the longest step size that produces a fourth-order solution within the requested level of precision. The sketch in Figure 4.14 shows how, after the first three iterations, the adaptive algorithm makes three attempts before it determines the proper step size for the fourth iteration. Despite the additional nested iteration, the adaptive method has two important advantages compared to fixed step size methods:

- The adaptive algorithm converges quickly in those intervals of the parameter t where the function $y(t)$ is "flatter".
- The adaptive algorithm usually guarantees the requested level of accuracy even if the function $y(t)$ "fluctuates" significantly.

The adaptive method used for track reconstruction has been configured to calculate the position of the extrapolated trajectories with an accuracy better than 0.5 μm.

The equations of motion have to be rewritten properly in order to use the Runge-Kutta procedure for their integration. A charged particle that moves in a magnetic field is subjected to the Lorentz force:

$$\frac{d\mathbf{p}}{dt} = q \cdot \mathbf{v} \times \mathbf{B} \tag{4.8}$$

The last equation is equivalent to the following system of differential equations:

$$\frac{d^2\mathbf{x}(s)}{ds^2} = \frac{q}{p} \cdot \frac{d\mathbf{x}(s)}{ds} \times \mathbf{B} \tag{4.9}$$

where s is the distance covered by the particle along its trajectory and the other symbols have a clear meaning. Such a system of three *second*-order differential equations can be expressed as a system of six *first*-order differential equations:

$$\frac{d\mathbf{x}}{ds} = \hat{\mathbf{p}} \tag{4.10}$$

$$\frac{\mathbf{p}}{ds} = q \cdot \hat{\mathbf{p}} \times \mathbf{B}(s) \tag{4.11}$$

To resemble Eq.4.1, let us express the last two equations in the compact form $\mathbf{A}'(s) = \mathbf{f}(s, \mathbf{A})$ where $\mathbf{A}(s) = (\mathbf{x}(s), \mathbf{p}(s))$. Then the six functions $\mathbf{A}$ can be resolved running six numerical integrations in parallel and using six RK4 relations (Eq. 4.2):

$$A_{i,n+1} = A_{i,n} + h \cdot (k_{i,1} + 2k_{i,2} + 2k_{i,3} + k_{i,4})/6, \quad i = 1 \ldots 6 \tag{4.12}$$

or the analogous formulas of the adaptive Cash-Karp method. In analogy with equations 4.4-4.7, each term $k_{i,j}$ is calculated using $f_i(s, \mathbf{A})$.

4.3.2 Performance

The benefit given by the new track propagator is shown first analyzing the bias[2] of the pull distribution of the reconstructed transverse momenta. This observable has been plotted as a function of particles pseudorapidity for three samples of isolated muons of simulated p_t equal to 1,10 and 100 GeV.

[2]In this context, the bias is the deviation of the distribution mean from the expected theoretical value. For pull distributions, the expected mean is zero.

These samples have been reconstructed using the original analytical propagator and using the new one. The usage of the analytical propagator produces a bias up to 1 sigma for charged particles traversing the end-caps of the Tracker (Fig. 4.12). In these regions, compared to the barrel, the magnetic field is more inhomogeneous and the assumption of an uniform field is less appropriate. Using the same configuration of the tracking software, except for the new track propagator, the three samples have been reprocessed. The bias of the analogous pull distributions is significantly reduced. It is smaller than 0.2 sigma in the full η range for the 10 and 100 GeV samples and it is reduced by approximately a factor of two for the 1 GeV sample (Fig. 4.15). For this last sample of low energy muons, the distribution of the transverse momentum bias has still some peaks of about 0.5 sigma. These deviations from zero are currently explained as the influence of some approximations used in modeling the energy loss during the track fit. This issue is not analyzed here, since it is outside the scope of this thesis.

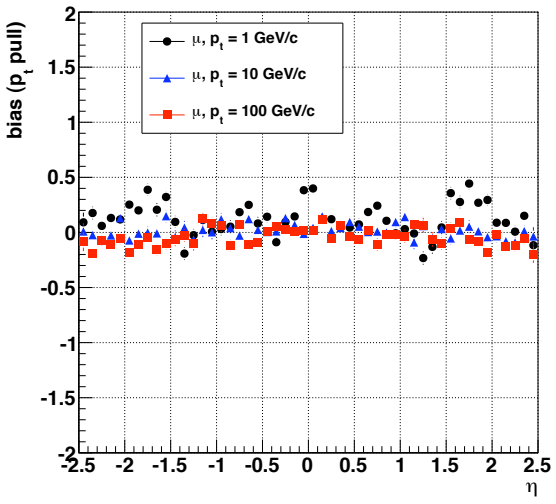

Figure 4.15: *Pull distributions of the reconstructed transverse momentum are evaluated separately for different values of pseudorapidity and using three samples of isolated muons. For each bin in η, the histogram shows the mean of the corresponding pull distribution. All the trajectories have been reconstructed using the new track propagator based on an adaptive Runge-Kutta method.*

Also a sample of simulated LHC collisions have been reconstructed twice using both propagators. Each event contains a Higgs boson that has a mass of 200 GeV/c^2 and decays into four muons through the $H \rightarrow ZZ \rightarrow 4\mu$ decay-chain. The masses of both the Higgs bosons and the pairs of Z bosons have been reconstructed summing up the 4-vectors of the four reconstructed

muons. These values have been compared to the invariant masses attained in the same way from the 4-vectors of the simulated muons. The mass resolution plots (Fig. 4.16) are obtained subtracting the invariant mass of the system of simulated muons from the invariant mass of the reconstructed system. The reconstructed mass values are not compared directly to the masses of the simulated Z and Higgs bosons in order to decouple the effects on the mass resolution due to track reconstruction from those due to radiation of internal bremsstrahlung photons[3].

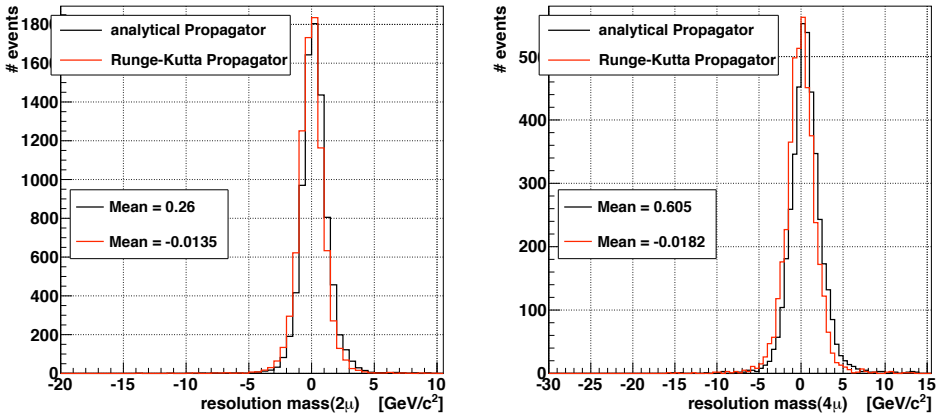

Figure 4.16: *Resolution plots of the reconstructed mass of Z (left) and Higgs (right) bosons. The simulated Higgs bosons are produced with a mass of 200 GeV/c^2 from LHC collisions and they are forced to decay in four final muons through the $H \to ZZ \to 4\mu$ decay-chain. The mass resolution is here defined as the difference between the invariant masses of the system of four reconstructed muons and the invariant mass of the corresponding system of simulated particles.*

A resolution of about 2.5 GeV/c^2 is achieved for the masses of both resonances, independently of the track propagator that was used in the reconstruction. However, the bias of the two resolution plots is significantly smaller when the new track propagator is used. Starting from a bias of about 300 and 600 MeV respectively for the Z and the Higgs boson, the new propagator reduces the bias of the reconstructed mass to approximately 10 MeV for both samples (Fig. 4.16).

[3]The tail in the mass distribution due to the internal radiation is visible in the left plots of Fig. 4.6.

Chapter 5

Reconstruction and analysis of cosmic muons

For the four weeks between the months of October and November 2008, the "Cosmic global Run At Four Tesla" (CRAFT) involved all CMS components in their final configuration. The excess of 400 million triggered cosmic rays was used both to start the calibration of the tracking and calorimeter systems and to test the reliability of the trigger and data acquisition chain. While other global runs were performed between May and August 2008 with the CMS solenoid turned off, CRAFT was the first long global run with the magnet operating at 3.8 Tesla. Hence, the CRAFT data provided the first statistically significant sample of reconstructed tracks whose momenta was measured.

Among all the CRAFT events saved on disk, about 15 million consisted of cosmic muons traversing the Tracker volume. These particles offered a valuable opportunity to understand the silicon tracking system and study its performance before the first LHC collisions. The CRAFT data were used for many types of studies, including:

- identification and correction of errors in the cabling of silicon modules.
- synchronization among read-out electronics, digitizer modules and trigger clock.
- measurement of detector modules efficiency and resolution from data.
- calibration of detector response.
- internal alignment of the tracking system and inter-alignment of the Tracker with the other sub-detectors of CMS.

From the perspective of the software for track reconstruction, the cosmic global run offered the opportunity to test on real data the efficacy of the tracking algorithm and to compare the performance with those obtained in simulations.

This chapter is structured as follows. Section 5.1 describes the modifications to the CTF algorithm that were implemented to reconstruct cosmic muons, as opposed to charged particles from LHC collisions. Section 5.2 describes a method to determine the track reconstruction efficiency from data and provides a measurement of it. Finally, Section 5.3 describes the method and provides a measurement of the resolution on transverse momenta of the reconstructed tracks. For both the efficiency and the resolution measurements, the results are compared with performance expected from simulations.

5.1 CTF algorithm for the reconstruction of cosmic muons

The different steps of the standard track reconstruction sequence were described and analyzed in Chapter 3. In particular, the reason to initiate the trajectory building from the inner pixel layers was motivated in terms of detector occupancy and tracking efficiency (Sec. 3.3.2). While these two arguments are valid for particles originating from LHC collisions, the topology of typical CRAFT events sets different constraints to the tracking software used for cosmics:

- Except for multi-track events due to the showering of cosmic particles inside the calorimeter or Tracker volume, the majority of CRAFT data consists of single particles producing less than about 30 hits in the whole Tracker volume. Assuming that the number of pixel and strip clusters due to noise is under control, the occupancy of the Tracker is not an issue independently of what layers are chosen for trajectory seeding.

- Cosmic particles that reach the Tracker volume have, in first approximation, uniform longitudinal and transverse impact parameter distributions. Therefore, the probability that a specific detector layer is crossed by a cosmic is proportional to the area of the layer "seen" by the incoming particle (Fig. 5.1). Of the 15 million CRAFT events containing a muon trajectory pointing to the Tracker volume, only 800 thousands have a cosmic particle traversing the volume of the pixel detector. For this reason, an order of magnitude higher efficiency is achieved if the track reconstruction of cosmic particles is initiated by the measurements of the outer (and larger) layers, as opposed to the pixel ones.

- Particles produced in LHC collisions, originate from the center of CMS and move outwards. Except for very low p_t tracks, each trajectory intersects a layer of the Tracker only once. Unlike particles from collisions, cosmics travel upside-down through the tracking system volume and, depending on their direction in η, they intersect twice the barrel layers (Fig. 5.1).

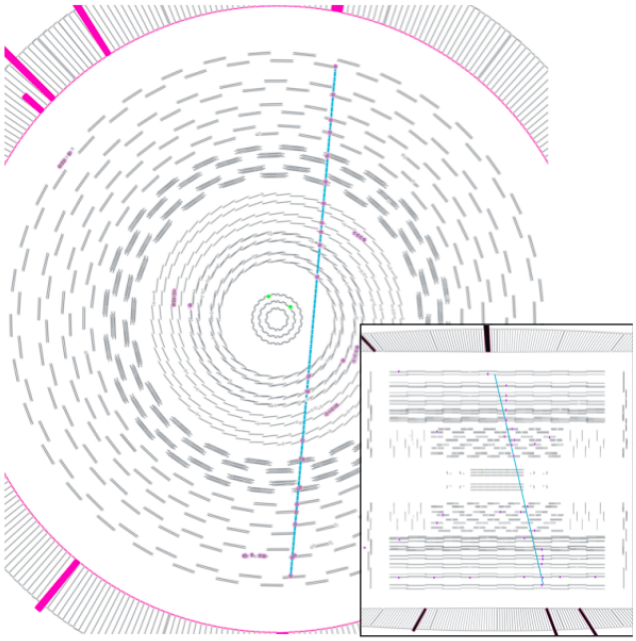

Figure 5.1: *Screenshot showing the reconstructed trajectory of a cosmic particle in the transverse (background) and longitudinal (small quadrant) view of the detector. The track is produced using exclusively measurements on strip modules, since the cosmic particle didnt intersect the small volume of the pixel detector. Like in most of the CRAFT events, the barrel layers are crossed twice.*

The CTF algorithm has been extended in order to use it in the reconstruction of cosmic particles. The major changes affected the trajectory seeding stage of the reconstruction sequence. A new seed generator module that makes use of measurements on the strip outer layers has been implemented. Also, the layer navigation service used by the trajectory builder (Sec. 3.3.3) has been generalized. Now it includes also a special link between a layer and itself to allow the complete reconstruction of particles traversing opposite hemispheres of the Tracker. Figure 5.1 shows a typical track reconstructed from CRAFT data, which intersects twice the strip barrel layers.

A special configuration of the tracking software, which uses the new seeding module and activates the more general layer navigation, was employed

throughout the processing of CRAFT data. Despite differences with the software for LHC collisions, the software for the reconstruction of cosmic trajectories utilizes exactly the same track-fitting algorithm and, except the difference in the layer navigation, the same implementation of the track-finding procedure. It has been evaluated that approximately 90% of the track reconstruction software that will be used for processing the LHC data has already been successfully employed in CRAFT.

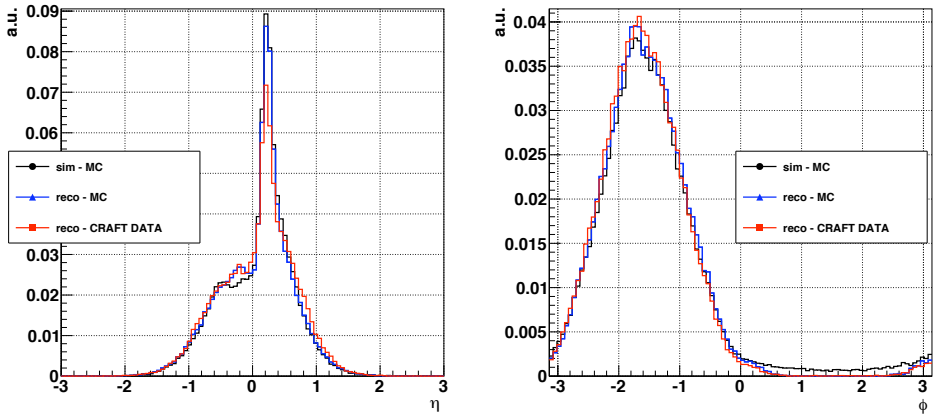

Figure 5.2: *Pseudorapidity (left) and phi (right) distributions for cosmic tracks traversing the Tracker. The comparison is among MC particles, tracks reconstructed in simulated samples and tracks reconstructed in CRAFT data. All histograms are normalized to unit area.*

Figure 5.2 shows the distributions of η and ϕ directions for tracks reconstructed in CRAFT and reconstructed in simulated events. As expected, both in simulation and in data, most of the particles move vertically from the earth surface toward CMS: the ϕ direction peaks at -$\pi/2$. The majority of the cosmic particles that reach CMS move through the shaft which was used to lower the detector components into the underground cavern. Except for very energetic cosmic particles, the approximate 100 meters of ground between the surface and the cavern operate as a shield. The asymmetry in the η distribution is a consequence of the fact that the detector is not placed exactly below the shaft, but is shifted in z. Apart from this effect, few of the reconstructed (and reconstruct-able) trajectories have pseudorapidity outside the [-1,1] range. It is again apparent the qualitative difference between the particles traced in CRAFT and the ones that will be reconstructed from LHC collisions: uniformly distributed in ϕ and extending beyond 1.5 in $|\eta|$.

5.2 Measurement of the tracking efficiency

The performance of both the Tracker software and the detector has been verified through the measurement, from CRAFT data, of the track reconstruction efficiency in the barrel region. The method used proceeds in two steps. First, events in which good quality *stand-alone* muons point into the Tracker volume are selected. Secondly, the subset of these events that contain a *Tracker* track matching the initial muon is determined. The ratio between the number of events in the two classes provides the estimate of the Tracker reconstruction efficiency. In the following, the first set of events is indicated as "denominator" class and the second as "numerator" class of events.

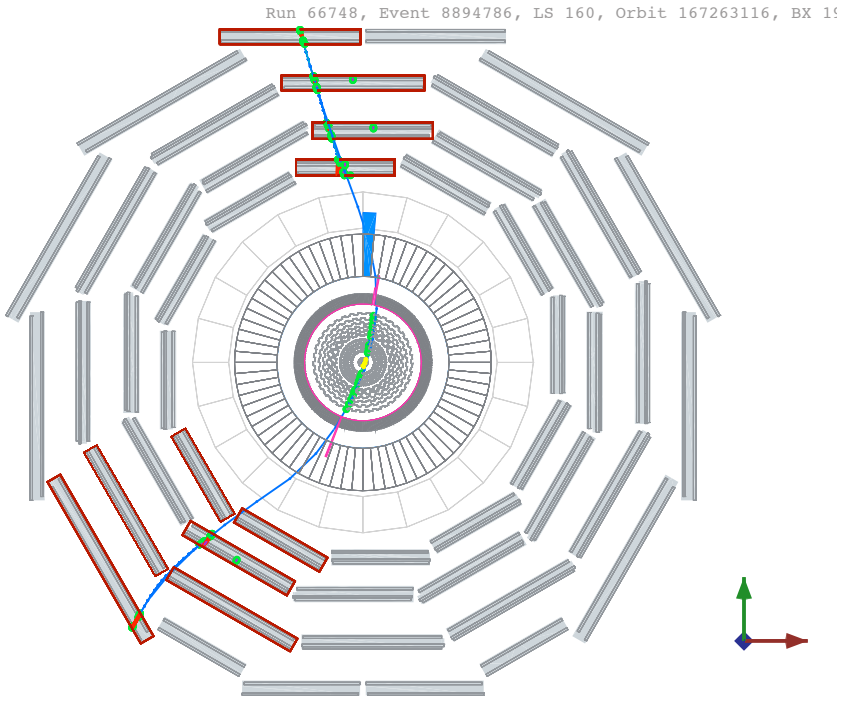

Figure 5.3: *Event display of a cosmic muon reconstructed both in the Muon tracking system and in the inner Tracker.*

The analysis is simplified restricting the efficiency measurement only to the events of a good quality run that were taken in stable conditions. Run 69912 has been chosen since the magnet operated constantly at 3.8 T and all the Tracker sub-detectors joined the data taking without any temporary problem of calibration, time synchronization or configuration. This run consists of more than 2 millions triggered events, of which about 200 thousands contain muon trajectories pointing to the Tracker volume. To avoid any bias in the

measurement of the Tracker efficiency, the data-set has been preliminary skimmed. Only those events corresponding to a special HLT trigger path, exclusively based on the information provided by the Muon stations, are considered.

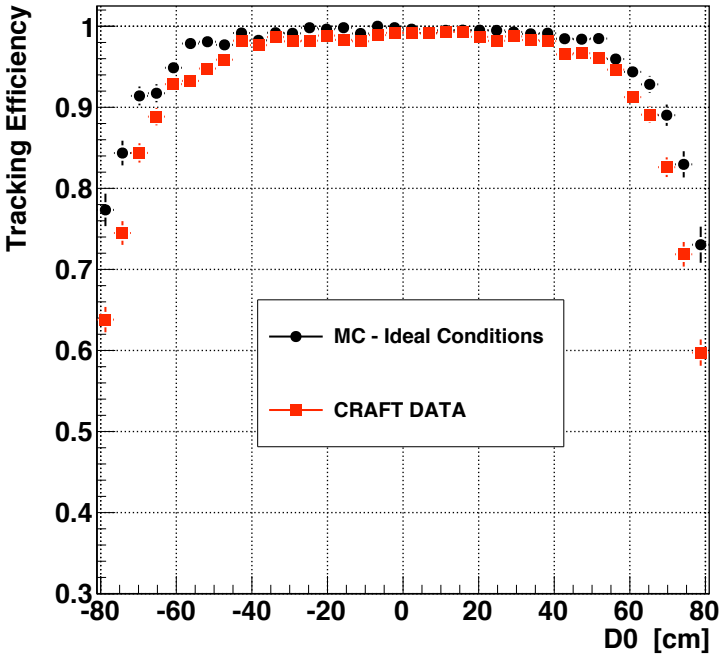

Figure 5.4: *Tracking efficiency measured from CRAFT data and compared to MC.*

The events of the denominator class are required to have one stand-alone muon, whose trajectory intersects the surface of the sixth layer of TOB both in the positive- and in the negative-y hemisphere of CMS. The crossing points must have $|z|$ less than 50 cm, so that the actual trajectory of the cosmic particle is expected to be safely far from the TIB longitudinal edges. The muons objects used are those belonging to the "1-leg stand-alone muons" collection. These tracks are obtained combining in a *single* object two stand-alone muons reconstructed on two stations on opposite sides of CMS (Fig. 5.4). As a quality cut, the 1-leg muons are required to have more than 52 valid hits among DT and RPC chambers.

The subset of events that enter the numerator class must satisfy the additional condition that a Tracker track is reconstructed. To avoid to overestimate the efficiency because of fake tracks reconstructed in the Tracker from hits of noise, a loose matching criterion is used. The tracks reconstructed in the muon system and those in the silicon Tracker are required to have

similar momenta at the same reference point: $\Delta R < 0.1$, $\delta(p_t)/p_t < 50\%$. As a reference, the trajectory point of closest approach to the beam-line is used.

For cosmic particles that traverse the Tracker volume passing sufficiently close to axes of the barrel layers, the measured tracking efficiency is higher than 99%. Figure 5.4 shows the measured efficiency as a function of the transverse impact parameter (d_0) of the initial muon track. Both in simulation and real data, the efficiency is close to 100% for transverse impact parameters up to about 40 cm. For larger values, the efficiency gradually decreases. This trend is expected since larger d_0 tracks cross fewer barrel layers and have a smaller probability to produce the minimum number of hits requested by the CTF algorithm (5 hits, in the configuration for CRAFT). Although all particles with a d_0 less than about 80 cm are supposed to cross the three outermost barrel layers twice, some of them still produce less than 5 measurements on the detector. The reason is that the Tracker layers are not completely hermetic for trajectories that don't originate from the center of CMS. In some cases the cosmic particles cross the layers at very shallow angle and the corresponding signal is spread on clusters of many strips. Since it is difficult to evaluate the position of these anomalous clusters, the trajectory builder occasionally fails to attach these hits to the track-candidates and the final track is not reconstructed.

5.3 Measurement of the transverse momentum resolution

The resolution with which the transverse momentum of charged particles is measured in the CMS Tracker can be estimated from data exploiting a characteristic of cosmic trajectories. Unlike the charged particles originating from LHC collisions, the cosmic particles can traverse the full volume of the tracking system as in Figure 5.5. In such a case, the upper and bottom parts of the trajectory can be fitted separately. The difference between the two distinctly measured transverse momenta is, under general assumptions, directly connected to the resolution of the transverse momentum measurement. Although this section is devoted uniquely to the transverse momentum measurement, the same "track-splitting" method can be employed to estimate also the resolution on the measurement of the other four track parameters.

Given a cosmic particle with a transverse momentum $p_{T(\text{true})}$, let's indicate with $p_{T(\text{estim,A})}$ and $p_{T(\text{estim,B})}$ the measurements of the transverse momentum using either the hits in the upper or in the lower part of the Tracker respectively (Figure 5.5, bottom). Then X_1 and X_2 are defined as the distributions of $p_{T(\text{estim,A})}$ and $p_{T(\text{estim,B})}$ obtained repeating the measurement

many times. The "track-splitting" method to estimate from data the resolution on the reconstructed transverse momentum assumes that:

1. The two distributions X_1 and X_2 are normally distributed around the true value of the particle's transverse momentum, i.e:

$$X_1 \sim N(p_{T(\text{true})}, \sigma_1^2) \tag{5.1}$$
$$X_2 \sim N(p_{T(\text{true})}, \sigma_2^2) \tag{5.2}$$

2. The resolutions of the two *independent* measurements are the same. This assumption is generally valid if the upper and lower part of the detector are aligned with a comparable precision and the two measurements of the transverse momentum are obtained using the same number of hits and a similar lever arm.

In probability theory, it can be shown that the sum of two independent random variables that are normally distributed is also normally distributed. The resulting distribution has variance (mean) equal to the sum of the variances (means) of the two initial distributions. Defined the difference between the two measurements as:

$$z = p_{T(\text{estim,A})} - p_{T(\text{estim,B})} \tag{5.3}$$

from the first assumption it descends that the random variable z is normally distributed and that $Z \sim N(0, \sigma_1^2 + \sigma_2^2)$. According to the second assumption

$$\sigma_1 = \sigma_2 = \sigma \tag{5.4}$$

and as a result

$$Z \sim N(0, 2 \cdot \sigma^2) \tag{5.5}$$

The Z distribution is estimated from data and its standard deviation, rescaled by $\sqrt{2}$, represents the resolution on the measurement of the transverse momentum.

The validity of the method has been examined on a simulated sample of cosmic particles. The resolution estimated with the "track-splitting" method has been compared with the actual resolution determined doing the difference between the parameters of the simulated particles and those of the corresponding reconstructed tracks. Figure 5.6 shows the resolution on the transverse momentum as a function of the particle p_T for two different scenarios of alignment of the Tracker: the initial alignment of the tracking system due to the precision of the detector assembly and before any correction based on alignment with tracks is applied; a perfectly aligned (ideal) detector. The resolutions, estimated using exclusively information that is

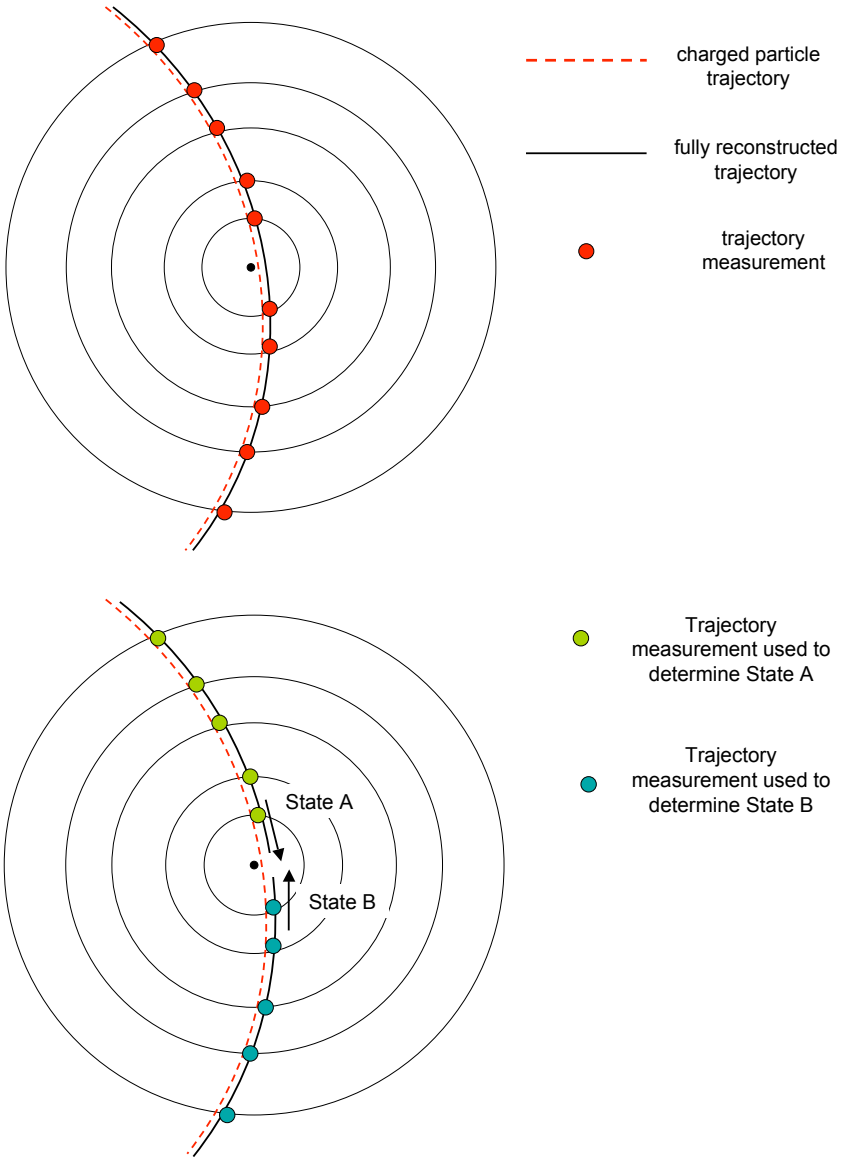

Figure 5.5: *Sketch showing the trajectory of a cosmic muon traversing the full volume of the Tracker system. The trajectory can be fitted using all the measurements provided by the detector (top) or using the measurements in the upper and lower parts of the Tracker separately (bottom). In the latter case, given that the original charged particle is the same, the difference between the two reconstructed values of the trajectory curvature is directly connected to the resolution on the measurement of the curvature itself.*

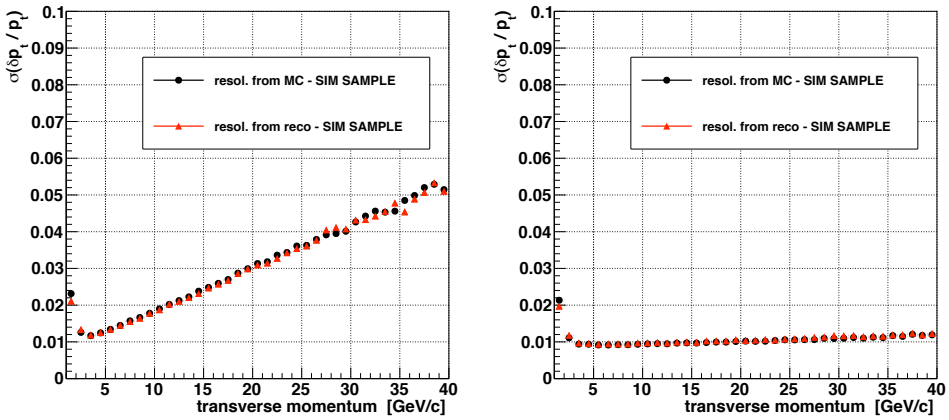

Figure 5.6: *Resolution on transverse momentum measured on simulated data using generator information (black circles) and using only reconstructed information that is available also in real data (red triangles). The sample has been simulated with initial alignment conditions (left) and ideal alignment of the detector (right).*

available also in real data, match very precisely the expected values for both scenarios of alignment. In particular, the "track-splitting" method is able to measure the deterioration of resolution at low energy due to the effects of multiple scattering and the deterioration at high energy due to the detector misalignment.

Only those cosmic particles that point within the volume of the pixel detector were considered in producing, for each bin of transverse momentum, the Z distribution as in Eq. 5.5. Thanks to this selection, the lever arm which is used for the two measurements of the transverse momentum of the cosmic trajectory is comparable with the lever arm that is used to reconstruct the tracks that originate from LHC collisions. The resolution estimated with the "track-splitting" method on a sample of cosmic particles is therefore very close to the resolution obtained for collision-like particles. The comparison between Figures 5.6 and 3.23 confirms it.

The full subset of about 800 thousands cosmic tracks reconstructed in CRAFT and pointing within the pixel detector has been used to estimate the transverse momentum resolution from real data. All the CRAFT tracks, including also those that don't cross the pixel layers, were used previously for the alignment of the tracking system. Starting from an initial uncertainty on the sensors position of about 100-150 μm, the silicon modules have been aligned with a precision that ranges between 10 and 40 μm depending on the Tracker sub-structure [109] where the modules are mounted. After the correction to the modules position, the data from CRAFT has been reprocessed and the

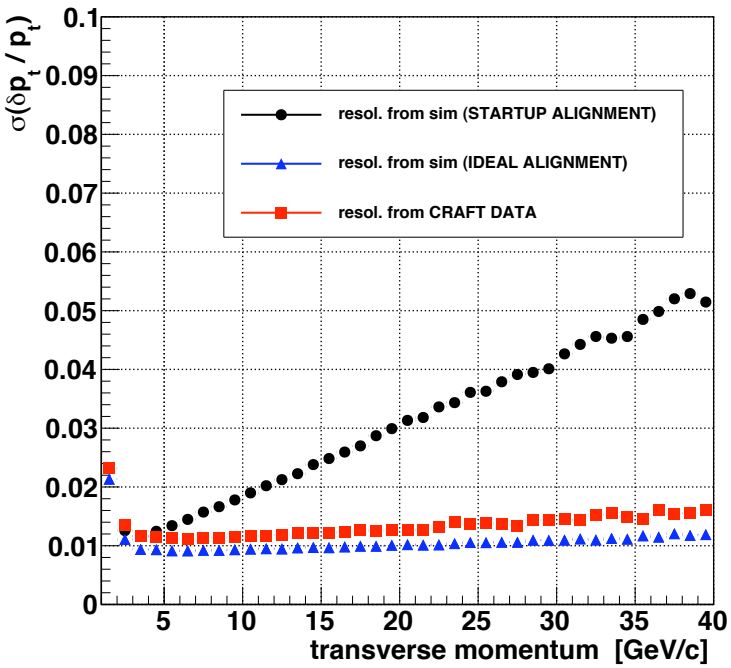

Figure 5.7: *Resolution on transverse momentum, as a function of transverse momentum, measured from CRAF data (red squares). The plot reports also the same kind of resolution evaluated on a simulated sample of cosmic particles that are reconstructed with two different scenarios of Tracker alignment: initial alignment after detector assembly (black circles) and perfect alignment (blue triangles).*

tracks have been reconstructed with improved precision. This second set of reprocessed tracks is the one used by the author for the estimation of the *current* resolution of the detector. The resolution is indeed tightly coupled with the level of alignment of the detector. When a larger sample of cosmic tracks (or tracks from LHC collisions) is used for a new round of alignment of the Tracker, the resulting detector resolution will necessarily improve.

The plot in Figure 5.7 shows the measured resolution as a function of the reconstructed transverse momentum. The same plot reports also the resolution evaluated on a simulated sample for both ideal and initial alignment of the Tracker. Given the current level of alignment of the tracking system, the resolution on the transverse momentum is better than 2% for tracks with p_T greater than 1 GeV/c and between 1% and 1.5% for tracks with p_T greater than 2 GeV/c and smaller than about 30 GeV/c. The current performance of the tracking system is already very close to that which is expected for a

perfectly aligned detector. In the considered range of transverse momenta ($p_T < 40$ GeV/c), a further refinement of the detector alignment cannot improve the resolution on the transverse momentum by more than 0.5%.

Summary and outlook

The silicon tracking system of the CMS detector has been described in detail explaining why it is a vital component for the physics reach of the experiment. The design of its pixel and strip detector modules has been analyzed and the procedure employed to test and qualify, prior to their installation, the latter kind of sensors has been described.

The current CTF (Combinatorial Track Finder) software for the reconstruction of charged particles in the CMS Tracker has been fully described analyzing all its components: trajectory seeding, trajectory building and track fitting modules. The performance of the tracking algorithm has been evaluated on simulated samples of isolated particles and LHC collisions.

The reconstruction of isolated muons results fully efficient within the whole geometrical acceptance of the detector. On the other hand, the reconstruction efficiency for hadrons and electrons is between 90% and 98% depending on the direction of the particle and the corresponding amount of material that the particle crosses before reaching the outer radii of the detector. The effect of the Tracker material budget is significant and affects not only the tracking efficiency, but also the precision with which the trajectory parameters are estimated. Compared to muons, the curvature of hadrons and electrons is usually reconstructed exploiting a shorter effective lever arm respectively because of inelastically nuclear interactions and energy loss by bremsstrahlung. In the barrel region of the Tracker, the transverse momentum is measured with a resolution between 0.5% and 1.5% for most of the charged particles with p_T in the range of 1-100 GeV/c. However, unlike for muons, the resolution distributions for electrons and hadrons contain long tails as a consequence of the large Tracker material budget. While a dedicated algorithm for the track fitting of electrons that radiate a significant fraction of their energy by bremsstrahlung has already been developed in CMS, an efficient nuclear vertex finder has not been implemented yet. The latter may be employed to improve the reconstruction of the transverse momentum for those hadrons that inelastically interact after they cross few layers of the tracking system.

The performance studies on simulated LHC collisions show that an efficiency

between 85% and 93% is achieved for the reconstruction of primary charged particles in multi-track events, while the overall fake rate is kept below 1%. For very collimated jets of particles, the tracking efficiency drops to less than 75% since the reconstruction algorithm is not able to unambiguously assign the positions measurements to the corresponding tracks. The reconstruction would benefit by the use of a more sophisticated algorithm that allows a *soft* assignment of hits to tracks. While two such algorithms have been already implemented in CMS, they have still to be tuned and they are not yet employed within the standard reconstruction software. The correct assignment of hits to tracks is also expected to benefit from the use of the cluster charge and shape information during trajectory building. Tools that verify the compatibility between the cluster shape and the incidence angle of the tracks are currently under investigation and may be employed by default in the next version of the tracking algorithm.

Three main extensions to the CTF algorithm have been implemented and the improvements of the tracking performance have been analyzed:

- A new seed generator that combines pixel and strip measurements has been developed. With respect to the original seed generator that was exclusively based on pixel measurements, the new seeder allows the recovery of all those tracks that were lost because of geometrical inefficiencies of the pixel sub-system. Also, the new seeder can be used to define a configurations of the CTF algorithm that doesn't rely on the pixel detector at all. Such setup has been used for the reconstruction of cosmic ray trajectories during the commissioning of the detector.

- A new trajectory builder that identifies more efficiently the hits produced by the charged particles in the Tracker has been implemented and described in detail. The hit-finding efficiency increases from the previous value of 87% to more than 98%. An important feature of the new builder is the possibility to associate to a track more than one hit per Tracker layer when the charged particle crosses two adjacent sensors. The correlation between such pair of measurements and the very short distance between them are vital for monitoring the sensor resolution from data. They also facilitate the alignment of the detector using tracks.

- A method for the numerical integration of the equations of motion within a non uniform magnetic field has been implemented and integrated inside the tracking software. The bias in the reconstructed transverse momentum is reduced from about 1 standard deviation to less than 0.1 for particles with $p_T = 10$ GeV/c and traversing the endcap region of the Tracker.

The same tracking software developed for the reconstruction of charged particles originating from LHC collisions has been adapted to reconstruct also the trajectories of cosmic ray particles. During the 2008 global runs of CMS, the CTF tracking algorithm has been employed for the first time in the reconstruction of tracks using the full tracking system in its final configuration. The sample of reconstructed tracks has provided a vital input for the commissioning and the alignment of the detector. Since the magnet of CMS was turned on during the data taking, the tracking system was able to determine for the first time the momenta of the reconstructed trajectories. The efficiency of the track reconstruction has been estimated from data and results higher than 99% for cosmic ray trajectories that have topologies similar to those of particles produced in LHC collisions. Also the resolution on the transverse momentum has been estimated from data and results between 1% and 2% for particle with p_T smaller than 40 GeV/c. The performance is consistent with the expected level of alignment of the detector.

Appendix A

Detailed description of CMS Tracker layout

A.0.1 Pixel Barrel system

The Pixel Barrel system consists of three cylindrical layers of radii 4.4, 7.3 and 10.2 cm, respectively. The three layers have different numbers of equidistant ϕ-faces, each containing 8 pixel modules on one ladder or 16 half-modules on two ladders. Because the barrel system is mechanically divided into two half-cylinders, a set of half-modules is necessary for the disjoint ϕ-faces at $\phi = \pm\pi/2$ (Fig. A.2).

All three layers incorporate the same type of barrel pixel modules (half-modules) shaped in 2 raws (1 raw) of 8 ROCs. The pixel cells are oriented with their 150- μm-large edges along the z-coordinate and the 100- μm sides along ϕ. Each ladder of modules is not fully hermetic in z, but contains 2.2-mm-wide gaps between the sensitive areas of contiguous detectors, for an overall longitudinal extension of $\pm$26.7 cm. On the other hand, adjacent ladders overlap in ϕ to assure complete azimuthal hermeticity. Altogether, the three barrel layers contain 768 detector elements (modules and half-modules) built from 300- μm-thick silicon wafers. The total number of read-out channels is approximately 48 milions.

layer	radius [cm]	# ϕ-faces	# detector elements	# ROC/mod	# channels
1	4.4	18	128 modules + 32 half-modules	8×2	9.6 M
2	7.3	30	224 modules + 32 half-modules	8×2	16.0 M
3	10.2	42	320 modules + 32 half-modules	8×2	22.4 M

Table A.1: Layout of the three layers of the Pixel Barrel system.

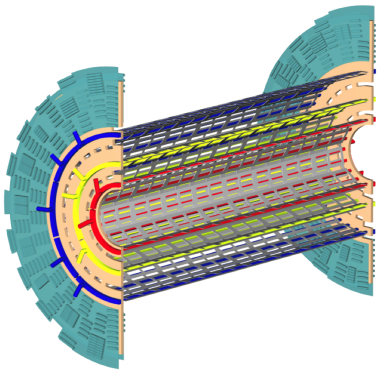

Figure A.1: Support structure of one of the two pixel barrel half-cylinders. Pixel modules are assembled on three layers of radii 4.4, 7.3 and 10.2 cm.

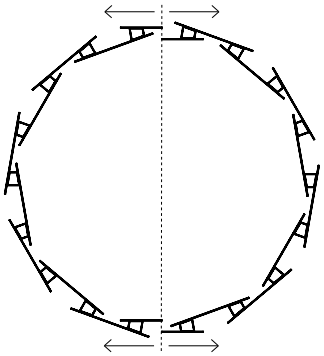

Figure A.2: Transverse section of the innermost pixel barrel layer, which consists of 18 faces overlapping in azimuth. Each of the two faces at $\phi = \pm\pi/2$ is made of ladders of half-modules.

total # detector elements	768
total # silicon wafers	768
total # ROCs	11 520
total # pixel channels	48 M

Table A.2: Summary table for the Pixel Barrel system.

A.0.2 Pixel Forward system

The Pixel Forward system completes the central barrel on each side with two pairs of disks, which are located perpendicularly to the z-axis at ± 34.5 and ± 46.5 cm from the IP. Each disk extends radially between 6 and 15 cm and has sensitive sensors placed on 24 blades arranged in a turbine-like geometry: blades are rotated by 20° around their radial axes to increase the charge sharing, and improve the position resolution, for particles coming from LHC collisions. Every turbine element hosts pixel modules, called "plaquettes", on both the surface facing the IP (front panel) and the opposite one (back panel). The front panel has 4 modules with ROCs grouped in 4 distinct raws × columns arrangements: 1×5, 2×4, 2×3 and 1×2. Conversely, the back panel has read-out chips grouped in three 2×5, 2×4 and 2×3 modules, for a total of 45 ROCs per blade. The pixel cells are oriented with their wider edges along the azimuthal coordinate and the shorter ones along the radial coordinate.

In order to assure the full hermeticity of the FPIX system, contiguous modules on opposite sides of the same blade overlap radially, while pairs of contiguous blades overlap in ϕ. To allow the insertion of the disks after the beam pipe has been baked out, each forward layer is mechanically divided

into two half-turbines.

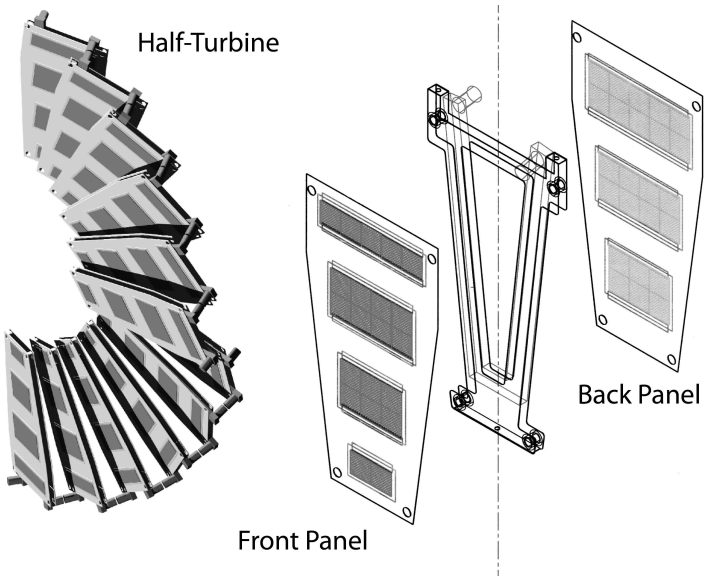

Figure A.3: One half of a forward pixel disk supporting sensitive modules in a turbine-like geometry (left). Exploded view of a single turbine blade that hosts sensors on both sides (right).

# disks	2 × 2
disks z-positions [cm]	±34.5 / ±46.5
disks inn./out. radii [cm]	6.0/15.0
# blades per disk	24
# "plaquettes" per blade	4+3
# ROCs per blade	21+24

Table A.3: Layout of the two pairs of Pixel Forward disks.

total # detector elements	672
total # silicon wafers	672
total # ROCs	4 320
total # pixel channels	18 M

Table A.4: Summary table for the Pixel Forward system.

Altogether the forward pixel system contains 672 modules, built from 300-μm-thick wafers, for a total of 18M of read-out channels.

A.0.3 Tracker Inner Barrel and Tracker Inner Disks

The Tracker Inner Barrel system consists of four ±73-cm-long cylindrical layers spanning, with active components, the radial range between 23.9 and 51.4 cm. To facilitate the integration of silicon modules, electronics and services, each layer is mechanically sub-divided into 4 sub-assemblies ($\pm z$, up-down) called half-shells. The silicon modules are directly placed on the half-shells by means of small aluminum ledges that are glued to the carbon fiber structure of the layer. The ledges serve also as thermal contact

between the modules and the cooling pipes. Every half-shell bears strings of 3 modules on both the inner and the outer surface. Contiguous strings, placed on the same side of a layer, have a ϕ-overlap between 1.1 and 3.4 mm (Fig. A.5). Similarly, the modules on the inner strings are staggered in z with respect to the ones on the outer strings to assure hermeticity of the whole layer. After all the sub-assemblies are connected, the active sensors of each layer are geometrically aligned in 6 inner and 6 outer ϕ-rings (Fig. A.4).

The TIB system mounts two types of detector elements: the double-sided (DS) modules consist of pairs of rectangular-shaped sensors that are rotated by 100 mrad and then glued back-to-back; single-sided (SS) modules are made of single rectangular sensors. All four layers have sensors with strips running along the $z - axis$ and, therefore, measuring the ϕ-coordinate. On the other hand, the DS modules situated on the two innermost layers are able to measure also the z-coordinate. All sensors on DD layers host 6 APV chips and 768 strips with a pitch of 80 μm; sensors on SS layers have only 4 APVs, 512 strips and a larger pitch equal to 120 μm.

Because of the less stringent requirements on resolution with respect to the pixel detector, in order to lower their occupancy the TIB modules are tilted by 9° to compensate for the Lorentz angle and reduce the spread of the induced signal over several strips (Fig. A.5). This arrangement of modules allows the detectors to have a more symmetric response to positive and negative charged particles and facilitate the overlap in ϕ among sensors of the same layer.

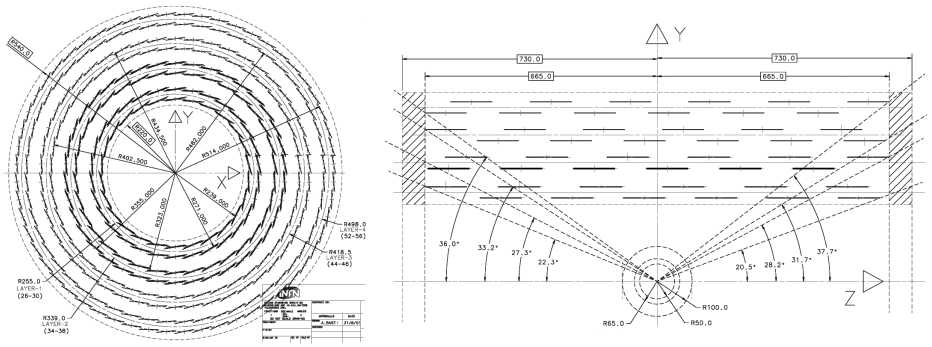

Figure A.4: Transverse (left) and longitudinal (right) views of the Tracker Inner Barrel system. Active sensors are placed on both the inner and the outer surfaces of all the four system layers.

The Tracker Inner Disks system consists of two groups of three disks extending on both sides of the TIB. Each disk contains 3 rings of modules spanning, between 23.3 and 50.3 cm, almost the same radial range of the inner barrel system. Similarly to TIB, the TID modules are directly mounted on the carbon fiber support structure of the layer by means of aluminum

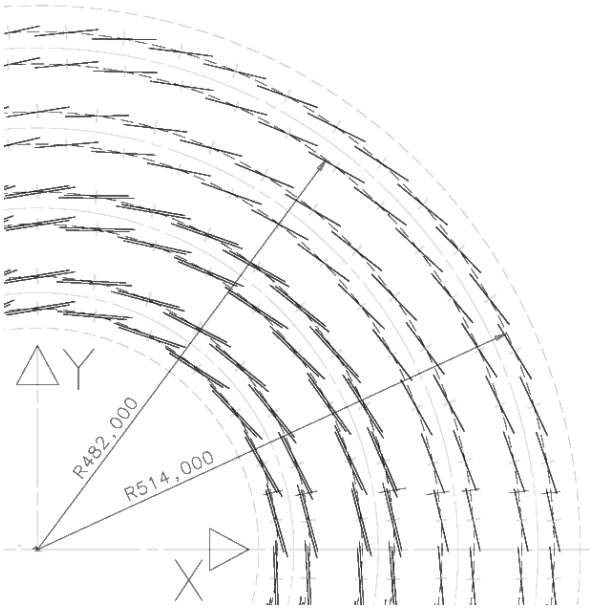

Figure A.5: Within the same layer, the Tracker Inner Barrel modules are tilted by 9° with respect to the radial coordinate r to compensate for the Lorentz angle and to facilitate the ϕ-overlap between contiguous sensors .

layer	Avg.radius [cm]	# ϕ-strings	# modules	APV/mod	pitch mono/stereo [μm]	thickness [μm]
TIB1	25.5	26-30	336	6+6	80/80	320
TIB2	33.9	34-38	432	6+6	80/80	320
TIB3	41.8	44-46	540	4	120/-	320
TIB4	49.8	52-56	648	4	120/-	320

Table A.5: Layout of the 4 cylindrical layers of the Tracker Inner Barrel system.

ledges: the innermost and the outermost rings are placed on the disk's face pointing toward the IP, while the middle ring is located on the other side of the layer (Fig. A.6). Contiguous modules alternate around the average z-position of the ring and overlap by 1-2 mm in ϕ. Adjacent rings overlap in r, but not enough to make the forward layers completely hermetic: a high energy particle coming straight from the IP has a a 5% probability to go through the innermost disk without crossing any sensitive module and a 7% probability to go through the whole TID system missing, at least, one layer.

All the TID modules are made of trapezoidal-shaped sensors with strips running radially. The two innermost rings of each disk have DS modules measuring both the ϕ and the r coordinates; conversely, the third outermost ring has single-sided modules measuring only ϕ. The sensors of DS layers have 6 APVs and 712 strips each with a pitch between 81 and 143 μm,

while the sensors of the SS layers have only 4 APV and 512 strips with wider pitch (Tab. A.6).

ring	radius in/out [cm]	modules in ϕ	# rings in z	# modules	APV/mod	mono/stereo pitch [μm]	thickness [μm]
R1	23.3/32.0	24	3 × 2	144	6+6	81-112/81-112	320
R2	32.3/41.1	24	3 × 2	144	6+6	113-143/113-143	320
R3	39.2/50.3	40	3 × 2	240	4	123-158/ -	320

Table A.6: Layout of the Tracker Inner Disks system.

	TIB	TID
total # modules	1 956 (1 188 SS + 768 DS)	528 (240 SS + 288 DS)
total # det. elements	2 724	816
total # silicon wafers	2 724	816
total # APVs	13 968	4 416
total # strip channels	1.8 M	0.6

Table A.7: Summary table for the Tracker Inner Barrel and Forward Disks system.

Altogether the TIB/TID system contains 2484 modules built from 3540 silicon wafers with a thickness of 320 μm. The total number of strips that are read out is equal to 2.4 millions.

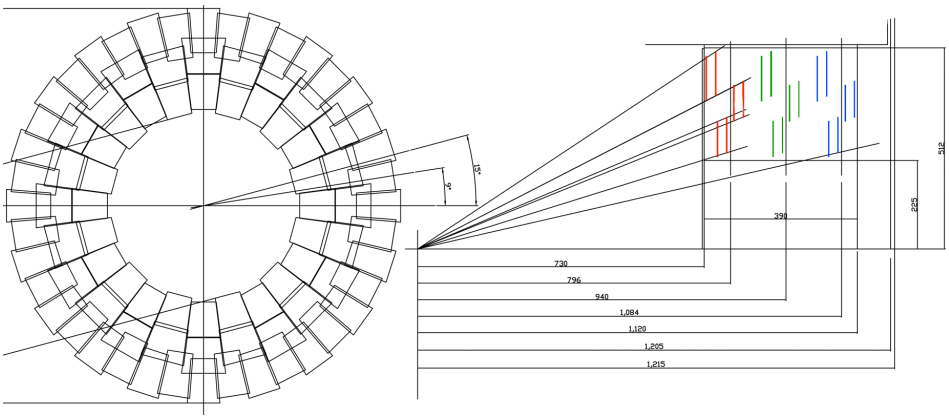

Figure A.6: Transverse (left) and longitudinal (right) views of the Tracker Inner Disks system. Each layer has sensitive trapezoidal-shaped detectors arranged in three concentric rings. The red, green and blue colors refer to rings of the first, second and third TID disks, respectively.

A.0.4 Tracker Outer Barrel

The Tracker Outer Barrel system consists of six cylindric layers spanning, with active components, the radial range between 59 and 110 cm. With

a longitudinal size of ± 110 cm, the TOB layers match the length of the TIB/TID system. A carbon fiber and aluminum *wheel* (Fig. A.7) is the single mechanical structure supporting all the system sub-assemblies, called *rod*s. After they are independently assembled and tested, all the rods are inserted along z into the support structure through rings of openings arranged on both side of the wheel.

Each rod has 6 positions where 6 detectors or 6 pairs of detectors are installed (Fig. A.9). The four outer layers of TOB have rods containing only single detectors with strips aligned along z and measuring the ϕ coordinate; the two innermost layers have rods supporting pairs of detectors with a small stereo angle and measuring both ϕ and z. These pairs of detectors are not glued together as the two elements of a TIB double-sided module, but are usually indicated as double-sided modules of TOB because of their proximity inside the same rod. Two rods, inserted from opposite sides of the support wheel, are enough to cover the full length of the TOB system with 12 modules.

Within the same layer, rods are positioned at two different radii in such a way that contiguous rods can overlap in ϕ and increase the hermeticity of the layers (Fig. A.8). Similarly, inside a rod, three modules are arranged above the central rod surface and three are positioned below it to implement an z-overlap between consecutive modules. Every detector of TOB is made out of two 500-μm-thick silicon wafers that are wire-bonded together into a single sensitive element. The gap between the pair of sensors is 0.1 mm, producing an inefficiency which is smaller than 0.1%.

layer	Avg.radius [cm]	# ϕ-slots	# modules	APV/mod	mono/stereo pitch [μm]	thickness [μm]
TOB1	60.8	42	504	4+4	183/183	500
TOB2	69.2	48	576	4+4	183/183	500
TOB3	78.0	54	648	4	183/-	500
TOB4	86.8	60	720	4	183/-	500
TOB5	96.5	66	792	6	122/-	500
TOB6	108.0	74	888	6	122/-	500

Table A.8: Layout of the 6 cylindrical layers of the Tracker Outer Barrel system.

Independently if they are used for single or double-sided modules, all the silicon detectors of the layers 1-4 have the same pitch (183 μm) and the same number of APVs (4) and strips (512). The detectors of the two outermost layers have a finer pitch (122 μm) and more strips (768 channels read out by 6 APVs).

Altogether the 6 layers of the TOB system contain 5 208 silicon detectors arranged in more than 4 thousands modules for a total of 3.1 M read-out channels (Tab. A.9). An excess of 10 thousands silicon sensors is used.

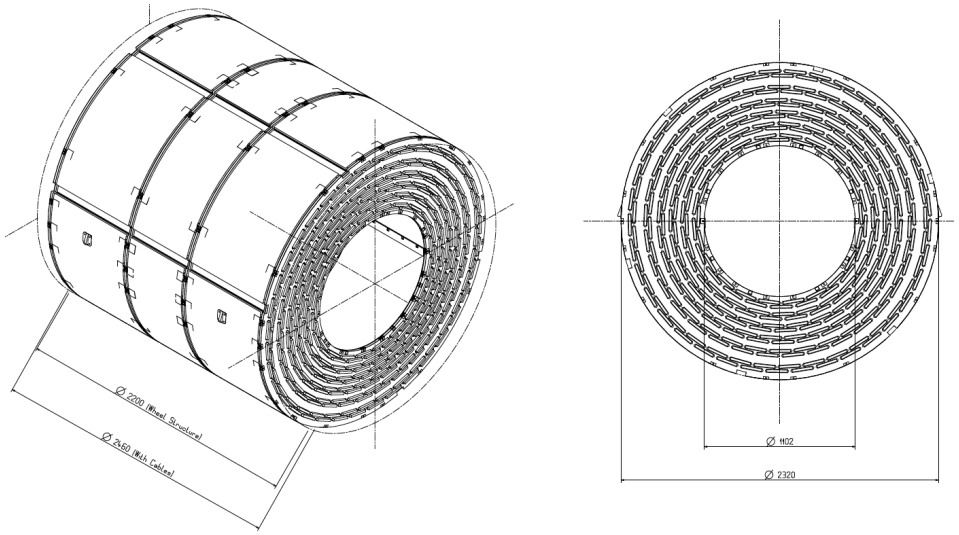

Figure A.7: Carbon fiber and aluminum mechanical structure that supports all the sub-assembies of the Tracker Outer Barrel system. Six rings of openings on both sides of the wheel allow the detectors insertion and their arrangement into six sensitive cylindrical layers.

total # modules	4 128 (3 048 SD + 1 080 DS)
total # detector elements	5 208
total # silicon wafers	10 416
total # APVs	24 192
total # strip channels	3.1 M

Table A.9: Summary table for the Tracker Outer Barrel system.

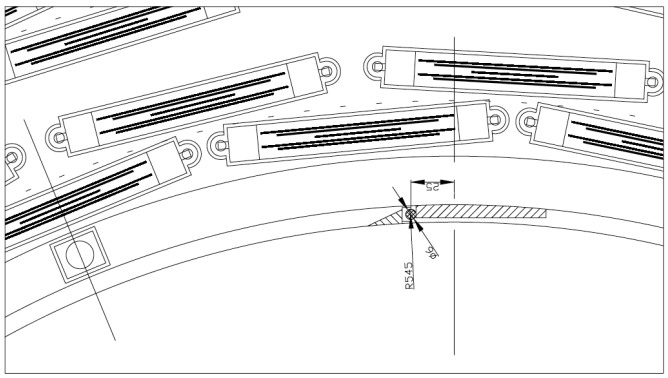

Figure A.8: Transverse view of few rods of modules inserted into the most internal layer of the Tracker Outer Barrel system. Within the same layer, the rods are installed at two different radii to allow a ϕ-overlap between contiguous sub-assemlies.

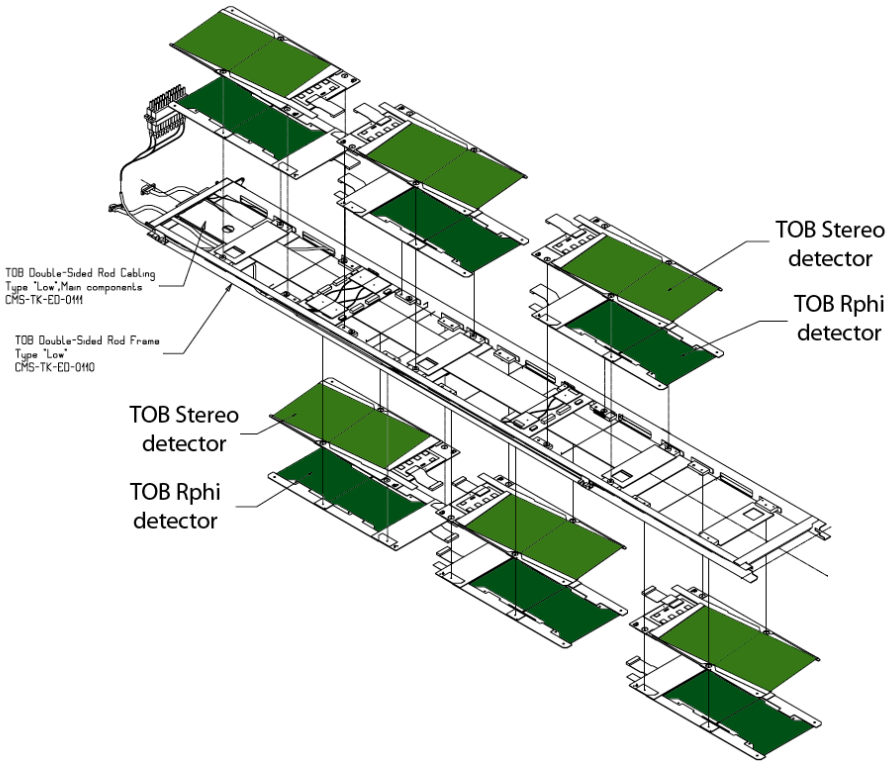

Figure A.9: Exploded view of a Tracker Outer Barrel rod containing 6 double-sided modules.

A.0.5 Tracker End-caps Disks

The Tracker End-caps Disks system extends the TOB/TIB/TID barrel on both z-sides with two sets of 9 disks (Fig. A.10 right), arranged between $+(-)132$ cm and $+(-)267$ cm. The silicon detectors are mounted on 3 levels of mechanical structures:

1. Each layer of the TEC consists of a carbon fiber disk that hosts 16 sub-assemblies on both sides: 8 front petals and 8 back petals. Each petal spans approximately 22° in azimuth and contiguous ones, on opposite sides of the support disk, have a small overlap in ϕ (Fig. A.10 left).

2. Both front and back petals have 4 wedges of detectors placed on the side facing the IP and 3 wedges on the opposite one. Nevertheless, the two type of petals have wedges of different sizes containing a different number of detectors (Figures A.11). Contiguous wedges are staggered in z so that they overlap radially by 1-3 mm.

3. Each wedge supports between 1 and 5 trapezoidal-shaped silicon modules. Consecutive detectors are staggered in z to allow a r-ϕ overlap of 1-2 mm.

With all the 16 petals integrated over a disk layer, the silicon modules are arranged in 7 concentric rings. The azimuthal overlap between modules is constant within the same ring, but it is generally higher for the outer rings respect to the inner ones. The transverse disposition of the detectors within the three innermost TEC rings equal that in the three rings of the TID system.

Double-sided modules, that are made out of pairs of detectors glued back-to-back with a small stereo angle, are positioned on rings 1,2 and 5. Because one of the two components of the DD modules has its strips disposed radially, these 3 rings provide an accurate measurement of the ϕ coordinate and a rougher one of the r coordinate. All other rings have SS modules measuring only ϕ. Following the different requirements of resolution and occupancy, the pitch and length of strips vary from ring to ring. Detectors with 768 strips each and finest pitch are positioned on rings R1 and R2, while detectors on R3 and R4 have only 512 strips with larger separation. Modules on the three outermost rings R5-R7 have even larger pitch and are made out of pairs of daisy chained 500-μm-thick sensors, practically doubling the strip length (Tab. A.10).

Only few disks of each TEC end-cap have their petals completely instrumented. While the three pairs of innermost disks TEC1,TEC2 and TEC3 are equipped with all 7 rings of modules, the number of instrumented rings progressively reduces to only three (R4-R7) moving towards the outermost

disks TEC9 (Tab. A.11). Anyhow, the setup of the TEC system assures that every disk mantains a pseudorapidity coverage up to $|eta| = 2.4 - 2.5$ for particles coming from the IP.

Altogether the 2×9 layers of the TEC system contain 6400 silicon detectors arranged in more than 5 thousands modules for a total of 3.9 M of read-out strips; All the single and daisy-chained sensors are made out of more than 10 thousands silicon wafers (Tab. A.10).

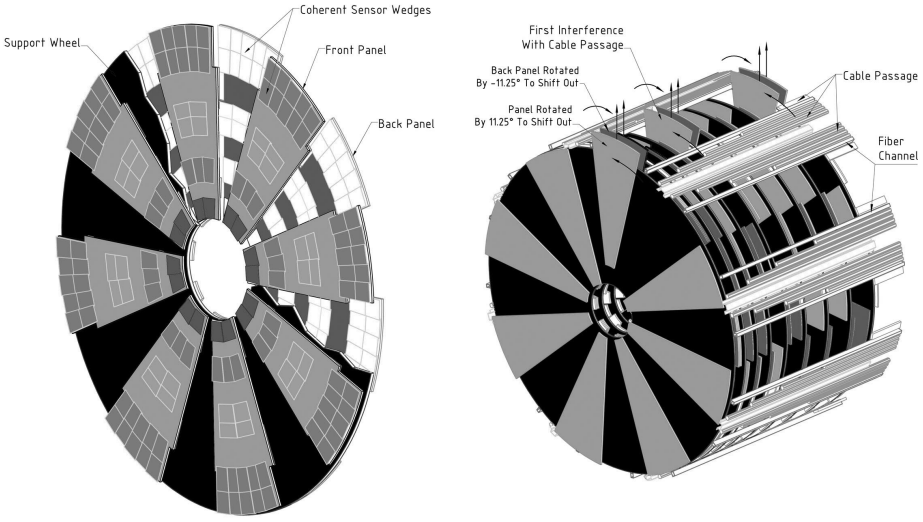

Figure A.10: Layout of theTracker Endcap Disks (TEC) system. Each disk supports 16 petals of detectors arranged uniformly on both sides (left); nine disks add up into each of the two TEC end-caps (right).

ring	radius in/out [cm]	modules in ϕ	# rings in z	# modules	APV/mod	mono/stereo pitch [μm]	thickness [μm]
R1	23.3/32.0	24	3×2	144	6+6	81-112/81-112	320
R2	32.3/41.1	24	6×2	288	6+6	113-143/113-143	320
R3	39.2/50.3	40	8×2	640	4	123-158/-	320
R4	50.4/61.9	56	9×2	1 008	4	113-139/-	320
R5	60.3/75.0	40	9×2	720	6+6	126-156/126-156	500
R6	72.7/91.1	56	9×2	1 008	4	163-205/-	500
R7	88.8/109.4	80	9×2	1 440	4	140-172/-	500

Table A.10: Layout of the Tracker End-caps Disks.

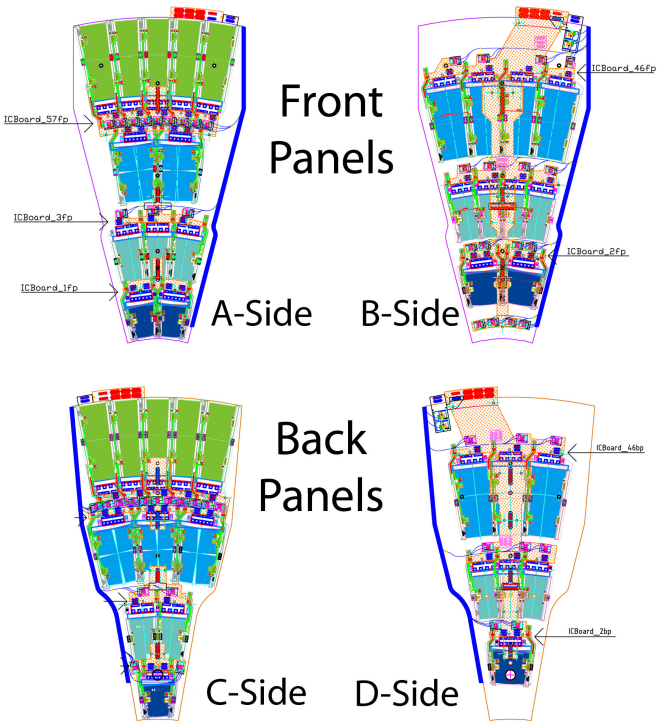

Figure A.11: Layout of petals used in the Tracker End-cap Disks (TEC) system. Each petal contains 4 wedges of silicon modules on the surface facing the interaction point (IP) and 3 wedges on the opposite side. Wedges on opposite faces of the petal are staggered to allow overlap in r. Similarly, modules within the same wedge overlap in ϕ. Front petals (upper) are installed on the surfaces of TEC disks looking toward IP; back petals (lower) are installed on the opposite side of the disks.

disk	avg. z [cm]	rings
TEC1	±132	R1-R7
TEC2	±146	R1-R7
TEC3	±160	R1-R7
TEC4	±174	R2-R7
TEC5	±188	R2-R7
TEC6	±206	R2-R7
TEC7	±225	R3-R7
TEC8	±245	R3-R7
TEC9	±267	R4-R7

Table A.11: Position in z and rings composition of the Tracker End-cap Disks.

total # modules	5248 (4096 SS + 1152 DS)
total # det. elements	6400
total # silicon wafers	10 288 (2 512 thin + 7 776 thick)
total # APVs	30 208
total # strip channels	3.9 M

Table A.12: Summary table for the Tracker End-cap Disks system.

Bibliography

[1] Lyndon Evans and Philip Bryant (editors). Lhc machine. *Journal of Instrumentation*, 3(08):S08001, 2008.

[2] M. Della Negra et al. CMS: The Compact Muon Solenoid: Letter of intent for a general purpose detector at the LHC. CERN-LHCC-92-03.

[3] CMS, the Compact Muon Solenoid: Technical proposal. CERN-LHCC-94-38.

[4] The CMS Collaboration and S Chatrchyan *et al.* The cms experiment at the cern lhc. *Journal of Instrumentation*, 3(08):S08004, 2008.

[5] S. L. Glashow. Partial Symmetries of Weak Interactions. *Nucl. Phys.*, 22:579–588, 1961.

[6] Steven Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19:1264–1266, 1967.

[7] Murray Gell-Mann. A Schematic Model of Baryons and Mesons. *Phys. Lett.*, 8:214–215, 1964.

[8] H. Fritzsch, Murray Gell-Mann, and H. Leutwyler. Advantages of the Color Octet Gluon Picture. *Phys. Lett.*, B47:365–368, 1973.

[9] D. J. Gross and Frank Wilczek. Ultraviolet behavior of non-Abelian Gauge Theories. *Phys. Rev. Lett.*, 30:1343–1346, 1973.

[10] H. David Politzer. Reliable Perturbative results for Strong Interactions. *Phys. Rev. Lett.*, 30:1346–1349, 1973.

[11] Gerard 't Hooft. Renormalization of Massless Yang-Mills Fields. *Nucl. Phys.*, B33:173–199, 1971.

[12] Gerard 't Hooft. Renormalizable Lagrangians for Massive Yang-Mills fields. *Nucl. Phys.*, B35:167–188, 1971.

[13] Gerard 't Hooft and M. J. G. Veltman. Regularization and Renormalization of Gauge Fields. *Nucl. Phys.*, B44:189–213, 1972.

[14] P.W. Higgs. Broken symmetries, massless particlees and gauge fields. *Physics Letters*, 12(2):132 – 133, 1964.

[15] Peter W. Higgs. Spontaneous symmetry breakdown without massless bosons. *Phys. Rev.*, 145(4):1156–1163, May 1966.

[16] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Phys. Rev. Lett.*, 13:321–322, 1964.

[17] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. Global Conservation Laws and Massless particles. *Phys. Rev. Lett.*, 13:585–587, 1964.

[18] G. Arnison et al. Experimental observation of isolated large transverse energy electrons with associated missing energy at. *Physics Letters B*, 122(1):103 – 116, 1983.

[19] G. Arnison et al. Experimental observation of lepton pairs of invariant mass around 95 gev/c2 at the cern sps collider. *Physics Letters B*, 126(5):398 – 410, 1983.

[20] G. Arnison et al. Further evidence for charged intermediate vector bosons at the sps collider. *Physics Letters B*, 129(3-4):273 – 282, 1983.

[21] G. Arnison et al. Observation of the muonic decay of the charged intermediate vector boson. *Physics Letters B*, 134(6):469 – 476, 1984.

[22] E. Radermacher. The experimental discovery of the Intermediate Vector Bosons W^+, W^- and Z^0 at the CERN $p\bar{p}$ collider. *Prog. Part. Nucl. Phys.*, 14:231–328, 1985.

[23] F. Abe and others. Observation of top quark production in $p\bar{p}$ collisions with the collider detector at fermilab. *Phys. Rev. Lett.*, 74(14):2626–2631, Apr 1995.

[24] The LEP Electroweak Working Group.
Site located at http://lepewwg.web.cern.ch/LEPEWWG/.

[25] ALEPH Collaboration; DELPHI Collaboration; L3 Collaboration; OPAL Collaboration; The LEP Working Group for Higgs Boson Searches. Search for the Standard Model Higgs boson at LEP . *Physics Letters B*, 565:61–75, 2003.

[26] Benjamin W. Lee, C. Quigg, and H. B. Thacker. Weak Interactions at Very High-Energies: The Role of the Higgs Boson Mass. *Phys. Rev.*, D16:1519, 1977.

[27] Jean Iliopoulos. Physics Beyond the Standard Model. 2008.

[28] Hitoshi Murayama. Physics Beyond the Standard Model and Dark Matter. 2007.

[29] Stephen P. Martin. A Supersymmetry Primer. 1997.

[30] John F. Donoghue. The fine-tuning problems of particle physics and anthropic mechanisms. 2007.

[31] A Morales. Experimental searches for non-baryonic dark matter: Wimp direct detection. *Nucl. Phys. B, Proc. Suppl.*, 110(astro-ph/0112550):39–54, Dec 2001.

[32] T. J. Sumner. Experimental searches for dark matter. *Living Rev. Rel.*, 5:4, 2002.

[33] Mariano Quirs. New ideas in symmetry breaking. Technical Report hep-ph/0302189, Feb 2003.

[34] (ed.) Armesto, N. et al. Heavy Ion Collisions at the LHC - Last Call for Predictions. *J. Phys.*, G35:054001, 2008.

[35] Robert Rathbun Wilson. The Tevatron. *Phys. Today*, 30N10:23–30, 1977.

[36] The ATLAS Collaboration and G Aad *et al.* The atlas experiment at the cern large hadron collider. *Journal of Instrumentation*, 3(08):S08003, 2008.

[37] The ALICE Collaboration and K Aamodt *et al.* The alice experiment at the cern lhc. *Journal of Instrumentation*, 3(08):S08002, 2008.

[38] The LHCb Collaboration and A Augusto Alves Jr *et al.* The lhcb detector at the lhc. *Journal of Instrumentation*, 3(08):S08005, 2008.

[39] Conceptual Design of the Superconducting Super Collider. SSC-SR-2020.

[40] G. Acquistapace et al. CMS, the magnet project: Technical design report. CERN-LHCC-97-10.

[41] C. Albajar et al. Test beam analysis of the first cms drift tube muon chamber. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 525(3):465 – 484, 2004.

[42] P. Lecoq, I. Dafinei, E. Auffray, M. Schneegans, M. V. Korzhik, O. V. Missevitch, V. B. Pavlenko, A. A. Fedorov, A. N. Annenkov, V. L. Kostylev, and V. D. Ligun. Lead tungstate (pbwo4) scintillators for

lhc em calorimetry. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 365(2-3):291 – 298, 1995.

[43] M. Kobayashi, Y. Usuki, M. Ishii, N. Senguttuvan, K. Tanji, M. Chiba, K. Hara, H. Takano, M. Nikl, P. Bohacek, S. Baccaro, A. Cecilia, and M. Diemoz. Significant improvement of pbwo4 scintillating crystals by doping with trivalent ions. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 434(2-3):412 – 423, 1999.

[44] A. A. Annenkov, M. V. Korzhik, and P. Lecoq. Lead tungstate scintillation material. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 490(1-2):30 – 50, 2002.

[45] W. Bertl et al. Feasibility of intercalibration of CMS ECAL supermodules with cosmic rays. *Eur. Phys. J.*, C41S2:11–17, 2005.

[46] F Ferri and P Govoni. The cms electromagnetic calorimeter precalibration with cosmic rays and test beam electrons. Technical Report CMS-CR-2007-012. CERN-CMS-CR-2007-012, CERN, Geneva, 2007.

[47] A Heister, O Kodolova, V F Konoplyanikov, Sergey V Petrushanko, J Rohlf, C Tully, and A Ulyanov. Measurement of jets with the cms detector at the lhc. Technical Report CMS-NOTE-2006-036. CERN-CMS-NOTE-2006-036, CERN, Geneva, Feb 2006.

[48] Haifeng Pi, P Avery, D Green, J Rohlf, and C Tully. Measurement of missing transverse energy with the cms detector at the lhc. Technical Report CMS-NOTE-2006-035. CERN-CMS-NOTE-2006-035, CERN, Geneva, Feb 2006.

[49] Nural Akchurin and Richard Wigmans. Quartz fibers as active elements in detectors for particle physics. *Review of Scientific Instruments*, 74(6):2955–2972, 2003.

[50] G. L. Bayatian et al. CMS physics: Technical design report. CERN-LHCC-2006-001.

[51] G. L. Bayatian et al. CMS technical design report, volume II: Physics performance. *J. Phys.*, G34:995–1579, 2007.

[52] L. Wendland. Hadronic tau jet reconstruction with particle flow algorithms at CMS. Technical Report CMS-CR-2008-093. CERN-CMS-CR-2008-093, CERN, Geneva, Dec 2008.

[53] (ed.) Sphicas, P. CMS: The TriDAS project. Technical design report, Vol. 2: Data acquisition and high-level trigger. CERN-LHCC-2002-026.

[54] Addendum to the CMS tracker TDR. CERN-LHCC-2000-016.

[55] CMS, tracker technical design report. CERN-LHCC-98-06.

[56] Gavril Giurgiu, D. Fehling, P. Maksimovic, M. Swartz, and V. Chiochia. Pixel Hit Reconstruction with the CMS Detector. 2008.

[57] L Borrello, A Messineo, E Focardi, and A Macchiolo. Sensor design for the cms silicon strip tracker. Technical Report CMS-NOTE-2003-020, CERN, Geneva, Aug 2003.

[58] J. L. Agram et al. The silicon sensors for the Compact Muon Solenoid tracker: Design and qualification procedure. *Nucl. Instrum. Meth.*, A517:77–93, 2004.

[59] G. Lindstrom and S. Watts. Second RD48 status report: R & D on silicon for future experiments. CERN-LHCC-98-039.

[60] H. Chr. Kastli et al. Design and performance of the CMS pixel detector readout chip. *Nucl. Instrum. Meth.*, A565:188–194, 2006.

[61] D. Kotlinski et al. The control and readout systems of the CMS pixel barrel detector. *Nucl. Instrum. Meth.*, A565:73–78, 2006.

[62] C. F. Parkman. VME bus at CERN: a brief review. Presented at Conf. Big Science Centers and their Role as Promoters of Industrial Development, Uppsala, Sweden, Oct 17, 1985.

[63] M. Raymond et al. The CMS tracker APV25 0.25-mu-m CMOS readout chip. Prepared for 6th Workshop on Electronic for LHC Experiments, Cracow, Poland, 11-15 Sep 2000.

[64] D Ricci, B Checcucci, G M Bilei, F Casinini, and V Postolache. Experience in the development of the cms inner tracker analog optohybrid circuits: Project, qualification, volume production, quality assurance and final performance. Technical Report CMS-NOTE-2005-015. CERN-CMS-NOTE-2005-015, CERN, Geneva, Nov 2005.

[65] M. Friedl. Analog optohybrids for the readout of the CMS silicon tracker. *Nucl. Instrum. Meth.*, A518:515–518, 2004.

[66] J A Coughlan, S A Baird, I Church, E Corrin, C P Day, C Foudas, E J Freeman, J R Fulcher, W J F Gannon, G Hall, R Halsall, G Iles, M Noy, M R Pearson, I Reid, G Rogers, J Salisbury, S Taghavirad, I R Tomalin, and O Zorba. The cms tracker front-end driver. 2003.

[67] Karl Aaron Gill, F Drouhin, R Grabit, L Gross, K Kloukinas, C Ljuslin, G Magazz, A Marchioro, E Murer, E Noah, C Paillard, J Troska, F Vasey, and D Vintache. Progress on the cms tracker control system. 2005.

[68] C Paillard, C Ljuslin, and A Marchioro. The ccu25: a network oriented communication and control unit integrated circuit in a 0.25 μm cmos technology. 2002.

[69] P. Placidi, A. Marchioro, P. Moreira, and K. Kloukinas. A 40-MHz clock and trigger recovery circuit for the CMS tracker fabricated in a 0.25-mu-m CMOS technology and using a self calibration technique. Prepared for 5th Workshop on Electronics for the LHC Experiments (LEB 99), Snowmass, Colorado, 20-24 Sep 1999.

[70] G. Magazzu, A. Marchioro, and P. Moreira. The detector control unit: An ASIC for the monitoring of the CMS silicon tracker. *IEEE Trans. Nucl. Sci.*, 51:1333–1336, 2004.

[71] F Benotto, R Ciaranfi, R D'Alessandro, N Demaria, G DeRobertis, and S Mersi. Design and test of the digital opto hybrid module for the cms tracker inner barrel and disks. Technical Report CMS-NOTE-2008-013. CERN-CMS-NOTE-2008-013, CERN, Geneva, Apr 2008.

[72] U Goerlach. Industrial production of front-end hybrids for the cms silicon tracker. 2003.

[73] Anthony A. Affolder et al. Test of CMS tracker silicon detector modules with the ARC system. *Nucl. Instrum. Meth.*, A535:374–378, 2004.

[74] R.Frühwirth. Application of Kalman Filtering to track and vertex fitting. *Nucl.Instr.Meth*, A262, 1987.

[75] A. Frohlich, H. Grote, C. Onions, and F. Ranjard. Marc: Track finding in the split field magnet facility. CERN-DD-76-5.

[76] Carsten Noeding. Track reconstruction and experience with cosmic ray data in cms. Technical Report CMS-CR-2008-006. CERN-CMS-CR-2008-006, CERN, Geneva, Dec 2007.

[77] S Cucciarelli, M Konecki, Danek Kotlinski, and T Todorov. Track-parameter evaluation and primary-vertex finding with the pixel detector. Technical Report CMS-NOTE-2003-026, CERN, Geneva, Sep 2003.

[78] G. W. Foster, J. Freeman, C. Newman-Holmes, and J. Patrick. A Fast Hardware Track Finder for the CDF Central Tracking Chamber. *Nucl. Instrum. Meth.*, A269:93, 1988.

[79] R. Frühwirth, M. Regler, R.K. Bock, H. Grote, and D. Notz. *Data Analysis Techniques for High-Energy Physics, Second Edition.* Cambridge Monographs on particle physics, nuclear physics and cosmology, Cambridge, UK.

[80] Rudolf Frühwirth. Matching of broken random samples with a recurrent neural network. *Nucl. Instrum. Meth.*, A356:493–497, 1995.

[81] John. Wolberg. *Data Analysis using the method of Least Squares.* Springer, Berlin, 2006.

[82] M. Regler. Influence of computation algorithms on experimental design. (talk). *Comput. Phys. Commun.*, 22:167–175, 1981.

[83] M.G. Kendall and A. Stuart. *The Advanced Theory of Statistics, Volume II.* Charles Griffin and Company Limited, London.

[84] M. Regler. Multiple scattering in the fitting procedure. (in german). *Acta Phys. Austriaca*, 49:37–45, 1978.

[85] M. Crozon et al. DELPHI, a Detector with Lepton, Photon and Hadron identification: Letter of Intent for an experimental program at LEP. DELPHI-82-1.

[86] Pierre Billoir. Track fitting with multiple scattering: A new method. *Nucl. Instr. Meth.*, A225:352, 1984.

[87] Pierre Billoir, R. Fruhwirth, and M. Regler. Track element merging strategy and vertex fitting in complex modular detectors. *Nucl. Instrum. Meth.*, A241:115–131, 1985.

[88] A. Gelb. *Applied Optimal Estimation.* MIT press, Cambridge, Massachusetts.

[89] A. Strandlie and W. Wittek. Propagation of covariance matrices of track parameters in homogeneous magnetic fields in cms. CERN-CMS-NOTE-2006-001.

[90] Pierre Billoir. Progressive track recognition with a kalman like fitting procedure. *Comput. Phys. Commun.*, 57:390–394, 1989.

[91] Pierre Billoir and S. Qian. Simultaneous pattern recognition and track fitting by the kalman filtering method. *Nucl. Instrum. Meth.*, A294:219–228, 1990.

[92] G. Kitagawa. Non-gaussian seasonal adjustment. *Computer and Mathematics with Applications*, 18:503, 1989.

[93] Rudolf Frühwirth. Track fitting with non-gaussian noise. *Computer Physics Communications*, 100:1–16, 1997.

[94] COBRA: Coherent Object-oriented Base Reconstruction and Analysis.
Site located at http://cobra.web.cern.ch/cobra.

[95] ORCA: CMS Reconstruction Package.
Site located at http://cmsdoc.cern.ch/orca.

[96] CMSSW: CMS SoftWare.
Site located at http://cmsdoc.cern.ch/cms/cpt/Software/html/General/.

[97] S Cucciarelli, Danek Kotlinski, and T Todorov. Position determination of pixel hits. Technical Report CMS-NOTE-2002-049, CERN, Geneva, Nov 2002.

[98] G. Cerati. Outlier rejection during the final track fit. Technical Report CMS-IN-2008-007, CERN, Geneva, Feb 2008.

[99] C. Amsler et al. Review of particle physics. *Phys. Lett.*, B667:1, 2008.

[100] S. Baffioni et al. Electron reconstruction in CMS. *Eur. Phys. J.*, C49:1099–1116, 2007.

[101] T Speer, W Adam, R Frhwirth, A Strandlie, T Todorov, and M Winkler. Track reconstruction in the cms tracker. oai:cds.cern.ch:884424. Technical Report CMS-CR-2005-014. CERN-CMS-CR-2005-014, CERN, Geneva, Jul 2005.

[102] W. Adam et al. Track reconstruction with cosmic ray data at the tracker integration facility. Technical Report CMS-NOTE-2009-003. CERN-CMS-NOTE-2009-003, CERN, Geneva, Jan 2009.

[103] W. Adam et al. Silicon strip tracker detector performance with cosmic ray data at the tracker integration facility. Technical Report CMS-NOTE-2008-032. CERN-CMS-NOTE-2008-032, CERN, Geneva, Mar 2008.

[104] W. Adam et al. Cms tracker alignment at the integration facility. Technical Report CMS-NOTE-2009-002. CERN-CMS-NOTE-2009-002, CERN, Geneva, Jan 2009.

[105] A. Rizzi. Search and reconstruction of heavy stable charged particles with the cms experiment at lhc. 2007. CERN Theses.

[106] A Giammanco. Particle identification with energy loss in the cms silicon strip tracker. Technical Report CMS-NOTE-2008-005. CERN-CMS-NOTE-2008-005, CERN, Geneva, Jun 2007.

[107] D.G. Bettis. Efficient embedded runge-kutta methods. *Numerical Treatment of Differential Equations*, 631:9–18, 1978.

[108] J. R. Cash and Alan H. Karp. A variable order runge-kutta method for initial value problems with rapidly varying right-hand sides. *ACM Trans. Math. Softw.*, 16(3):201–222, 1990.

[109] The CMS collaboration. Alignment of the cms silicon vertex tracking detector with cosmic ray particles. CERN-CMS-NOTE-2009-xxx (in preparation).

Acknowlegments

I would like to thank Prof. Lorenzo Foà for his guidance and for the constant dedication in following the progress of my work. I am very grateful to Dr. Theodore Todorov who supervised my initial work at CERN and gave invaluable advice. His help was fundamental in defining the scope of my research.

I thank Dr. Wolfgang Adam for the instructive conversations about tracking and also Prof. Shahram Rahatlou, Dr. Tommaso Boccali and Dr. Andrea Rizzi for their assistance in understanding many details of the CMS software.

I also thank Dr. Fabrizio Palla and Prof. Guido Tonelli for having motivated me during the construction of the CMS Tracker. I now fully recognize the value of all those days (holidays included) spent in the INFN laboratory of San Piero.

I am very grateful to Prof. Vivek Sharma for having created the conditions to complete the writing of the thesis in the comfortable environment of San Diego. I'm also grateful to graduate students Ryan Kelley and Matthew LeBourgeois for having read the draft of the thesis and helped me in improving the quality of the text.

This work was made possible by Scuola Normale Superiore and INFN that supported me both in Pisa and at CERN during the three years of the doctoral program.