Diss. ETH Nr. 18232

CMS Pixel Module Qualification
and
Search for $B_s^0 \to \mu^+\mu^-$

A dissertation submitted to

ETH Zurich

for the degree of

Doctor of Science

presented by

Christina Eggel
Dipl. Phys. ETH
born on November 4, 1978
citizen of Naters VS

accepted on recommendation of
Prof. Dr. Urs Langenegger, examiner
Prof. Dr. Felicitas Pauss, co-examiner

2009

In loving memory of my father, Richard Eggel.

"Everything can change at any moment, suddenly and forever."

Paul Auster

Abstract

This work is divided into two parts. The first part focuses on the qualification and performance optimisation of the barrel pixel detector hardware in the CMS experiment. The barrel part of the pixel system holds 768 detector modules on three cylindrical layers around the interaction region. In the period of April 2006 to March 2008, 971 fully assembled detector modules were produced at the Paul Scherrer Institute (PSI) in Villigen, Switzerland. To meet the performance and lifetime requirements of the CMS pixel system, an elaborate test procedure was developed and a grading system was established. From the total of 813 modules that qualified for use in the CMS pixel system, 75 % were attested an excellent performance and 25 % held deficiencies with only minor impact on the detector performance. The remaining 158 modules exhibited serious flaws and were rejected.

The second part of this work is on the development of a physics analysis aiming at the measurement of the rare decay $B_s^0 \rightarrow \mu^+\mu^-$. This decay mode provides good sensitivity to $\tan\beta$, a central parameter of the Minimal Supersymmetric extension of the Standard Model (MSSM). The MSSM and many other Standard Model extensions predict a (very) large increase of the branching fraction expected in the Standard Model. This rare decay mode therefore offers an excellent opportunity to explore new physics—possible with only a small data sample gathered from the very first running period at the LHC. With the first $1\,\mathrm{fb}^{-1}$ of integrated luminosity, an upper limit on the branching fraction of 1.3×10^{-8} at the 90% C.L. is expected. In this analysis, the most effective selection criteria to discriminate the large number of background events from the signal events are based on the long lifetime of the B mesons. The pixel detector allows a precise determination of the displaced vertices and will also therefore play a crucial role in this part of thesis.

Zusammenfassung

Im ersten Teil dieser Arbeit wird die Qualifizierung und Optimierung der Module für den Einsatz im zentralen CMS Pixel Detektor beschrieben. Dieser besteht aus 768 Modulen, die auf drei zylindrischen Lagen um den Kollisionspunkt befestigt sind. Von April 2006 bis März 2008 wurden am Paul Scherrer Institut (PSI) in Villigen über 900 solcher Module zusammengebaut. Um die Anforderungen bezüglich Leistung und Lebensdauer zu erfüllen, wurde eine umfangreiche Testprozedur entwickelt, in der jedes Modul eingehend geprüft und eingestuft wurde. Als Teil der Testprozedur wurden unter anderem die charakteristischen Eigenschaften von jedem Modul bestimmt, und verschiedene Kalibrationsalgorithmen sowie Algorithmen zur Leistungsoptimierung durchgeführt. Entsprechend den Testergebnissen wurde jedes Modul in eine von drei möglichen Qualitätskategorien eingeteilt. Am Ende qualifizierten sich 813 der insgesamt 971 Module für den Einsatz im CMS Pixel Detektor. In 75 % der Fälle wiesen diese Module sogar eine ausgezeichnete Qualität auf.

Der zweite Teil der Arbeit präsentiert eine Monte Carlo Studie zur Messung des seltenen Zerfalls $B_s^0 \to \mu^+\mu^-$. Dieser Zerfallskanal bietet eine exzellente Möglichkeit, um nach neuer Physik zu suchen. In vielen Erweiterungen des Standard Modells, wie zum Beispiel in der Minimalen Supersymmetrischen Erweiterung des Standard Modells (MSSM), vergrössert sich das Verzweigungsverhältnis erheblich in Abhängigkeit von $\tan\beta$—einem zentralen Parameter vom MSSM. Schon mit relativ kleinen Datenmengen lässt sich der Parameter Raum der Modelle jenseits des Standard Modells eingrenzen. Bei einer integrierten Luminosität von $1\,\mathrm{fb}^{-1}$ wird eine obere Grenze von 1.3×10^{-8} mit 90% C.L. für das Verzweigungsverhältnis erwartet.

Contents

Introduction

Nearly everything we currently know about the elementary constituents of matter and their interactions can be described by a relativistic quantum field theory known as the Standard Model of particle physics. In the past decades, the predictions of the Standard Model have been confirmed to a high precision in a wide variety of experiments—making it one of the most stringently tested scientific theories. The only unobserved particle of the Standard Model is the elusive Higgs boson. The existence of this scalar particle is required by the Higgs mechanism, which was introduced ad hoc to explain how the gauge bosons of the weak force acquire their masses through spontaneous symmetry breaking. Despite its stunning success in describing the experimental data, the Standard Model has its deficiencies: in its original design, neutrinos had been assumed to be massless, but the neutrino oscillations—first observed in 1998—required an adjustment of the Standard Model to accommodate massive neutrinos. Moreover, a quantum description of gravity is not included in the Standard Model. The mass scale where gravitational effects become important, and the mass scale of the electroweak force are highly disparate, which leads to the so-called hierarchy problem. The electroweak scale is sensitive to higher energy scales, where quadratically divergent quantum corrections to the Higgs mass arise. To cancel the lowest order contributions, an unnatural fine-tuning of parameters is required. In addition, the particles of the Standard Model merely account for four percent of the energy density in our universe; the rest is made of mysterious dark matter and dark energy. The Standard Model also fails to explain the matter-antimatter asymmetry in the present universe.

Given the many shortcomings, the Standard Model is generally considered to be only a low-energy effective theory and new physics is expected to enter at the TeV scale. A wide variety of theoretical approaches beyond the Standard Model have been proposed. The new hadron accelerator facility—the Large Hadron Collider (LHC)—in Geneva will play a significant role in providing evidence for physics beyond the Standard Model. At unprecedented energies and interaction rates, the LHC will open the door to a new energy regime, putting the Standard Model and its extensions to the test. In three experimental areas, counter-rotating beams of protons will collide at a center-of-mass energy of 14 TeV. A fourth experiment will look at heavy ions collisions at a center-of-mass energy of 1148 TeV (2.76 TeV per nucleon).

Chapter 1 reviews the design and operation of the LHC and provides an outlook on possible future colliders at the high-energy frontier. Chapter 2 is dedicated to one of LHC's general purpose experiment—the Compact Muon Solenoid (CMS). Given the topic of this work, the emphasis of this chapter is on the Silicon tracker, the innermost device of CMS. Chapter 3 focuses on the barrel part of the pixel detector and its basic component—the CMS barrel pixel module. In an elaborate procedure comprising of all functionality, calibration and performance tests, the quality of each module was assessed. The algorithms used in the test procedure are explained in Chapter 4, including a summary of the test results. Chapter 5 describes the different steps of the qualification procedure and concludes the first part of the thesis with the results of the module qualification. The second part of this work starts by outlining the basic concepts of the Standard Model in Chapter 6, with emphasis on flavour mixing and CP violation. The major goals of B physics and the prospects at the LHC are also highlighted. Chapter 7 provides a detailed Monte Carlo study of the physics analysis aimed at the measurement of the rare decay $B_s^0 \to \mu^+\mu^-$. This analysis chapter concludes with the expected upper limit for the branching fraction in $1\,\mathrm{fb}^{-1}$ of integrated luminosity.

Chapter 1

The Large Hadron Collider at CERN

The Large Hadron Collider (LHC) is a new hadron accelerator facility at the European Organisation for Nuclear Research (CERN) near Geneva. The LHC aims to explore physics beyond the standard model by colliding protons into protons at unprecedented high energies and interaction rates [1; 2]. It is designed to collide proton beams at a center-of-mass energy of 14 TeV and a nominal luminosity of $\mathcal{L} = 10^{34}\,\mathrm{cm^{-2}\,s^{-1}}$. The resulting event rate of approximately 10^9 inelastic interactions per second is achieved by colliding bunches of approximately $1.15 \cdot 10^{11}$ protons every 25 ns. The highly complex and challenging two-ring accelerator was installed in the existing 26.7 km tunnel of the Large Electron-Positron collider (LEP) and reuses the existing proton accelerator facilities of CERN as injectors. Proton beams were successfully circulated at the injection energy of 450 GeV in both directions of the LHC for the first time on September 10, 2008. Due to an incident, that occurred nine days later and the time required to repair the resulting damage, the restart of LHC is scheduled for September 2009, with the first particle collisions following in late October 2009.

The length scale probed in a collision experiment is given by the de Broglie wavelength $\lambda = h/p$ and decreases with the momentum of the colliding protons. In the energy regime of the LHC, the constituents of the incident protons—quarks and gluons—can interact directly with each other. As each of the constituents only carries a fraction of the total energy of the proton, a wide spectrum of effective collision energies is available. This enables the LHC to be a powerful discovery machine with a very high mass reach for direct discovery of new particles. Although several precision measurements are also possible with the LHC itself, electron-positron colliders are much better suited for that purpose due to the very clean experimental environment and the known collision energy of the point-like particles. Advanced research on new types of linear electron-positron accelerators, that complement the capabilities of the LHC, is being conducted at present [3; 4].

In addition to proton-proton collisions, shorter runs with completely ionised lead nuclei (Pb^{82+}) are planned before each annual machine shutdown. With the nominal dipole field strength, the center of mass energy will be 1148 TeV (2.76 TeV per nucleon). Bunches containing 7×10^7 nuclei will collide every 100 ns reaching a design luminosity of $\mathcal{L} = 10^{27}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

The LHC envisages a very rich and diverse physics program that will be covered by the six different experiments located in the four experimental caverns: the two general purpose experiments, CMS [5; 6] and ATLAS [7; 8], will elucidate the nature of electroweak symmetry breaking and search for its agent—the Higgs boson. CMS and ATLAS will also conduct b-physics studies. The main part of the LHC b-physics program, however, will be covered by another dedicated experiment, LHCb [9; 10]. The fourth experiment, ALICE [11; 12], has been conceived for heavy ion physics and will study the properties of quark-gluon plasma. The TOTEM experiment [13] is designed to study phenomena in the very forward region, including elastic and diffractive scattering and will provide a precise measurement of the total cross-section σ_{pp}. LHCf [14; 15] is a special purpose detector that will measure the production cross-section of neutral pions emitted in the very forward direction of proton-proton collisions, providing the input for models used in ultra-high energy cosmic ray studies.

The following sections serve as an introduction to the LHC: section 1.1 outlines the design and operation of the LHC; section 1.2 discusses the machine luminosity and the staged upgrade scenario of the LHC and its injector chain; section 1.3 describes the properties of proton-proton collisions at the LHC. Section 1.4 concludes the chapter by providing an outlook on possible future colliders at the high-energy frontier.

1.1 The LHC Design

Despite the inherent design advantages of a $p\bar{p}$ collider allowing two beams with opposite charge to use the same vacuum and magnet system, the requirement of high beam intensities and the difficulty to produce sufficient amounts of antiprotons preclude any obtainable advantage. This choice does not effect the physics potential of the LHC, as most of the interactions are gluon-initiated [16] and the distributions of gluons in protons and antiprotons are the same. Cost saving reasons and, of overriding importance [2], the lack of space in the LHC tunnel resulted in the adoption of a twin bore magnet design, where both beam pipes and superconducting coils are combined within the same mechanical structure and cryostat. The maximum beam energy at LHC is limited by the peak dipole field that can be achieved with the dipole magnets. The envisaged energy of $E = 7\,\mathrm{TeV}$ for each proton beam requires a magnetic field of 8.33 T, following from

$$E[\,\mathrm{TeV}] = B[\,\mathrm{T}] \times 0.84\,\frac{\mathrm{TeV}}{\mathrm{T}}. \tag{1.1}$$

Based on the layout of the LEP tunnel, the LHC has eight arcs and eight straight sectors. An arc contains 23 regular cells with six main dipole magnets bending the beam and two main quadrupole magnets focusing the beam, as well as various multipole corrector magnets. Each straight region either serves as an experimental or a utility insertion: four are dedicated to the experiments, one to the radio-frequency (RF) system, two to beam cleaning and one to beam dumping. Dispersion suppressors (DS) are located at the transitions between the arc and straight sections to adapt the LHC reference orbit to the geometry of the tunnel, to correct horizontal dispersion and to help matching insertion optics. In total, the LHC magnet system contains over 9000 magnets. These superconducting magnets are generally operated at 1.9 K in a static bath of superfluid helium, the exceptions being a few that are operated at 4.5 K. The 1232 main dipoles are bent in their horizontal plane. At 1.9 K the bending angle is 5.1 mrad per dipole and the magnetic length of a dipole is 14.3 m. This corresponds to a radius of curvature of 2804 m.

The LHC will be supplied with protons from the existing complex of proton accelerators at CERN (see Figure 1.1). These accelerators having been in use for decades for other experiments were extensively upgraded in order to adapt them to the requirements of the LHC. The acceleration starts in the duoplasmatron proton source of the linear accelerator facility (LINAC2). In LINAC2, the protons are accelerated to 50 MeV before they are injected into the Proton Synchrotron Booster (PSB). In order to reduce space charge, the 1.4 GeV protons are transferred in two batches from the four PSB rings to the Proton Source (PS). The PS accelerates the protons to an energy of 25 GeV and forms the bunch train with the correct LHC spacing of 25 ns. At the last stage of the injector chain, the Super Proton Synchrotron (SPS), the protons are accelerated to the LHC injection energy of 450 GeV. The two counter-rotating proton beams are delivered to the LHC through two separate transfer lines (Tl 2 and Tl 8). The rise times of the kicker magnets at the different injection stages lead to missing bunches in the LHC beam structure. Including the 3 μs gap foreseen for the rise time of the beam dumping magnets, 2808 out of 3654 available bunches will be filled. Filling one LHC ring takes approximately three minutes.

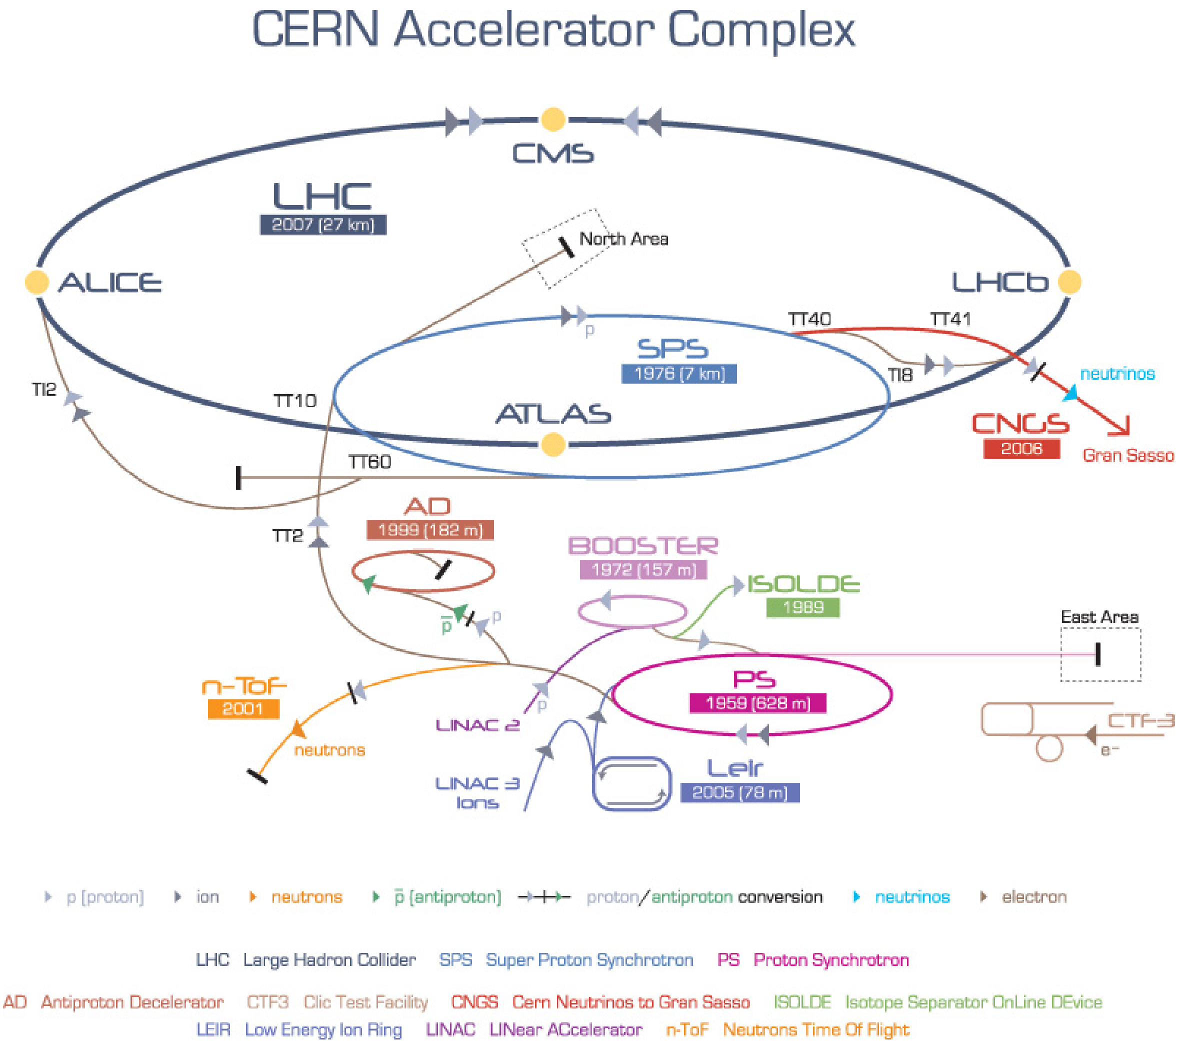

Figure 1.1: LHC accelerator and injection chain at CERN and the location of the four main experimental insertions (utility insertions are not shown).

The radio-frequency (RF) acceleration systems and the beam feedback systems are located at the insertion region (IR) at Point 4 (between ALICE and CMS). The main 400 MHz RF Acceleration System (ACS) captures, accelerates and stores the injected beam, and at the same time damps the longitudinal injection errors. There are two independent ACS systems for each beam, both containing eight Superconducting Cavities (SC) operated at 4.5 K. Driven by 300 kW klystrons that are connected to an SC through a variable power coupler, each cavity can provide a tunable acceleration voltage of 1 MV at injection to 2 MV at nominal energy. The latter corresponds to a field strength of 5.5 MV/m. Each turn, the beam energy is increased by 485 keV, giving a ramp-up time of 20 minutes. The transverse damping and feedback system (ADT) uses electrostatic deflectors that damp the transverse oscillations.

1.2 The Machine Luminosity and Upgrades

The event rate of a certain physics process is given by

$$R = \mathcal{L}\sigma \tag{1.2}$$

where σ is the cross-section of the physics process under study and $\mathcal{L}$ is the machine luminosity. The machine luminosity depends only on the beam parameters and can be written as

$$\mathcal{L} = \frac{n_b N_b^2 \gamma_r f}{4\pi\varepsilon_n\beta^*} F \tag{1.3}$$

where N_b is the number of particles per bunch, n_b is the number of bunches per beam, γ_r is the relativistic gamma factor, f is the bunch crossing frequency, ε_n is the normalised transverse beam emittance, β^* is the beta function at the collision point and F is the geometric luminosity reduction factor due to the crossing angle at the interaction point. F can be calculated from the crossing angle θ_c, the RMS bunch length σ_z and the transverse RMS beam size σ^*: $F = 1/\sqrt{1 + (\theta_c\sigma_z/2\sigma^*)^2}$.

The existing complex of injectors at CERN can provide the beam for reaching the nominal and "ultimate" luminosity of $10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and $2.3 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ respectively. The proton accelerators in the injection chain, however, were built decades ago and are not optimised for the purpose of LHC. Also, there is call for a luminosity upgrade within less than ten years of operation due to the radiation damage limit for the low-β quadrupoles in the IR and the evolution of the statistical error halving time.

In terms of upgrades, a staged upgrade of the LHC and its injectors to achieve a luminosity of $10^{35}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ has been proposed [17] and is still under study [18; 19]. Phase I of a planned machine upgrade has been approved by CERN and aims at enabling reliable operation at a luminosity of $2 - 3 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. The major improvements during this stage will be the replacement of LINAC2 with a 160 MeV H^- proton linac (LINAC4) and the upgrade of the low-β triplets in the IR. Phase II, or

Super LHC (SLHC), is still under discussion and aims at reaching a peak luminosity of $10^{35}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. This tenfold increase in luminosity with respect to nominal performance extends the discovery range for new particles by 20 to 30 % in mass [20]. The upgrades for SLHC include major modifications in the injector chain and the LHC insertion regions. The large increase in event rate also requires significant changes to the experiments in the two high luminosity insertion regions, CMS and ATLAS [20]. The enormous particle rates present at the SLHC—both in terms of instantaneous and integrated rates—are particularly challenging for the detector elements close to the interaction point (such as the CMS pixel detector).

1.3 Proton-proton Collisions at the LHC

The total cross-section of proton-proton interactions at the LHC energy of 14 TeV can be extrapolated from previous experiments at lower energies [21] or extracted from cosmic ray data [13; 22]. The total cross-section σ_{pp} has contributions from elastic and inelastic scattering. The inelastic processes can be subdivided into diffractive and non-diffractive scattering, and therefore $\sigma_{pp} = \sigma_{\mathrm{el}} + \sigma_{\mathrm{di}} + \sigma_{\mathrm{nd}}$.

In the *elastic* process, two protons are only slightly deflected, interacting mostly via photon exchange (Coulomb scattering) at low four-momentum transfer or predominantly via Pomeron exchange in the region of high momentum transfer. The dominant contributions in *diffractive* processes come from single and double diffractive dissociation, in which the exchange of a colourless Pomeron leads to the fragmentation of one or both protons respectively, giving rise to hadronic activity at large pseudorapidities on one or both sides of the detector. Most of the *non-diffractive* inelastic interactions are soft and happen at low four-momentum transfer. They are often referred to as *minimum bias events*, indicating a minimal bias in the online selection that is defined by the minimum bias trigger. In literature, different definitions of minimum bias interactions exist; historically, the double diffractive inelastic processes are also included[1] and minimum bias events therefore approximately comprise the non-single diffractive inelastic (NSD) interactions.

The prediction of the total *pp* cross-section depends on the model used for the extrapolation. A discrepancy between the two final results from Tevatron and the

[1] from experiments where the minimum bias trigger is based on a a two-arm coincidence that suppresses single diffractive events

large uncertainties in the cosmic ray data leaves a broad interval for the expected value, typically ranging from 90 to 130 mb depending on the model. The large uncertainty will be resolved by the precise measurement of σ_{pp} by the TOTEM experiment. The cross-section estimate given by the Monte Carlo (MC) event generator PYTHIA amounts to about 101 mb, of which 22 mb come from elastic scattering and 55 mb is due to non-diffractive inelastic interactions.

At the initial low luminosity there will be 3.5 non-diffractive inelastic interaction per bunch crossing on average. At high luminosity conditions this number will increase to an average of about 20 events per bunch crossing (implying around 1000 charged particles). These minimum bias interactions in addition to the event under study are termed *pile-up events*. Most of these events are *soft*, which means they happen at low four-momentum transfer Q^2. In *hard* scattering processes the interaction takes place between the constituent partons of the protons (quarks and gluons). Soft partonic interactions in the remnants of the proton that are not associated with the hard scattering processes, contribute to the so-called *underlying event* (UE). The production rate and event properties of hard interactions can be estimated with good precision using perturbative QCD [16]. At LHC energies, the partons involved in the interaction carry a small momentum fraction x. The predominant processes at LHC are therefore sea quark and gluon scatterings (as opposed to Tevatron, where valence quark scatterings prevail).

1.4 The High Energy Frontier

The "Livingston Plot" in Figure 1.2 exhibits the immense exponential growth in the constituent energy reach of lepton and hadron colliders during the last decades. This fast ascent owes itself mainly to the progress in accelerator technology—in particular superconducting magnet technology. Simultaneously, the plot indicates a much slower progression of the energy frontier at which new physics can be probed by future colliders.[1]

A major drawback of circular electron-positron colliders is the energy loss due to synchrotron radiation. This energy loss has been the limiting factor for the center-of-mass energy at the Large Electron-Positron Collider (LEP) at CERN.

[1] Besides the energy of an accelerator (which defines the threshold for new discoveries), machine luminosity is equally important, as it determines the interaction rate and hence the probability of new discoveries.

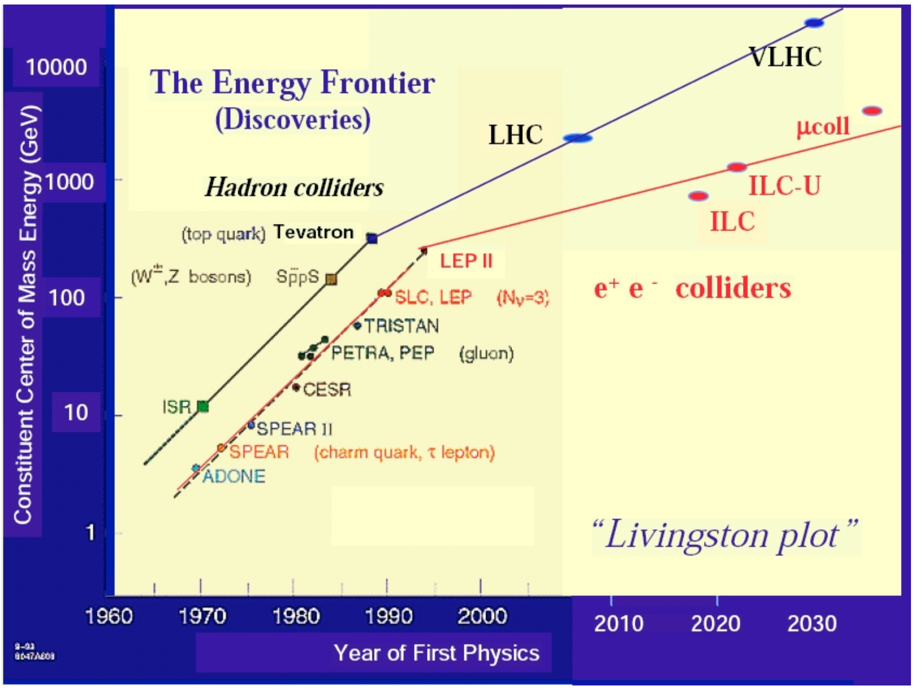

Figure 1.2: The Livingston plot showing the constituent energy reach of past, present and future colliders.

For the highly relativistic particles in an accelerator

$$\beta = \frac{v}{c} \approx 1 \qquad \text{and} \qquad \gamma = \frac{E}{mc^2} \tag{1.4}$$

and the energy loss due to synchrotron radiation per revolution is given by [23]:

$$\Delta E = \frac{4\pi\alpha}{3R}\beta^4\gamma^4 \sim \frac{E^4}{Rm^4}, \tag{1.5}$$

where R is the accelerator radius and α is the electromagnetic fine-structure constant. One option to reduce the energy losses in circular lepton-lepton colliders is to increase the radius. Building a circular electron-positron collider beyond LEP energies, however, would result in unwarranted costs incurred due to the sheer size of the machine and its power consumption. On the other hand, synchrotron radiation losses vanish in case of a linear collider, where $R \to \infty$. At present two new linear accelerators, aiming to collide electrons into positrons at the TeV scale, are in development: the International Linear Collider (ILC) [3] and the Compact Linear Collider (CLIC) [4]. Another option to realise a circular lepton-lepton collider is to accelerate muons instead of electrons. Muons are 200 times heavier than electrons and thus synchrotron radiation becomes negligible. This particular approach is currently studied by the Muon Collider Task Force (MCTF), proposing muon collisions at a center-of-mass energies above 1 TeV [24]. The rest lifetime $\tau_0 = 2\,\mu$s of such highly relativistic muons is stretched by a Lorentz factor γ of the order of 10^5. Similarly, synchrotron radiation does not pose a serious problem at the LHC. A proton is approximately 1800 times heavier than an electron which reduces the energy loss by a factor 10^{13}. The LHC requires, however, very strong magnetic field of 8.33 T to keep the protons on a circular track. The proposed Very Large Hadron Collider (VLHC) [25] foresees a staged construction of a 233 km storage ring, increasing the dipole field of the bending magnet from 2 T to 10 T and a center-of-mass energy of 40 TeV and 175 TeV respectively.

To date, it remains unclear whether any of the above accelerators will reach realisation.

Chapter 2

The Compact Muon Solenoid

The Compact Muon Solenoid (CMS) is one of the two general purpose detectors, designed to explore a wide range of physics in the 14 TeV proton-proton collisions at LHC. The detector requirements are driven by the aim to measure the decay signatures of the hypothetical Higgs boson and the search for new physics at the TeV scale. The strong magnetic field generated by a super-conducting solenoid allows for a simple and relatively compact design. CMS is composed of several sub-detector systems, arranged cylindrically around the beam-pipe as illustrated in Figure 2.1. Closest to the collision region, a silicon tracking device determines the trajectories of charged particles. This tracker is surrounded by the electromagnetic and hadronic calorimeters that measure the energy of charged and neutral particles. Except for the hadron forward calorimeter, the tracker and other calorimeters are contained within the super-conducting coil. Outside the magnet, the muon system is interleaved with the iron plates of the flux-return yoke of the solenoid.

Before outlining the design and operation of the different detector parts in this chapter, the CMS coordinate conventions are introduced in section 2.1. The subsequent sections describe the different detector parts and sub-detector systems of CMS: the solenoid is discussed in section 2.2, the silicon tracker in section 2.3, the electromagnetic calorimeter in section 2.4, the hadronic calorimeter in section 2.5 and the muon system in section 2.6. According to the topic of this thesis, the emphasis has been put on the tracking device and a description of the main algorithm, used to reconstruct the tracks of charged particles in the CMS tracker, can be found at the end of this chapter, in section 2.7.

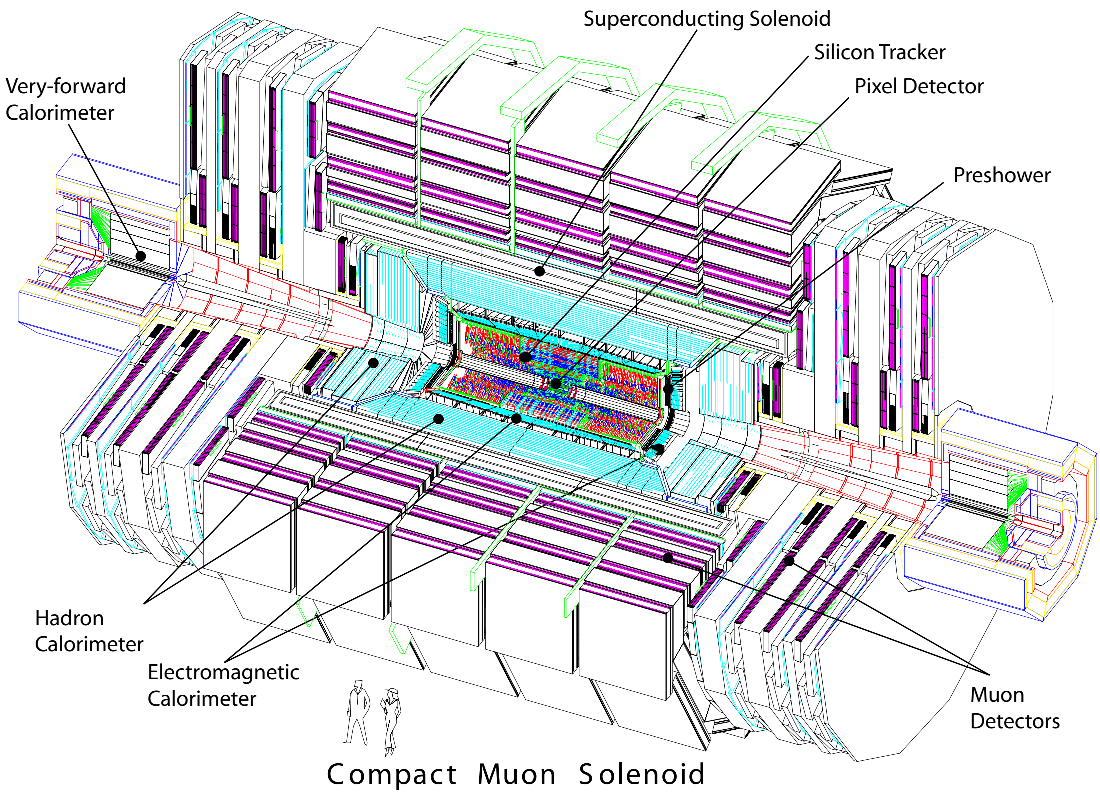

Figure 2.1: The detector layout of the Compact Muon Solenoid (CMS).

2.1 Coordinate conventions

- (x, y, z): The z-axis is placed along the beam direction toward the Jura mountain, and the x- and y-axis define the transverse plane perpendicular to the beam. x points toward the center of LHC and y points vertically upward.
- (r, θ, ϕ): The azimuthal angle ϕ is measured from the x-axis in the x-y-plane and the polar angle θ is measured from the z-axis. r is the radial distance.

Since the actual interaction in the collision at the LHC happens between the constituents of the protons, the rest frame of the hard collision will be boosted along the beam axis. To study the proton collisions in a coordinate frame that is invariant under Lorentz boosts along the beam axis, the polar angle θ is commonly replaced by the pseudorapidity η, defined as

$$\eta = -\ln\left(\tan\frac{\theta}{2}\right). \tag{2.1}$$

In (r, η, ϕ) coordinates, the transverse quantities as well as the differences in η are Lorentz invariant under longitudinal boosts.

2.2 Solenoid

The design of the CMS solenoid [5; 26] is driven by the large bending power needed to precisely measure the momenta of charged particles. An innovative design, featuring amongst other things four winding layers of reinforced self-supporting conductors, was necessary to build the CMS super-conducting magnet. The coil has a diameter of 6.3 m and a length of 12.5 m. The large dimensions and a strong magnetic field of 4 T distinguish the solenoid notably from previous experiments. A high bending power is crucial for unambiguous charge determination and good momentum resolution (and thus sharp trigger thresholds): since the momentum is determined by measuring the sagitta of the particle track, the momentum resolution is proportional[1] to $1/BL^2$, where B is the magnetic field and L is the distance between the inner- and outermost measurement layer in the tracker. At the same time, the strong field increases the occupancy at low radii. Therefore, a highly granular device, such as a pixel detector, is required in the region closest to the interaction point. The calorimeter system benefits from the strong magnetic field, since most of the minimum bias events are confined to low radii and therefore the trapping of particles in the barrel region is reduced.

2.3 The Silicon Tracker

In the dense charged particle environment of the interaction region, the CMS silicon tracker [5; 27] plays an essential role, ensuring an efficient and ghost free track reconstruction with CMS. An important aspect of the CMS tracker is the measurement of tracks close to the interaction point, allowing precise determination of the secondary vertices of long-lived objects and distinguishing them from the large background of light quarks and gluon jets.

To ensure efficient reconstruction and correct bunch crossing allocation of the charged particle trajectories, a high granularity and a fast response are essential. A highly granular device also implies a large amount of material needed for support, cooling, electronics and cabling. Stringent constraints on the material budget of the tracker

[1] if multiple scattering effects are not considered

(see Figure 2.2) are, however, imposed by multiple scattering, bremsstrahlung, photon conversions and nuclear interactions. This leads to an inevitable compromise limiting the number of active layers and the choice of materials used. A radiation hard design is furthermore imperative, considering the high charged particle fluxes in the vicinity of the interaction point. Based on these considerations, a tracker device entirely based on silicon detector technology was selected. Its specification are summarised in Table 2.1 at the end of this section.

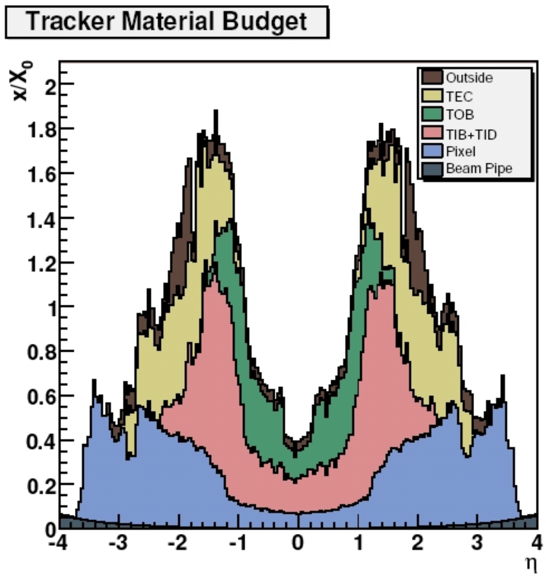

Figure 2.2: Tracker material budget for the different sub-detectors as a function of η in units of radiation length [5].

The tracker along with the muon system will also allow muon pair invariant masses to be reconstructed in heavy ion collisions, which is useful in studying quark-gluon plasma physics.

At the design luminosity of $10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, an average of 20 overlapping proton-proton interactions per bunch crossing are expected, producing approximately 1000 charged particles. In addition, the strong magnetic field confines the low $p_{\perp}$ charged particles to helical trajectories with small radii. Together with the steeply falling $p_{\perp}$ spectrum of minimum events (see Figure 2.3), this leads to a charged particle density that rapidly decreases with the radius. This decrease deviates from the $1/r^2$ law due to the presence of the 4 T magnetic field. In order to keep the occupancy below the level of a few percent, the architecture of the tracker is determined by the three particle flux regimes present at high luminosity (see Figure 2.4).

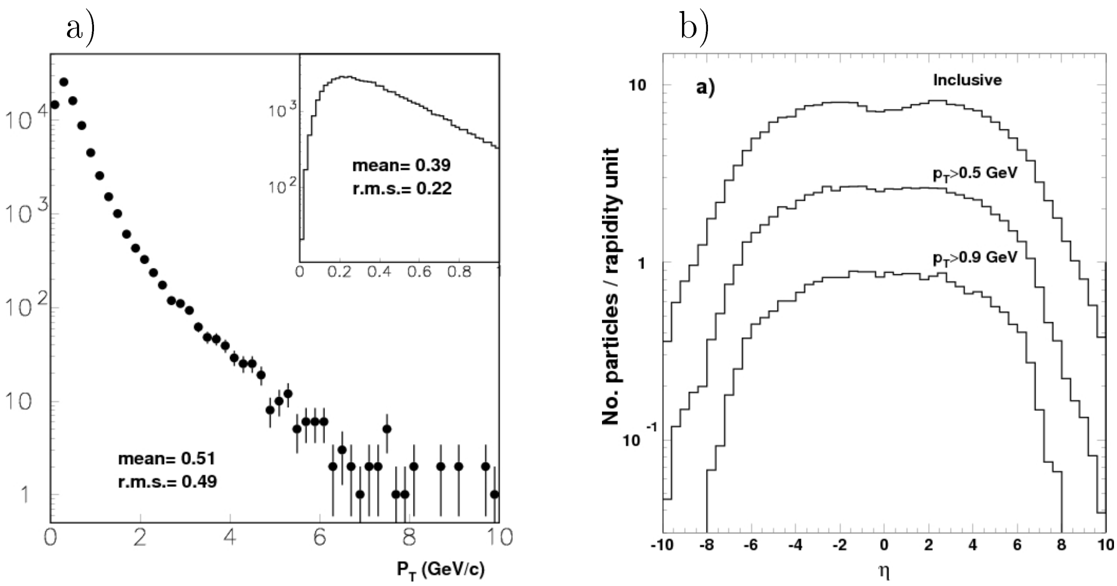

Figure 2.3: Distributions of charged particles in minimum bias events [27]: a) transverse momentum, b) pseudorapidity.

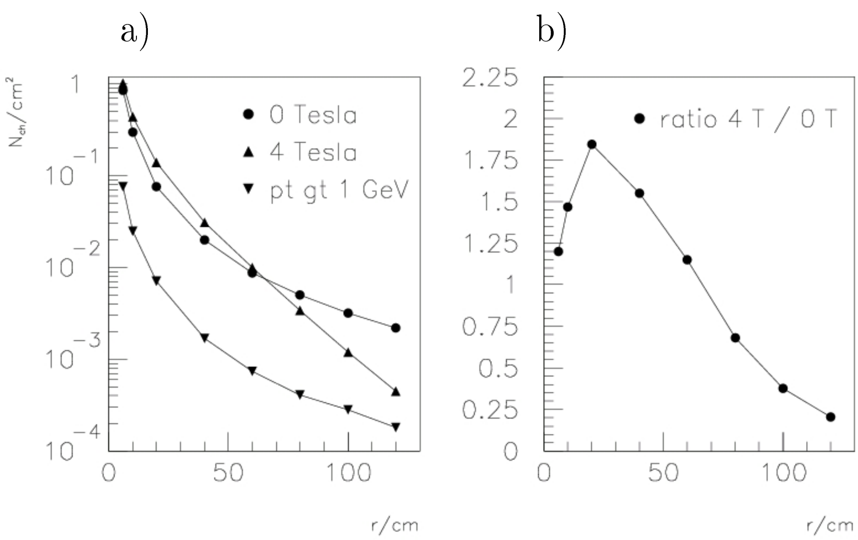

Figure 2.4: a) Charged particle densities per cm^2 at $\eta = 0$ as a function of the distance from the interaction point for 20 superimposed minimum bias events [27], with the magnetic field on (4 T) and off (0 T). b) Ratio of the two charged particle densities.

Close to the interaction point at radii $r < 20\,\text{cm}$, the high particle fluences are amplified by the presence of the magnetic field and require a highly granular device such as a pixel detector. The pixel size of $100\,\mu\text{m} \times 150\,\mu\text{m}$ is driven by the desired impact parameter resolution and gives an occupancy of the order of 10^{-4} per bunch crossing. The resulting track resolution is similar in both $r-\phi$ and z direction and allows for 3D vertex reconstruction in space (important for secondary vertex reconstruction with low track multiplicity). The pixel detector is arranged in three cylindrical layers of hybrid pixel detector modules at radii of 4.4, 7.3 and 10.2 cm complemented by two disks at $|z| = 34.5$ and 46.5 cm. The barrel pixel geometry leads to charge sharing across neighbouring pixels due to a Lorentz angle of 23° of the electrons in the magnetic field. This large Lorentz effect improves the hit resolution in $r-\phi$, allowing a spatial resolution in the range of $15-20\,\mu\text{m}$ to be achieved. In a similar way, charge sharing is induced in the forward disks by arranging them in a turbine-like geometry with blades rotated at 20°. The pixel system covers a pseudorapidity range of $-2.5 < |\eta| < 2.5$. Since the hit reconstruction in the pixel detector has a very low inefficiency (0.5 %) and a low ghost rate (0.01 %), the pixel detector is particularly useful for track seeding (see section 2.7). It also plays an important role in primary vertex finding and in High Level Trigger (HLT) algorithms, as for example in the displaced dimuon trigger used in the physics analysis part of this work (see section 7.2).

The intermediate region $20\,\text{cm} < r < 55\,\text{cm}$ is instrumented with a four layer microstrip silicon detector in the barrel region complemented by three disks at each side. The strips are 10 cm in length with a minimum pitch of $80\,\mu\text{m}$ in the barrel region and $100\,\mu\text{m}$ in the endcaps—giving an occupancy of up to $2-3\,\%$ per bunch crossing. Depending on the $r-\phi$ pitch, the single point resolution in $r-\phi$ is 23 or $35\,\mu\text{m}$.

In the outermost region at radii $r > 55\,\text{cm}$ the magnetic field enhances the rapid decrease of charged particle rates. The particle flux is sufficiently low enough to increase the strip length to 25 cm and a maximum pitch of $180\,\mu\text{m}$ and $184\,\mu\text{m}$ in the barrel and endcap region respectively. The increase in strip size is also necessary to limit the number of readout channels covering the large area. However, electronic noise grows linearly with the strip length[1] and in order to keep the signal to noise ratio above 10, the thickness of the sensors was increased to $500\,\mu\text{m}$ in the outer region of the tracker. The resulting higher depletion voltage can be reduced by choosing a higher initial resistivity, so that the initial depletion voltages of the thick and thin sensor are within the same

[1] noise scales with the capacitance $C \propto \frac{A}{d}$.

range.[1] The outer tracker consists of six barrel layers of silicon microstrip detector that surround the inner tracker, supplemented by nine disks on both sides, each carrying up to seven rings of silicon microstrip detectors (see Figure 2.5). The occupancy of the outer tracker amounts to approximately 1 % per bunch crossing. Depending on the $r-\phi$ pitch, the single point resolution in the outer barrel is between 35 and 53 μm in $r-\phi$.

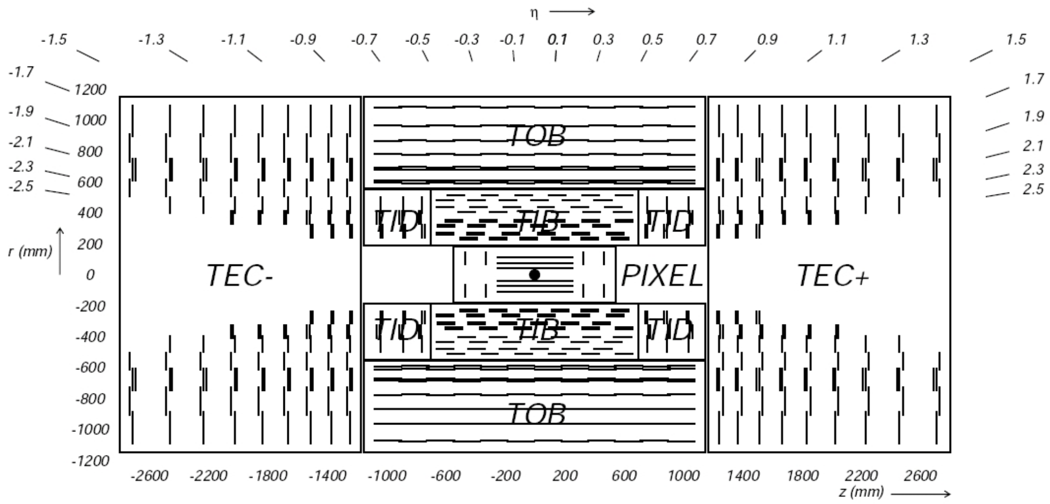

Figure 2.5: Schematic layout of the CMS tracker. The innermost detector consists of the barrel and forward pixel detector. The intermediate region holds the Tracker Inner Barrel and Disks (TIB/TID). The outer parts are the Tracker Outer Barrel (TOB) surrounding TIB/TID and the Tracker EndCaps (TEC) [5].

As indicated by the double lines in Figure 2.5, some layers are equipped with stereo-modules. In that case, two modules are mounted back-to-back at a stereo angle of 100 mrad, hence providing a measurement in (r, z) as well as in (r, ϕ). This enables single point resolutions of 230 and 530 μm in z to be achieved in the inner and outer barrel respectively. As shown in Figure 2.6, the layout of detector components ensures ≈ 9 hits up to $|\eta| < 2.4$ and the ultimate tracker coverage ends at the $|\eta| < 2.5$. The complete tracking system is 5.8 m long and has a diameter of 2.5 m. The total active silicon area, embodying 75 million readout channels, covers an area of 200 m^2, which makes the CMS tracker the largest silicon detector device ever built.

[1] With respect to radiation damage this is only advisable in the outer region of the tracker where the radiation levels are lower.

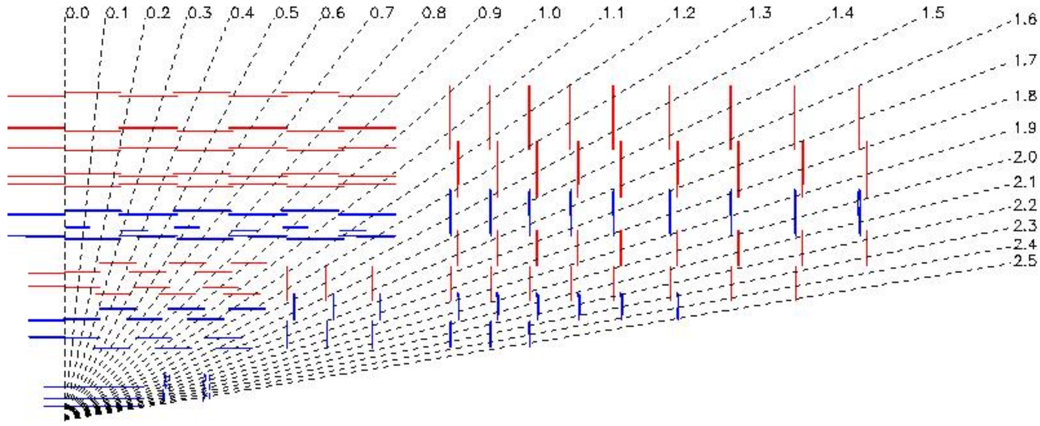

Figure 2.6: Layout of one quarter of the CMS tracker components.

The second part of this work, which treats the discovery potential of the rare decay $B_s^0 \rightarrow \mu^+\mu^-$ in CMS, relies largely on the tracker. Besides the precise determination of the secondary decay vertex of the B_s^0 meson, the tracker also ensures a good muon momentum resolution of the low momentum muons from the $B_s^0 \rightarrow \mu^+\mu^-$ decay: at $p_\perp$ below 200 GeV the resolution in the muon chambers is dominated by multiple scattering and the resolution from the tracker system is improved by an order of magnitude (see also Figure 2.14 in section 2.6). Figure 2.7 illustrates the expected tracker resolution of the transverse momentum, the transverse impact parameter and the longitudinal impact parameter as a function of pseudorapidity for single muons with different transverse momenta.

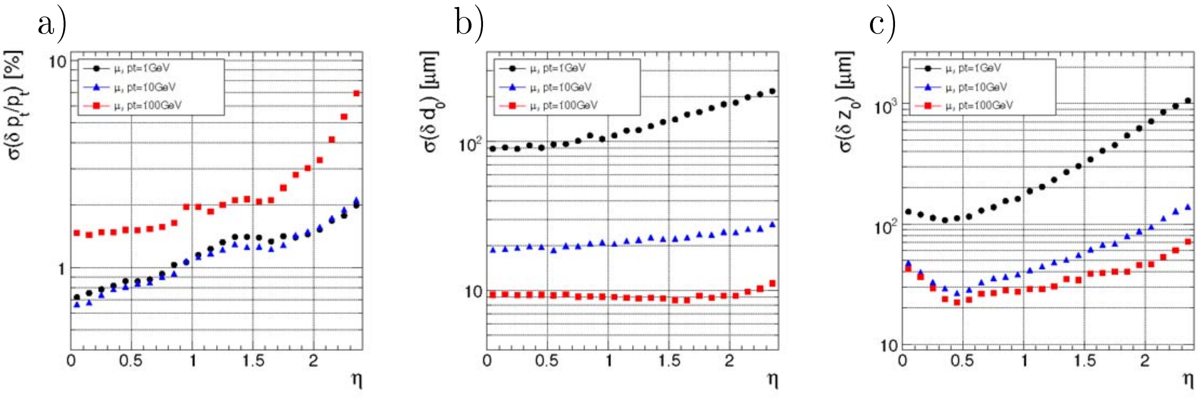

Figure 2.7: Tracker resolution of several track parameters for single muons with transverse momenta 1, 10 and 100 GeV. a) transverse momentum, b) transverse impact parameter and c) longitudinal impact parameter [5].

Table 2.1: Design parameters of the CMS tracker.

	Pixel	Inner tracker	Outer tracker
Active area	$1\,m^2$	$198\,m^2$ (combined)	
Channels	66 Mio	9.3 Mio (combined)	
Occupancy	$1\,\%$	$2-3\,\%$	$1\,\%$
Sensor thickness	$285\,\mu m$	$320\,\mu m$	$500\,\mu m$
Length	$150\,\mu m$	$10\,cm$	$25\,cm$
Barrel	$4\,cm < r < 11\,cm$	$20\,cm < r < 55\,cm$	$55\,cm < r < 110\,cm$
Dose in $500\,fb^{-1}$	$840-190\,kGy$	$70-11\,kGy$	$11-1.8\,kGy$
Layers	3	4	6
$r-\phi$ pitch	$100\,\mu m$	$80\ (120)\,\mu m$	$180\ (122)\,\mu m$
Resolution	$15\,\mu m$	$23\ (35)\,\mu m$	$53\ (35)\,\mu m$
Resolution in z	$20\,\mu m$	$230\,\mu m$	$530\,\mu m$
Stereo layers	-	1,2	1,2
Disks	$34.5 < \|z\| < 46.5\,cm$	$80\,cm < \|z\| < 90\,cm$	$124\,cm < \|z\| < 284\,cm$
Layers	2	3	9
$r-\phi$ pitch	$100\,\mu m$	$100-141\,\mu m$	$97-184\,\mu m$
Stereo rings	-	1,2	1,2,5

2.4 The Electromagnetic Calorimeter

The Electromagnetic Calorimeter (ECAL) [28] is designed to identify and precisely measure the energy and direction of electrons and photons in the experimental environment of CMS. It surrounds the tracker and, in combination with the Hadronic Calorimeter, allows the determination of jet energies with high precision. The 61200 lead tungstate ($PbWO_4$) scintillating crystal in the barrel region and the 7324 crystals in each of the two endcaps provide a hermetic, homogeneous coverage up to $|\eta| = 3$. The geometrical configuration of one quarter of the crystals is illustrated in Figure 2.8 by a transverse section through of the ECAL.

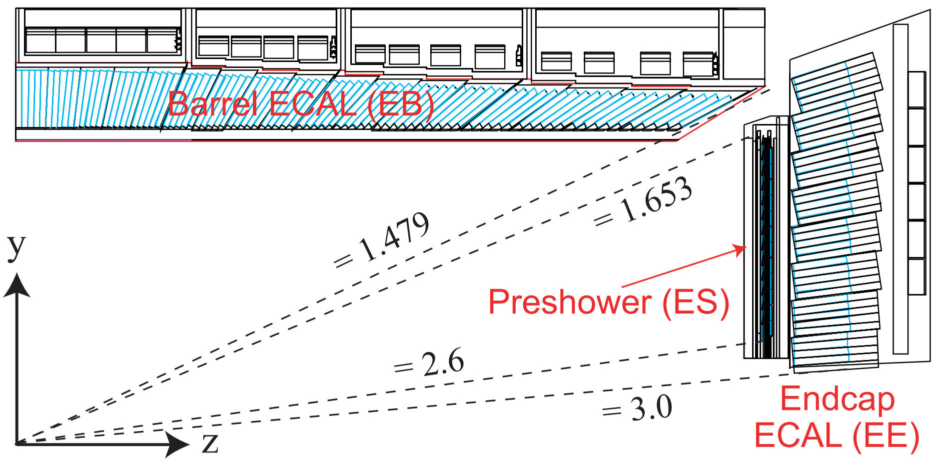

Figure 2.8: Partial transverse section through the ECAL, showing the geometrical configuration.

The high density $PbWO_4$ crystals ($8.2\,\mathrm{g/cm^3}$) have a short radiation length ($X_0 = 0.89\,\mathrm{cm}$) and a small Molière radius[1] (2.2 cm). This allows for the ECAL to be a very compact and highly granular device that is fast and radiation tolerant. The crystals are shaped like truncated pyramids, with a lateral size close to the Molière radius and a length corresponding to approximately $25X_0$ in terms of radiation thickness. The axes of the crystals in the barrel (EB) are inclined at an angle of $3°$ with respect to the vector originating from the nominal interaction vertex. The axes of the endcap (EE) crystals point to a focus point 1300 mm beyond the interaction point. The relatively low light yield requires the use of photo-detectors with an intrinsic gain even in the presence of a high magnetic field. Therefore the scintillation light is collected

[1] i.e. the radius of a cone containing 90 % of the energy of the shower

by Avalanche Photodiodes (APD) and Vacuum Phototriodes (VPT) in the barrel and endcaps respectively.

The main purpose of the two-layered preshower device (ES) placed in front of the EE is to reject $\gamma - \pi^0$ background to $H \rightarrow \gamma\gamma$, where the two closely spaced photons from the decay $\pi^0 \rightarrow \gamma\gamma$ fake a single photon. A lead absorber disk of $2X_0$ initiates an electromagnetic shower of incoming photons (and electrons). Two planes of silicon strip detectors measure the energy and transverse shower profiles. Besides improving the position determination of the incident photons, this also helps to distinguish electrons from minimum ionising particles (MIP).

The results from a test beam in 2004 [29], in which the energy resolution of the ECAL was determined using electrons with energies between 20 to 250 GeV, are shown in Figure 2.9. The intrinsic energy resolution can be parametrised as a function of the energy, and can be described as

$$\left(\frac{\sigma}{E}\right)^2 = \left(\frac{2.8\,\%}{\sqrt{E}}\right)^2_S + \left(\frac{0.12}{E}\right)^2_N + (0.30\,\%)^2_C, \tag{2.2}$$

where E is in units of GeV and the different contributions are given by the stochastic term, the noise term and the constant term respectively.

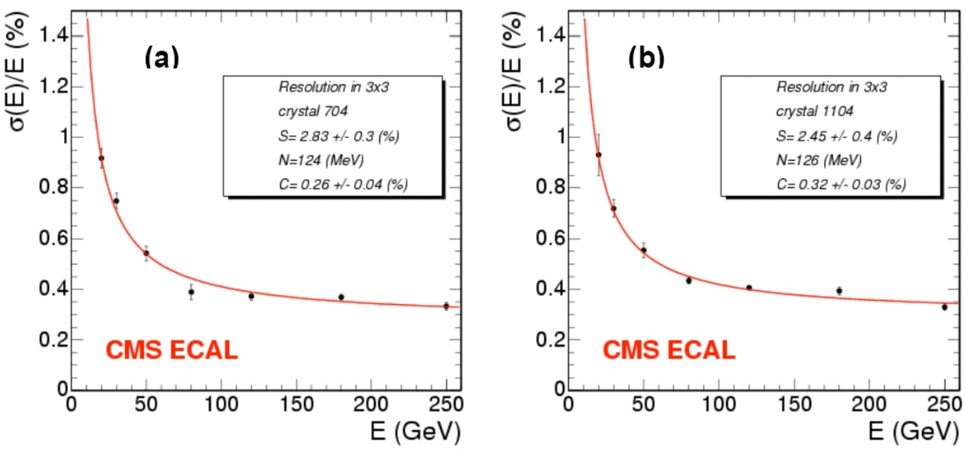

Figure 2.9: The energy resolution, σ_E/E as a function of the electron energy, for a 3×3 array of two reference crystals: a) 704 and b) 1104 [29].

2.5 The Hadronic Calorimeter

The Hadronic Calorimeter (HCAL) [30] is placed around the electromagnetic calorimeter and aims to measure the energy and direction of hadron jets. The barrel part (HB) and the endcaps (HE) hermetically cover the pseudorapidity range up to $|\eta| = 3$ and are entirely immersed in the magnetic field of the solenoid. The HB and HE are segmented into 2304 oriented towers, consisting of respectively 17 and 19 tiles of active plastic scintillator (readout with wave-length shifting fibres), interspersed between brass absorber plates. Figure 2.10 illustrates the tower segmentation of one quarter of the HCAL.

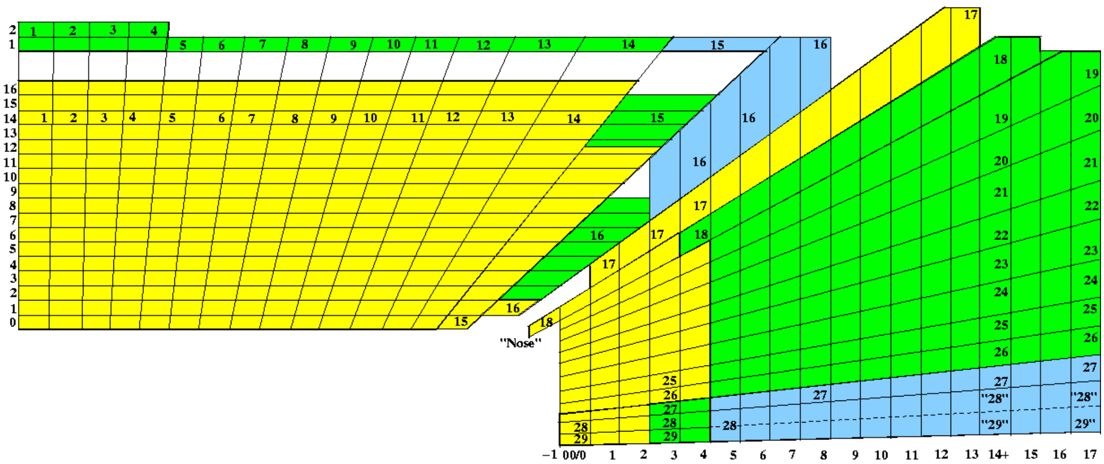

Figure 2.10: A schematic of the tower mapping in $r-z$ of the HCAL barrel and endcap regions.

Brass was selected as it has a relatively short interaction length and is non-magnetic. Additional structural strength in the barrel is provided by stainless steel plates in the inner- and outermost layers. Since the amount of absorber material that can be placed in the HCAL is constrained by the inner radius of the solenoid, a "tail catcher" is placed outside the solenoid to reduce the tails in the energy resolution function. This hadron outer detector (HO) consists of two scintillator layers on either side of an iron absorber. Taking into account the material of the magnet coil, the effective thickness of the HB amounts to over 10 hadronic interaction lengths.

The hadron forward calorimeters (HF), located 11.2 m from the interaction point, provide an extended hermetic coverage up to $|\eta| = 5.2$ for measuring missing transverse energy. The hostile environment in the forward region with very high charged hadrons rates lead to Cherenkov-based technology consisting of steel absorbers and embedded radiation hard quartz fibers.

The energy resolution of the combined barrel calorimeters was measured in a test beam with hadrons, electron and muons in the energy range 2 to 350 GeV. The optimised energy resolution[1] of the combined EB+HB system was found to be $\sigma/E = 84\,\%/\sqrt{E} \oplus 7\,\%$ [31], where the first and second terms in the equation represent the stochastic and constant term respectively. The results from a test beam with π^- are shown in Figure 2.11.

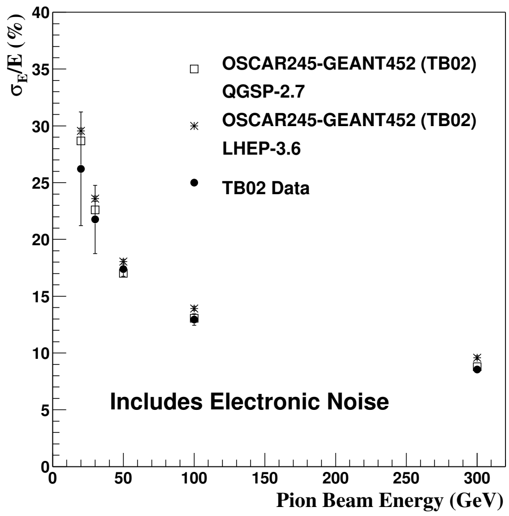

Figure 2.11: Energy resolution σ/E_{π^-} measured in a π^- test beam as a function of the beam energy [32].

[1] after applying corrections to compensate for the different intrinsic electron to hadron response (e/h) in the ECAL and HCAL

2.6 The Muon System

Muons are prominent signatures in most final states of the physics probed by the LHC. Muons are cleanly measurable objects due to their long lifetime, high mass, high penetration power and low radiative losses.[1] As shown in Figure 2.12, the muon spectrometer [5; 33] is hosted in the magnet return yoke and provides a full geometric coverage up to $|\eta| = 2.4$. At least 16 interaction lengths of material are present over the whole η range (Figure 2.13), ensuring efficient muon identification by absorption of other charged particles before the muon system (in the HCAL and ECAL) and inside the muon system (in the iron yoke).

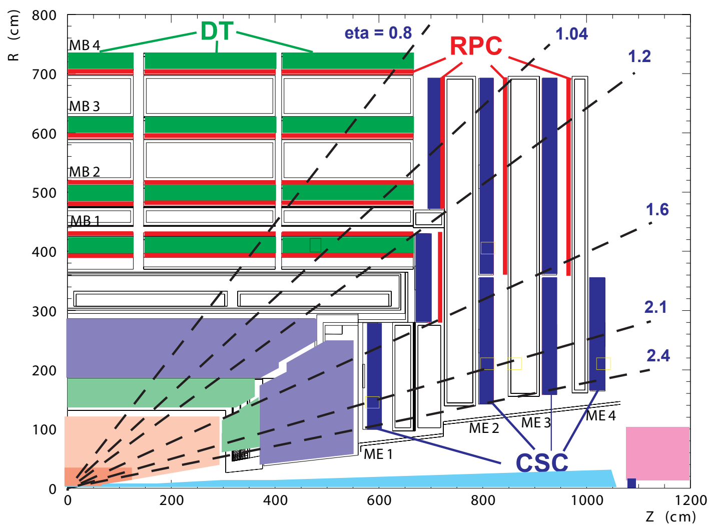

Figure 2.12: Layout of one quarter of the muon system for initial low luminosity running.

The muon system uses three types of gaseous particle detectors. Drift Tubes (DT) are used in the barrel region ($|\eta| < 1.2$), where the magnetic field is confined to the iron yoke, the muon rate is low and the neutron induced background rates are small. The endcap discs are instrumented with Cathode Strip Chambers (CSC) in order to deal with the strong, non-uniform magnetic field and the high charged particle rates in the forward region ($0.9 < |\eta| < 2.4$). The DTs and CSCs provide precise time and position measurements and are both complemented by the Resistive Plate Chambers

[1]Unlike most particles, high energy muons are not stopped in any of the calorimeters and they are less affected by radiative losses in the tracker than electrons.

(RPC). The fast response of the RPCs and a time resolution of 3 ns allow unambiguous assignment of a muon track to the correct bunch crossing.

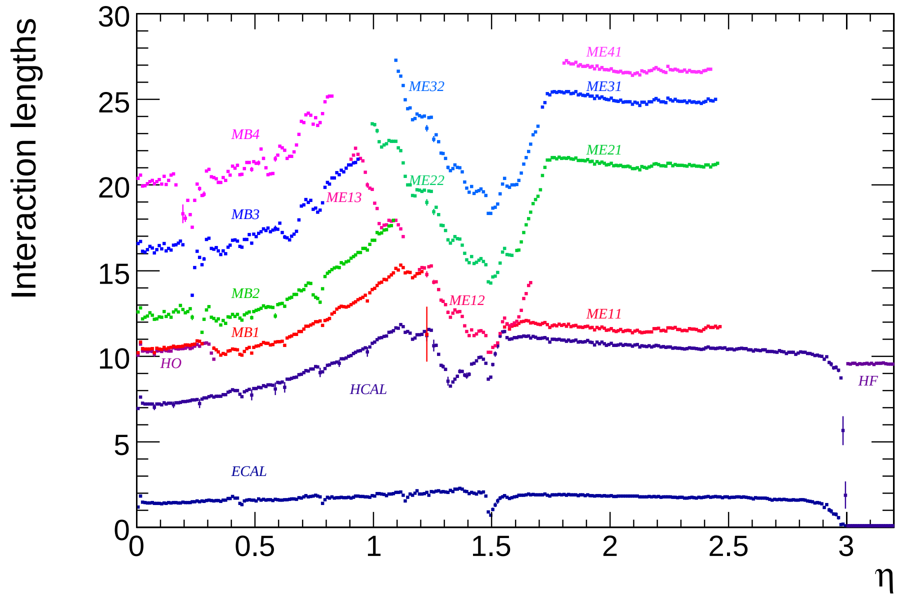

Figure 2.13: Material thickness in interaction lengths after the ECAL, HCAL, and at the depth of each muon station as a function of pseudorapidity. The thickness of the forward calorimeter (HF) is only partially shown and remains approximately constant over the range $3 < |\eta| < 5$ [32].

As a function of transverse momentum, the muon momentum resolutions using the tracker only, the muon system only and combinations of parts of the tracker and muon system are illustrated in Figure 2.14. For p_T values below 200 GeV, where the resolution in the muon chambers is dominated by multiple scattering, the best momentum resolution is given by the resolution obtained in the silicon tracker. If multiple scattering and energy losses are negligible the muon trajectory after the coil extrapolates back to the beam line. This fact can be used to improve the muon momentum resolution at high momentum when using the full system.

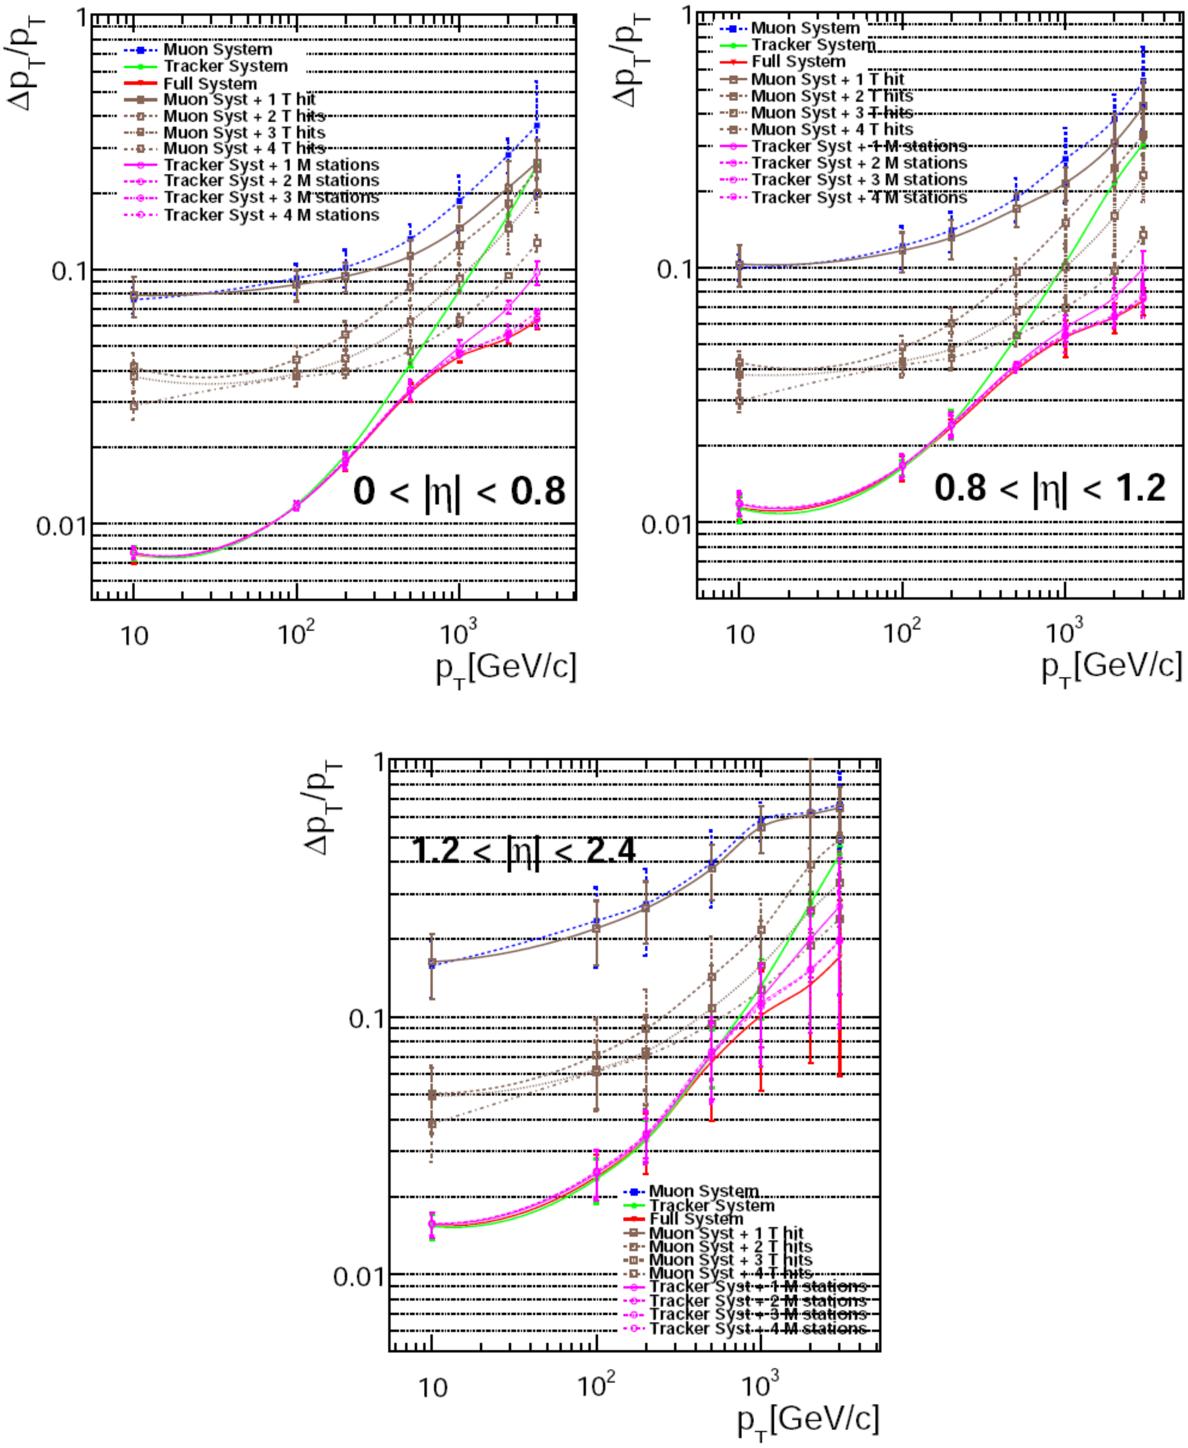

Figure 2.14: Momentum resolution in three different pseudorapidity regions as a function of the transverse momentum, for the tracker system, the muon system, tracker and muon system combined and for combinations of parts of the muon system (M) and parts of the tracker system (T) [34].

2.7 Track Reconstruction

Track reconstruction can be divided into five logical parts:

- Hit reconstruction
- Seed generation
- Pattern recognition (trajectory building)
- Ambiguity resolution (trajectory cleaning)
- The final track fit (trajectory smoothing)

At least three reconstructed hits or two reconstructed hits compatible with the beam spot are required for seed generation. Seeding provides the initial five-parameter description of the helical trajectory to start track building: starting from the position of a seed, trajectory building propagates each seed to the next detector layer (taking into account multiple scattering and energy losses) and a trajectory candidate is formed for each compatible hit. The pattern recognition is based on the combinatorial Kalman filter method [35], using the trajectory updated with the corresponding hit before searching for a compatible measurement in the next layer. The procedure is repeated for all trajectories until the outermost layer is reached or until a stopping condition [1] applies. This creates a large number of tracks, many of which partially share the same hits. If the fraction of shared hits between two trajectories is too large, the ambiguity has to be resolved to avoid double counting and thus only the highest-quality trajectory is kept. Since the full information is only available at the last hit and constraints applied during trajectory building can bias the estimate of the track parameters, all valid tracks are refitted with a standard Kalman filter and a second filter (smoother) running from the exterior towards the beam line.

The reconstruction efficiency of single muons tracks with transverse momenta 1, 10 and 100 GeV is illustrated in Figure 2.15 as a function of η. The efficiency is 99 % except in the regions $\eta < 0.1$ and $\eta > 2.0$. At low η the drop is because of the gaps between the sensors on the ladders of the pixel barrel at $z = 0$, and at high η the efficiency decreases due to the lack of coverage by the forward pixel detector.

[1] e.g. to limit the CPU time in the HLT, where only a partial track reconstruction with less than $5-6$ hits is necessary to achieve the required accuracy

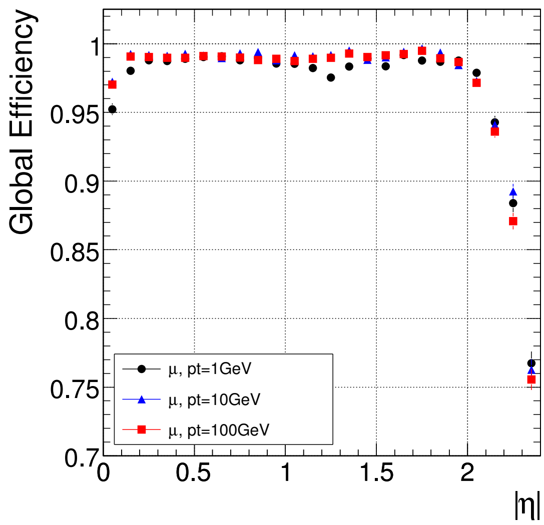

Figure 2.15: The track reconstruction efficiency for muons with transverse momentum as a function of η [32].

Part I

CMS Pixel Module Qualification

Chapter 3

The CMS pixel detector

At the LHC design luminosity of $10^{-34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, approximately 1000 particles from over 20 minimum bias events are expected on average per bunch-crossing. As shown in chapter 2, the presence of the 4 T field influences the charged particle densities and additionally enhances the charged particle fluences close to the interaction region. Therefore, a highly granular device with a fast response is required, not only to achieve a precise and efficient measurement of charged particle trajectories, but also to allocate each track to the correct bunch crossing. The region closest to the interaction point (i.e. at radii $r < 20\,\mathrm{cm}$) is instrumented with a pixel detector consisting of three cylindrical layers in the barrel region and two endcap disks on each side. The geometrical layout and properties of the pixel detector have already been discussed in detail in section 2.3 as part of the CMS silicon tracker.

The following sections will focus on the barrel part of the pixel system. Section 3.1 outlines the basic layout of the support structure and the supply system. The active silicon area of the barrel pixel detector consists of 768 entities mounted on the cylindrical half shells of the support structure. These basic entities of the pixel detector—called modules—are discussed in section 3.2. The readout of the silicon sensor pixels proceeds through a readout chip, which is connected to the pixels on the sensor by $\sim 20\,\mu$m indium bumps. The properties of this readout chip are described in section 3.3. Finally, the analogue readout sequence that will be sent from the module to the front end electronics is explained in section 3.4.

3.1 The Pixel Barrel System

The three layers of the barrel pixel detector, each divided into two half-cylinders, are located at radii of 4.4, 7.3 and 10.2 cm. Each half-cylinder contains ladders and half-ladders that provide the support structure and the cooling for pixel modules. The half-ladders at the edge of each half-cylinder have a small overlap and ensure hermetic coverage in $r - \phi$. Each ladder consists of eight modules. The normal direction of the module on each ladder alternates, pointing either towards the beam or away from it. The $-z$ and $+z$ sides, ranging from -285 mm to 285 mm around the interaction region, are electrically separated. Each side of a half-cylinder is divided into eight independently operating sectors (with the exception of the slow control). As shown in Figure 3.1, the detector half shells are completed by support frames on both sides, that fix the three detector layers. Printed circuit boards are mounted onto these frames, holding the connectors for the module cables and providing power to the modules of the eight individual sectors. Services from patch panels located outside the tracker volume are carried to the barrel through the supply tubes on each side of the detector. To allow installation in the presence of the beam-pipe, the supply tubes are also split into two halves. They carry the cooling fluid and the electric power lines as well as the optical fibers and electronics for read-out and control. The length of the full system amounts to 5.6 m, where the detector itself with a length of 570 mm makes up only a small part of the whole pixel barrel system.

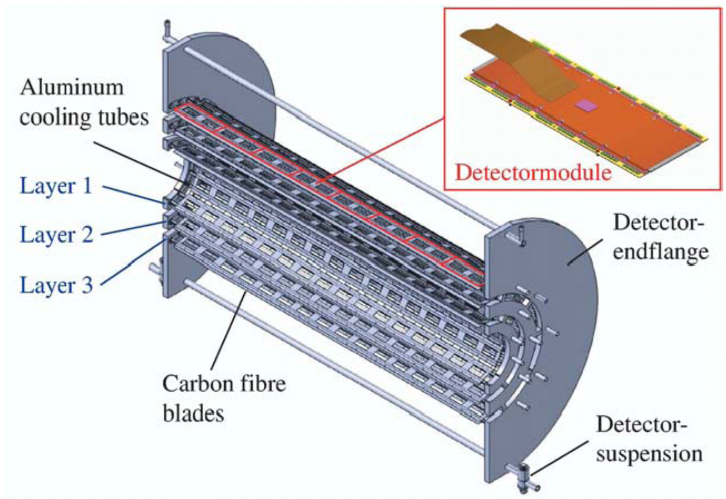

Figure 3.1: Support structure of a barrel pixel half shell [5].

3.2 The Detector Modules

The pixel barrel detector contains 768 modules in total, of which 672 modules are full modules and 96 are half-modules. The half-modules come in left-handed and right-handed versions, and are mounted on the edges of the six half-cylinders. With 66560 (33280) pixels on a full (half) module, the total number of channels on the barrel pixel detector is 48 million. A full module weighs 3.5 g and has dimensions of 66.6 mm $\times$ 26.0 mm.

A full (half) module consists of the following components (see Figure 3.2):

- Two (one) silicon nitride base strip to fix the module on the support structure.
- An n-on-n silicon sensor device: The pixels are formed by a high dose $n+$-implant introduced into a n-doped silicon substrate with high resistance. The pixels are isolated using the moderated p-spray technique. The small gap between the collecting electrodes (i.e. the n-implant) ensures a homogeneous drift field and also leads to a high capacitance. The rectifying pn-junction is placed on the back of the sensor and is surrounded by a multi guard ring structure that allows all sensor edges to be kept to a ground potential. To perform an on-wafer measurement of the current-voltage characteristics, each pixel is connected to a bias grid through a high resistance punch through connection (bias dot). The sensor thickness is 285 μm, giving an ionisation charge of approximately 23 ke^- for a (perpendicularly incident) Minimum Ionising Particle (MIP). The sensor is fully depleted at a reverse bias voltage of $50 - 60$ V. It will be operated at a voltage of 150 V initially. After irradiation at high particle fluences, higher bias voltages of up to 600 V will be needed to compensate for the irradiation damage in the sensor.
- 16 (8) readout chips (ROC): The ROCs are thinned down to 180 μm and contain 52×80 pixel unit cell (PUC) per ROC. Each PUC is connected to a pixel on the sensor through an indium bump bond with a diameter of approximately 20 μm. Since the required bump size could not be achieved with the standard industrial technology, a procedure using reflown indium bumps was developed at the Paul Scherrer Institute (PSI) [36].
- A High Density Interconnect (HDI): The HDI distributes the power and control signals to the chips and transmits the readout from the double column periphery

of the ROCs (see section 3.4) to the Token Bit Manager (TBM). The TBM is located on top of the HDI and consists of two identical entities that control the readout of a group of ROCs (up to 24 ROCs per TBM unit). The connections between the ROCs and the HDI as well as those between the HDI and the TBM are formed by wire-bonds (see section 5.1.1).

- A two layer Kapton/copper compound cable: The Kapton cable has 21 traces and transmits the readout and control signals. It is connected to the HDI through wire-bonds.
- A power cable: The cable consists of six copper coated aluminium wires each soldered to the corresponding pads on HDI. It provides the bias voltages for the sensor depletion and the voltages for the digital and analogue parts of the ROCs.

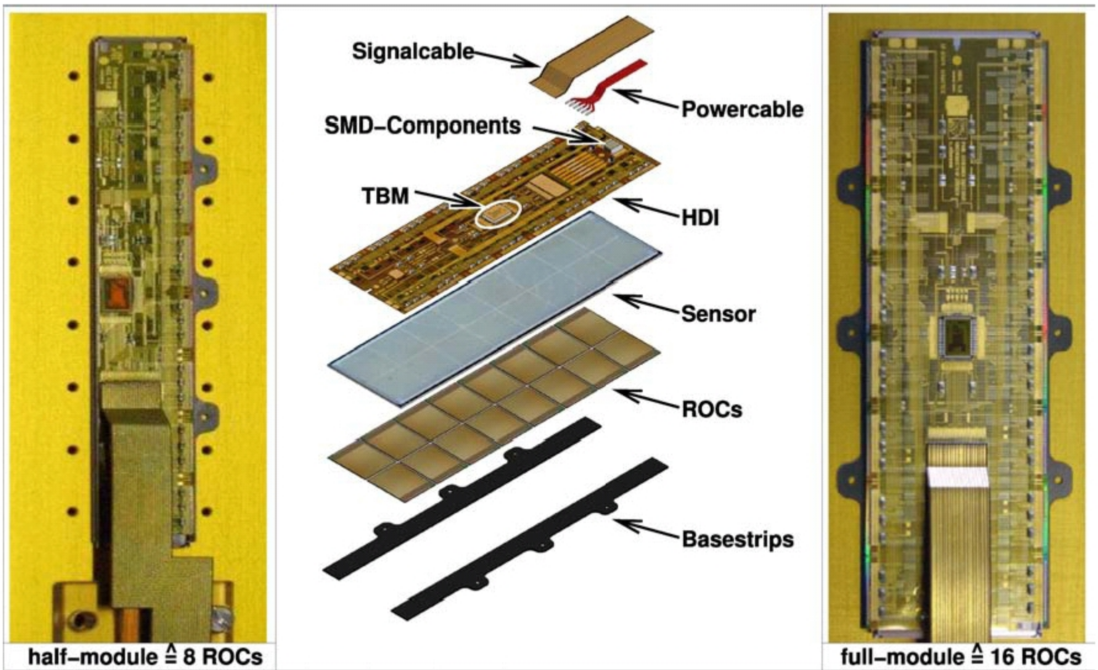

Figure 3.2: View of a half-module (left) and a full module (right) fully assembled. Middle: Exploded view of a barrel pixel module showing the two silicon nitride base strips, the 16 readout chips (ROCs), the silicon sensor, the High Density Interconnect (HDI) with the Token Bit Manager (TBM) and the power and Kapton cables.

The TBM and ROCs are both produced with the radiation hard 0.25 μm CMOS technology. The lifetime of a module is therefore mainly limited by the radiation damages in the sensor. The double-sided processed n-on-n sensor design allows operation of the modules with a partially depleted sensor and maintains a high charge collection efficiency at moderate voltages (< 600 V).

3.3 The Readout Chip

The charge produced by an ionising particle traversing the silicon sensor is collected at the electrode formed by the n-implant and a signal voltage is induced in the PUC through a capacitor. The pixels in a ROC are read out with the column drain mechanism, where the readout starts with the pixel closest to the periphery and proceeds by going up the left side of a double-column and coming back down the right column. The pixel detector will be operated in a zero suppression mode, therefore only pixels with a signal above a certain threshold will be read out. To optimise the signal processing and readout, and to compensate chip-to-chip variation, there are 26 Digital to Analogue Converters (DAC) and three registers controlling the voltages and currents on the PUC and the double-column periphery (see Figure 3.3). These DACs can be set individually for each ROC and apply to all PUCs and double-columns on the same ROC. A table listing the complete set of DACs and their application can be found in Appendix A.

The generated signal voltage is processed in a preamplifier (*VwllPr*, *VrgPr*) and a shaper system (*VwllSh*, *VrgSh*). If the signal exceeds the reference voltage in the comparator (*VthrComp*), it is passed on to the sample and hold capacitor with an adjustable delay (*VhldDel*), and the double-column periphery is notified. The signal is stored in the capacitor until the double-column periphery starts the readout and writes the address of the hit pixel, the pulse height and the bunch crossing in a data buffer.

For testing and optimisation, an internal calibration signal can be injected directly into the pixel readout chain. The amplitude of the calibration voltage can be varied with the *Vcal* DAC and a delay time can be set with the *CalDel* DAC. Two voltage ranges are available for the calibration signal. The range can be selected by setting the corresponding bit in the control register (*CtrlReg*). For the same value of *Vcal*, the amplitude of the injected signal in the high range is about seven times higher than in the low range. In the low range, one unit of *Vcal* corresponds to approximately 65 electrons, and in the high range to approximately 455 electrons.

The mask bit and four trim bits can be programmed separately for each pixel: the comparator of a PUC can be disabled with the mask bit and the threshold of a pixel can be adjusted individually with the trim bits. The impact of the trim bit value on the threshold depends on the value of the *Vtrim* DAC.

The ROC is programmed using a ‘fast I^2C’ interface in a 10 bit format with two synchronisation bits that are ignored. In the following text, a byte refers to 8 bits. The first byte is always the command byte (*ClrCal, Prog_DAC, ProgPix* and *Cal_Pix*), specifying the chip address and the type of command. The calibrate mode can be removed from all pixels with the corresponding command byte *ClrCal*. The DACs are programmed with a three byte command containing the command byte *Prog_DAC*, the DAC address and the DAC value. Pixel individual settings such as the mask bit and trim bit use a four byte command, starting with the command byte (*ProgPix*) followed by two byte for the column and row address, and the data byte with the required bit settings of five bits (one mask bit and four trim bits). The same four byte command structure (with *Cal_Pix*) is used to change between calibration through the sensor bumps or directly through a capacitor, by setting the two bits in the last byte. In case of the last two commands, multiple programming of pixels in the same double-column is possible and speeds up download times.

3.4 The Analogue Chain

The readout of a module is initiated by the TBM emitting a token bit to the ROCs for an incoming Level-1 trigger. The token bit is passed on from ROC to ROC and finally back to the TBM. Through a Kapton cable the analogue readout is sent from the TBM to the end ring of the pixel barrel. There, the analogue signal is separated from the digital one and is transmitted to the Analogue Optical Hybrids (AOH). The analogue signal is then sent to the pixel Front End Digitisers (pxFED) through 40 MHz optical links. The pxFED, located in the electronics room, digitises and formats the data before sending it to the CMS-DAQ event builder.

Figure 3.4 illustrates an analogue readout sequence for a module with a hit in the first chip. The very low levels (here at -700 ADC units) are called ultra black levels (UBL). The level around zero is called the black level and defines the zero level of the analogue signal. The start and the end of a readout sequence are marked by the the TBM header and trailer, respectively. Both consist of eight clock cycles. The TBM header starts with three UBLs followed by a black level and four clock cycles encoding an 8-bit hit counter. The TBM trailer contains two UBLs, two black levels and four status bits. The TBM header and trailer confine the readout of the 16 ROCs. The

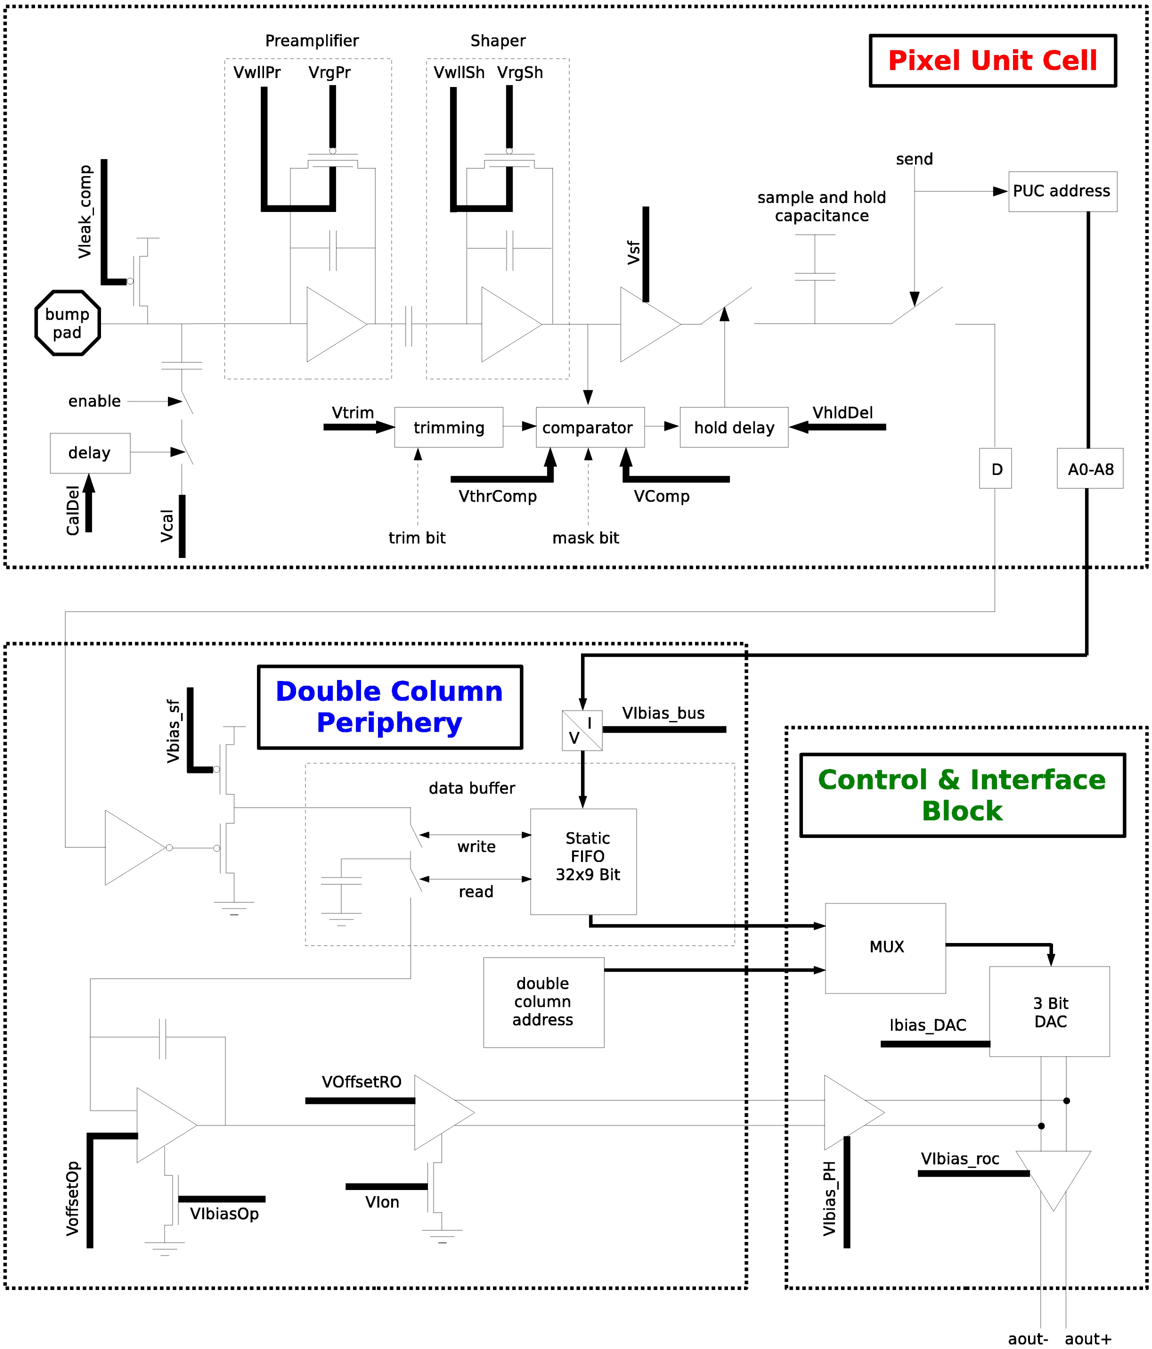

Figure 3.3: Diagram of the ROC control and readout system.

minimum readout of a ROC consists of three clock cycles: an UBL, a black level and a level called 'last DAC', representing the value of the most recently programmed DAC.[1] Therefore, a valid analogue readout for an empty bunch crossing always consists of 64 clock cycles. If one or several pixels have been hit, six more clock cycles per hit pixel will be appended to the associated ROC readout sequence: two for the double-column index, three for the row index and one for the signal charge.

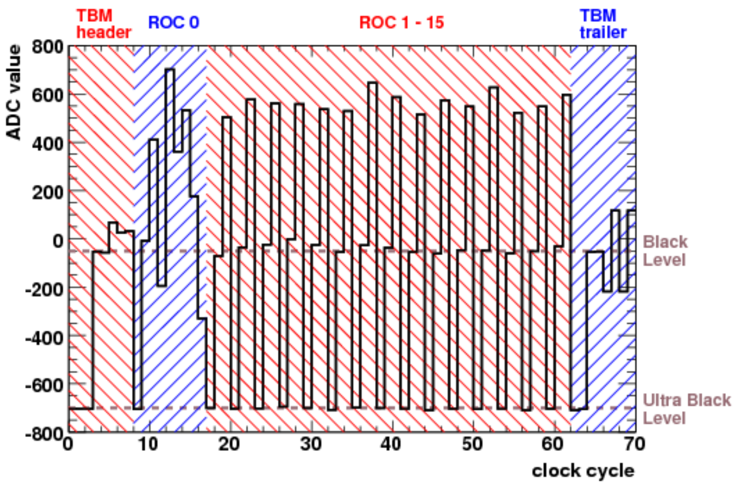

Figure 3.4: The analogue readout of a module with a hit in the first ROC (ROC 0).

[1] By default it contains the encoded output voltage of the built in temperature sensor.

Chapter 4

Test and Optimisation Algorithms

The pixel detector will be operated in a dense charged track environment with a bunch crossing every 25 ns. To ensure a highly efficient and precise track reconstruction even up to the highest luminosity, the functionality and performance of every module has been examined in an extensive test procedure before being released for use in the final detector system. The calibration and performance optimisation are likewise an important part of this procedure, as they ensure a uniform response over a whole ROC and an accurate charge measurement. Test and optimisation algorithms dedicated to the different operational aspects of a ROC were developed at the Paul Scherrer Institute (PSI) in Villigen. The algorithms were implemented in the module qualification procedure, which will be described in next chapter. Since threshold measurements are an essential part of many of the test and optimisation algorithms, the different types of thresholds are explained in section 4.2.

The test algorithms can be divided into three main categories: *functionality tests*, *calibration tests* and *performance tests*. The *functionality tests* include simple routines to validate the TBM readout, the programmability of a ROC or the pixel readout. More elaborate tests check for example the correct address decoding of each pixel or examine the quality of the bump bond between a PUC and a sensor pixel. The readout of a module is analysed by two 12-bit Analogue to Digital Converters (ADC). The purpose of the *calibration tests* is to convert a measured ADC signal into physical units, for instance the conversion of a given pulse height into an ionisation charge in units of electrons, or the temperature measurement with the built-in sensor by use of the 'last DAC'. The third test category establishes the *performance* of a module, such as the sensor leakage current or the pixel noise. The individual test algorithms of each of the three categories are described in sections 4.4—4.6, and in more detail in [37].

Based on a study [38] focusing on the optimisation of the ROC performance by use of the appropriate DAC settings, most of the DAC parameters are initialised to a default value, that was nonrecurringly optimised for all modules. Nine DAC parameters, however, need to be adjusted separately for each ROC in order to achieve optimum performance. The dynamic optimisation of these parameters is described in sections 4.3 and 4.7, and in more detail in [38]. A table listing the default values and optimisation criteria of each DAC can be found in Appendix B.

Section 4.8 details the results of the functionality, calibration and performance tests, and the performance optimisation in the 47.9×10^6 pixels and 1.15×10^4 ROCs of the 768 modules selected for implementation in the final detector system. The challenging task of processing and analysing the vast amount of data produced by the test and optimisation algorithms has been a major part of this work.

For the sake of clarity, a colour scheme, to display the test results obtained in the module qualification procedure, has been introduced and is explained in section 4.1.

4.1 Colour Codes for the Results from the Module Qualification Procedure

Each module was tested and qualified in two separate test suites, that will be described in detail in chapter 5. In this chapter, *test suite I* and *test suite II* refer to the first and second series of tests respectively. The former represents a comprehensive, 6 hours test series, focusing on the qualification and characterisation of every module emerging from module production and the latter a rather short test series of 3 hours, re-ensuring the basic functionality of the modules that had qualified for use in the detector system and establishing their optimum DAC settings. In the course of module testing, some of the results of the optimisation study [38] were already implemented into *test suite I*. The second test suite made sure that the optimised settings were also established for modules tested at the early module production stage. In addition, *test suite II* comprised of a calibration of the *Vcal*-DAC with two different radioactive sources. In both test suites measurements were carried out at $-10\,°\mathrm{C}$ and repeated at $17\,°\mathrm{C}$.

The colour codes for the different test suites and different temperatures are explained in Figure 4.1. The different temperatures are coded red and blue for $17\,°\mathrm{C}$ and $-10\,°\mathrm{C}$ respectively. The results of *test suite I* are denoted with the subscript I (e.g. $T_I = 17\,°\mathrm{C}$) and are represented in bright colours; whereas the results of *test suite*

II are denoted with the subscript *II* (e.g. $T_{II} = 17\,°\mathrm{C}$) and are displayed in muted colours. An eventual supplement “(m)” indicates that only the selection of modules, which were mounted onto the final detector, is shown in the figure.

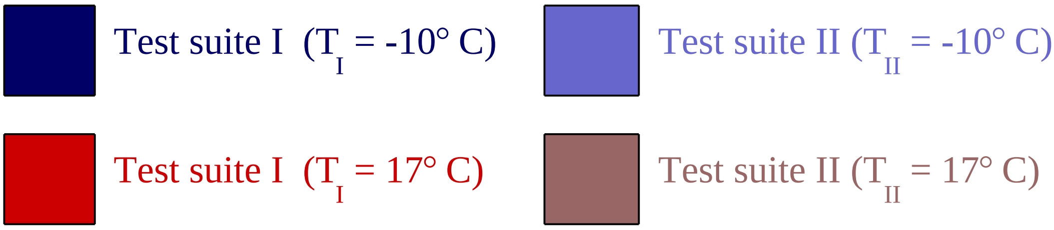

Figure 4.1: The colour codes for the two test suites: Bright colours will be used for *test suite I* and muted colours for *test suite II*. In both cases red and blue correspond to the temperature of $17\,°\mathrm{C}$ and $-10\,°\mathrm{C}$ respectively, at which the tests were performed.

4.2 Threshold Measurements

Depending on the context, a ‘threshold’ can describe a different physical property of a pixel. Usually the following types of thresholds are distinguished between:

- ***VthrComp*-threshold**: The calibration signal is injected with a constant *Vcal* value. The response efficiency is measured for decreasing *VthrComp*-values and the threshold is determined by the *VthrComp*-value at which the efficiency reaches 50%. Note that *VthrComp* is inverted, i.e. a higher value of the *VthrComp* DAC corresponds to a lower comparator threshold.

- ***Vcal*-threshold**: The comparator threshold is set to a fixed value of *VthrComp*. The response efficiency is measured for increasing *Vcal*-values. The threshold is given by the *Vcal*-value at which the measurements efficiency reaches 50 %.

- **In-time threshold**: The previously described thresholds are usually determined by searching for hits in a fixed bunch-crossing and are therefore referred to as in-time thresholds.

- **Timing independent threshold**: Due to different rise times, signals with different amplitudes can end up in different bunch-crossings. In particular signals with a low amplitude can end up in a later bunch-crossing, than signals with a high amplitude. The determination of the in-time and absolute thresholds is discussed in detail in [39].

4.3 Start-Up Tests

Before testing a module, several DAC parameters of a ROC and the TBM have to be adjusted, in order to be within the operational regime of the ROC. At start-up, all 26 DACs and the three registers are initialised to their default values. The default settings and optimisation criteria of each DAC can be found in Appendix B. The readout of a module is analysed by two 12-bit Analogue to Digital Converters (ADC), which sample the analogue signal in the interval $[-2048, +2047]$, with 1 ADC unit corresponding to 0.1275 mV. The following settings need to be established separately for each ROC.

4.3.1 Analogue Current

The nominal analogue current of 24 mA is set by adjusting the *Vana*-DAC. The distributions of the established *Vana*-values are illustrated in Figure 4.2.

4.3.2 Address Levels

To correctly decode the analogue readout of a module, the ultra black levels (UBLs) of the TBM and ROCs need to be adjusted. First, the ultra black levels of both TBM channels are set to a user-defined value of -1000 ADC units using the TBM DAC *Dacgain*. This also limits the highest TBM level to +1000 ADC units. In a second step, the ROC ultra black levels are adjusted to the same level as the TBM ultra black level. This can be achieved by variation of the *Ibias_DAC* (which concurrently appoints the position of the address levels). The distributions of the established *Dacgain*-values and *Ibias_DAC*-values are illustrated in Figures 4.3 and 4.4 respectively.

4.3.3 Threshold and Timing

To perform pixel tests with the internal calibration signal, the delay of the calibration signal with respect to the 40 MHz clock cycle and the signal threshold of the pixels must be tuned relative to each other. For this reason the readout efficiency over the entire *VthrComp-CalDel* parameter space is scanned with a *Vcal* of 200 in the low range. A stable working point is extracted by choosing a combination of *VthrComp* and *CalDel*, that lies in the center of a region with a high readout efficiency. The distributions of the *VthrComp*-values and *CalDel*-values established by this method are illustrated in Figures 4.5—4.7.

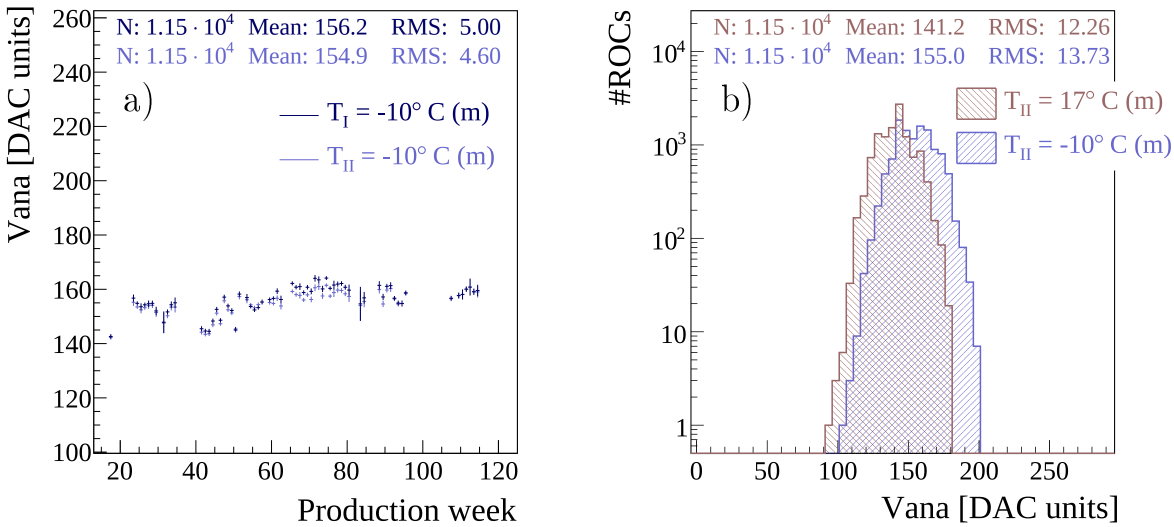

Figure 4.2: Distributions of the *Vana*-values of the 1.15×10^4 ROCs in the final system: a) The average *Vana*-value of modules produced in the same week as a function of time, determined at $-10\,°$C during *test suite I* (T_I) and *test suite II* (T_{II}). b) Distribution of the *Vana*-values from *test suite II* at $-10\,°$C and $17\,°$C.

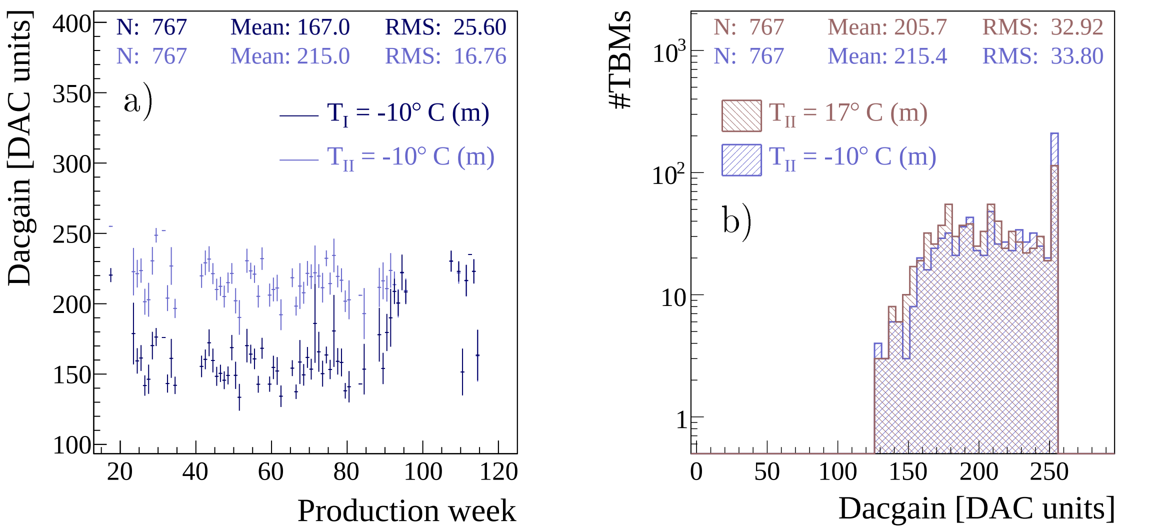

Figure 4.3: Distributions of the *Dacgain*-values of the 768 TBMs of the final system: a) The average *Dacgain*-value of modules produced in the same week as a function of time, determined at $-10\,°$C during *test suite I* (T_I) and *test suite II* (T_{II}). b) Distribution of the *Dacgain*-values from *test suite II* at $-10\,°$C and $17\,°$C.

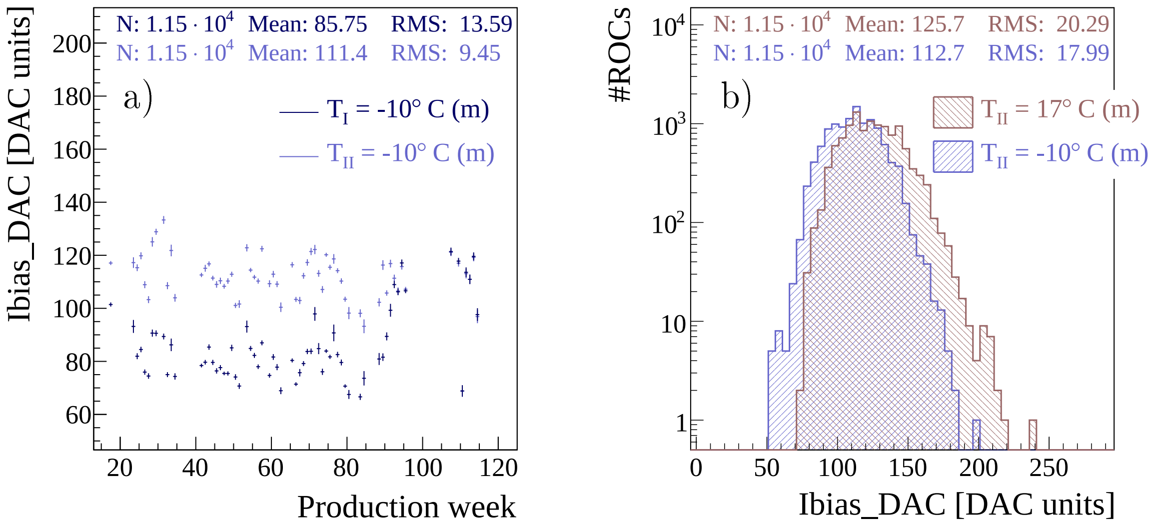

Figure 4.4: Distributions of the *Ibias_DAC*-values of the 1.15×10^4 ROCs in the final system: a) The average *Ibias_DAC*-value of modules produced in the same week as a function of time, determined at $-10\,°$C during *test suite I* (T_I) and *test suite II* (T_{II}). b) Distribution of the *Ibias_DAC*-values from *test suite II* at $-10\,°$C and $17\,°$C.

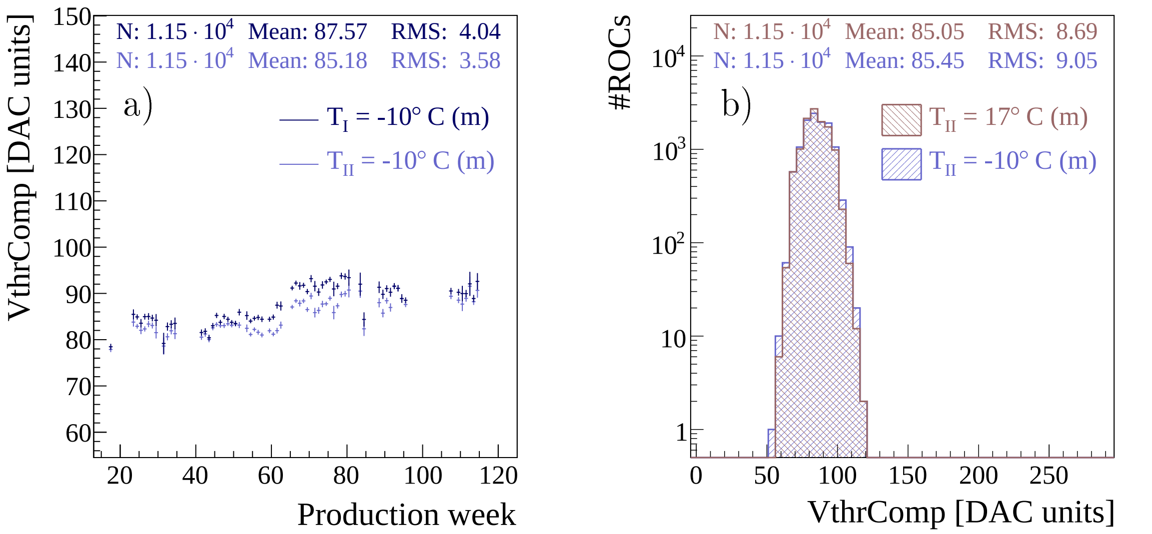

Figure 4.5: Distributions of the *VthrComp*-values of the 1.15×10^4 ROCs in the final system: a) The average *VthrComp*-value of modules produced in the same week as a function of time, determined at $-10\,°$C during *test suite I* (T_I) and *test suite II* (T_{II}). b) Distribution of the *VthrComp*-values from *test suite II* at $-10\,°$C and $17\,°$C.

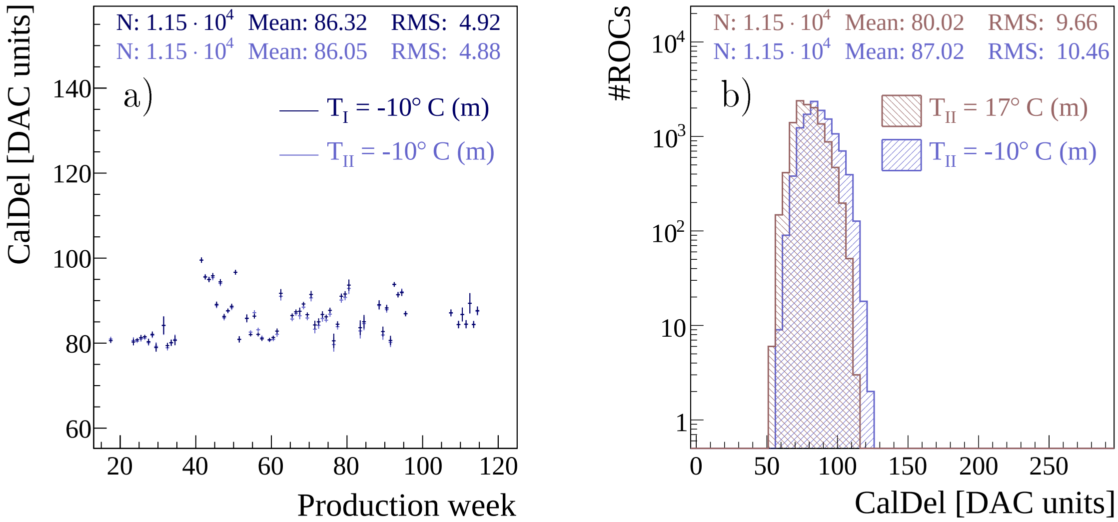

Figure 4.6: Distributions of the *CalDel*-values of the 1.15×10^4 ROCs in the final system: a) The average *CalDel*-value of modules produced in the same week as a function of time, determined at $-10\,°\mathrm{C}$ during *test suite I* (T_I) and *test suite II* (T_{II}). b) Distribution of the *CalDel*-values from *test suite II* at $-10\,°\mathrm{C}$ and $17\,°\mathrm{C}$.

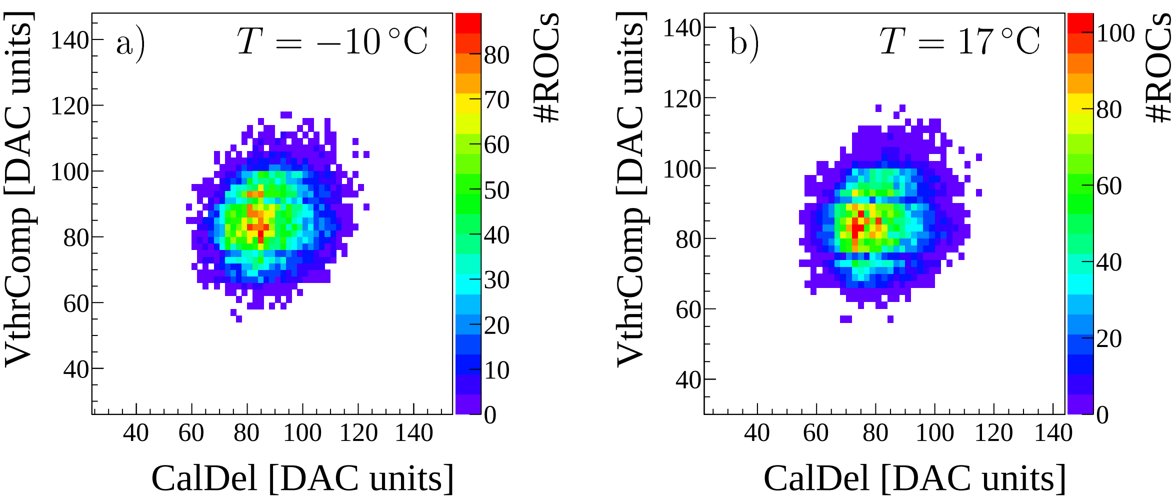

Figure 4.7: Distributions of the (VthrComp, CalDel) working points established in *test suite II* for the 1.15×10^4 ROCs in the final system, at a) $-10\,°\mathrm{C}$ and b) $17\,°\mathrm{C}$.

4.4 Functionality Tests

4.4.1 TBM Test

The TBM on the modules in the third layer will be operated in the single mode, i.e. the readout of all ROCs is sent to the same analogue channel of the TBM. For the modules on the first and second layer the TBM will be operated in a dual mode, where the ROC readouts are split between the two analogue readout channels of the TBM. By checking the length of an empty readout, the TBM test checks that the modules can be operated in both modes. In case of a failed TBM test, a new TBM was placed on top of the faulty TBM and connected to the HDI.

4.4.2 ROC Programmability Test

To check whether the DACs of each ROC can be programmed, the *Vcal* DAC is set to its extreme values 0 and 255. If the difference between the 'last DAC' (see section 3.4) is less than 20 ADC, the ROC is considered to be defective.

4.4.3 Pixel Readout Test

The pixel readout test allows the identification of pixels with a deficient readout. There are several types of defects that can occur in this test: *dead pixel, mask defects* and pixel with a *noisy readout*. To test the functionality of the pixel readout, a calibration signal with a *Vcal*-value of 200 in the low range is sent to an enabled pixel and the analogue signal is read out. During the test, only one pixel at a time is enabled and all other pixels are disabled. It is therefore sufficient to check for any hit in the analogue readout. The test is repeated ten times for each pixel. If less than ten hits were recorded, the pixel is classified as *dead pixel*. If for some reason more than ten hits were counted, the pixel would be considered to have a *noisy readout*. This type of defect, however, is very rare and occurred only in two pixels, in conjunction with other problems. The most serious out of the three types of defects listed in the beginning is the *mask defect*: the purpose of the mask bit is to have a handle on noisy[1]pixels. Such pixels can flood the buffers of a double-column with fake hits—thus inhibiting the chip from working properly. Suppressing the readout of such pixels with the mask bit is therefore crucial. The functionality of the mask bit is tested by trying to readout

[1]in this context referring to *electronic* noise (described in section 4.5.1)

a pixel with the mask bit enabled (i.e. a disabled comparator). Defective mask bits occurred very rarely—in less than 23 pixels in the entire module production.

In the final system, the fraction of *dead pixels* is $2.3 \cdot 10^{-5}$ at $-10\,°C$. Figure 4.8 shows the location of these 1086 dead pixels on the modules of the final system. The modules containing the *mask defects* and the pixels with a *noisy readout* did not comply with the requirements for use in the final system.

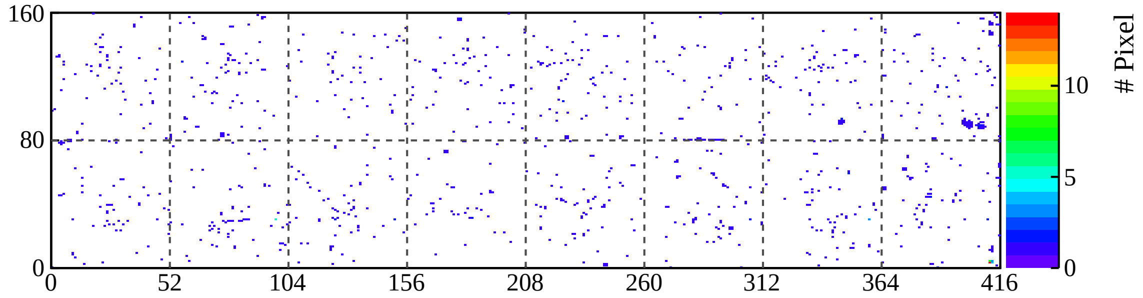

Figure 4.8: Superposition of the 768 modules in the final system (47.9×10^6 tested pixels) showing the location of the 1086 dead pixels, measured at $-10\,°C$. The resulting fraction of dead pixel in the final system is $2.3 \cdot 10^{-5}$.

4.4.4 Bump Bonding Test

To check the quality of the bump bond connecting the sensor pixel to the PUC, several methods were proposed and discussed in [40]. It is possible to mimic a hit in the sensor pixel by diverting the calibration signal to a pad on the ROC surface instead of the preamplifier. That in turn induces a charge in the sensor through the air capacitance between the ROCs and sensor. In principle, missing bump bonds can be identified by measuring the *Vcal*-threshold of the pixel. The shortcoming of this method is that, for large enough amplitudes, a parasitic capacitance between the voltage calibration line and the preamplifier leads to cross-talk, making it impossible to distinguish between bonded and unbonded bumps. The workaround for this problem is to measure two *Vcal*-thresholds for each pixel, one for charge injection through the sensor and one through the parasitic cross-talk. The difference of the two thresholds allows identification of bump bonds that are of poor quality or missing: if the two *Vcal*-thresholds differ less than five DAC values, the bump bond is considered to be defective. To ensure an optimum discrimination, the *VthrComp* DAC is set as low as possible, but still high enough to detect the cross-talk threshold.

In the final system, the fraction of *defective bump bonds* is $1.3 \cdot 10^{-4}$. Figure 4.9 shows the location of these 6289 defective bump bonds on the modules of the final system.

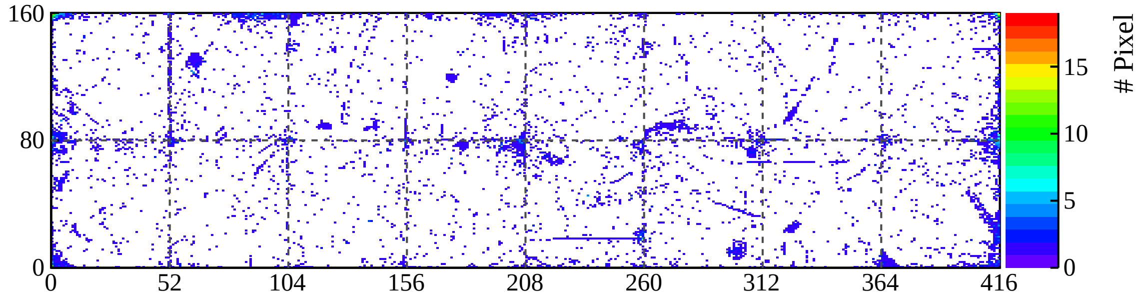

Figure 4.9: Superposition of the 768 modules in the final system (47.9×10^6 tested pixels) showing the location of the 6289 defective bump bonds, measured at $-10\,°C$. The resulting fraction of defective bump bonds in the final system is $1.3 \cdot 10^{-4}$.

4.4.5 Trim Bits Test

To adjust the thresholds of all pixels on a ROC, the threshold of each individual pixel can be fine-tuned by setting the four trim bits to a value between 0 and 15, and adjusting the *Vtrim* DAC (see section 4.6). By default all four trim bits are enabled (highest threshold). To test their functionality, the *Vcal*-threshold is measured first with all trim bits enabled and then by turning off each trim bit separately. A trim bit is considered to be defective, if the trimmed threshold has decreased less than two DAC units with respect to the untrimmed threshold. To obtain similar threshold differences from all four trim bit tests, the impact of the trim value on the threshold is scaled by setting *Vtrim* accordingly. Defective trim bits occurred very rarely. In total, 35 pixels with one or more defective trim bits were found in the entire module production. The defects are shared equally among the four trim bits (see Table 4.1).

In the final system, the fraction of pixels with one or more *defective trim bits* is $6.3 \cdot 10^{-7}$ at $-10\,°C$ (which corresponds to a total of 30 pixels with defective trim bits).

Table 4.1: Defective trim bits, measured at −10 °C and 17 °C.

Trim Bit	−10 °C	17 °C
	# Pixel	# Pixel
1	11	11
2	12	11
3	12	11
4	12	9
Any	35	30

4.4.6 Address Decoding

The pixel address is encoded in five clock cycles, each able to be set to six different analogue levels. Two clock cycles contain the double-column index and the other three, the index within the double-column. A range for each address level is extracted from the distribution of the address levels of all pixels on a ROC. The address decoding is tested by decoding the generated address of a pixel in the analogue readout and comparing it to the physical position of the only enabled pixel on the ROC. Only a few modules exhibited problems with address decoding. The problems generally occurred in combination with more severe defects and often concerned the whole ROC.

In the final system, the fraction of pixels, that exhibited *address decoding problems* at −10 °C, is $1.0 \cdot 10^{-7}$ (which corresponds to 5 pixels with address decoding problems). This fraction is somewhat enhanced at 17 °C, where contributions from two additional ROCs give rise to address decoding errors in 185 pixels.

4.5 Performance Tests

4.5.1 Noise Measurements

The electronic noise[1] in a pixel leads to a smearing of the *Vcal*-threshold. Assuming a Gaussian distribution of the noise, the response efficiency as a function of the amplitude of the calibration signal is described by an error function (in the absence of noise

[1] major sources of noise are fluctuations in the sensor leakage current or biasing networks (parallel) and noise in the amplifier system (series)

this would be a simple step function, that changes from zero to full efficiency at the threshold). Figure 4.10 shows the readout efficiency as a function of the calibration voltage. The so-called "S-curve" is fitted with an error function. For better precision, the number of injected calibration signals is increased in the window around the threshold. The width of the error function is proportional to the noise and the threshold is defined as the position where the response efficiency reaches 50%. The conversion of the *Vcal* value into a calibration voltage is necessary, as sometimes a higher *Vcal* value results in a lower calibration voltage. Here, the dependence of the calibration voltage on the value of *Vcal* was extracted from a measurement for one ROC.

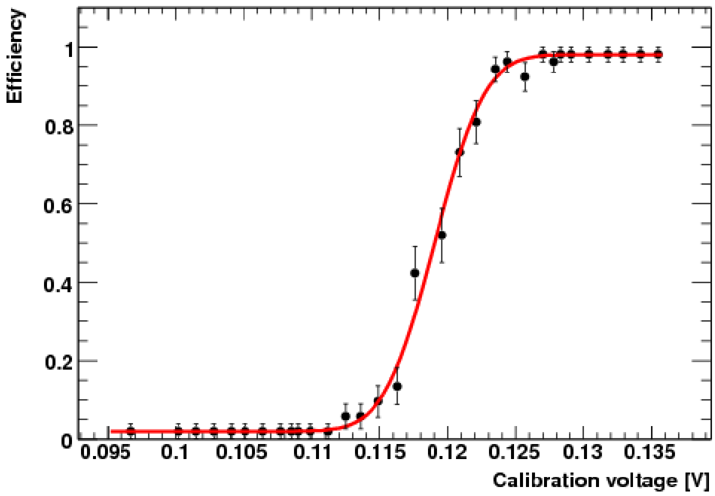

Figure 4.10: S-curve fit with an error function. The noise is given by the width of the S-curve, the threshold is defined by the calibration voltage at 50 % efficiency.

The average pixel noise in the modules of the final system is illustrated in Figure 4.11. Being twice and four times the size of a standard pixel, the pixels on the three inner edges and on the corners of a ROC, respectively, exhibit higher noise levels. At $-10\,^{\circ}$C, the average noise of all pixels in the final system amounts to $155\,e^-$, compared to $184\,e^-$ on the edge pixels and $228\,e^-$ on the corner pixels (see Figure 4.12).

The results from the S-curve method were confirmed by a measurement of the pulse height distribution at a fixed signal amplitude. The RMS of this distribution depends on the noise in a pixel. Taking into account the RMS of the black level distribution, the RMS of the pulse height distribution is converted into electrons using the gain and pedestal from the pulse height calibration. The noise determined with the pulse height measurement lies on average 20 e^- above the noise from the S-curve measurement. The spread in noise difference of the two methods is about 14 e^- [37].

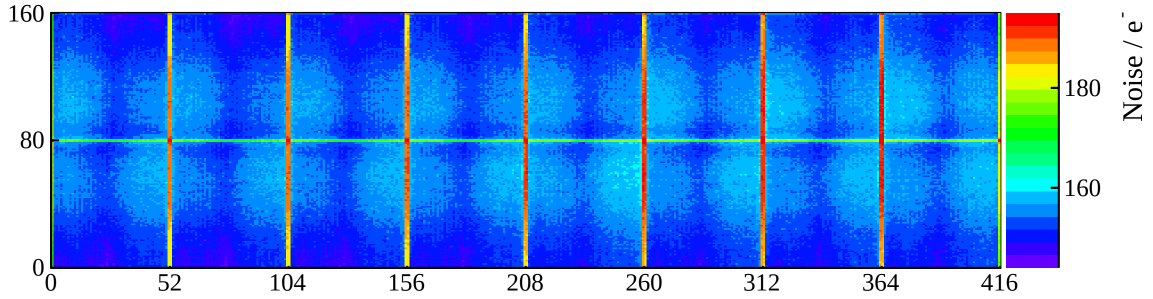

Figure 4.11: Average pixel noise at −10 °C in the modules of the final system. The mean noise of the corner pixels lies outside the z-range and amounts to approximately 230 e^-.

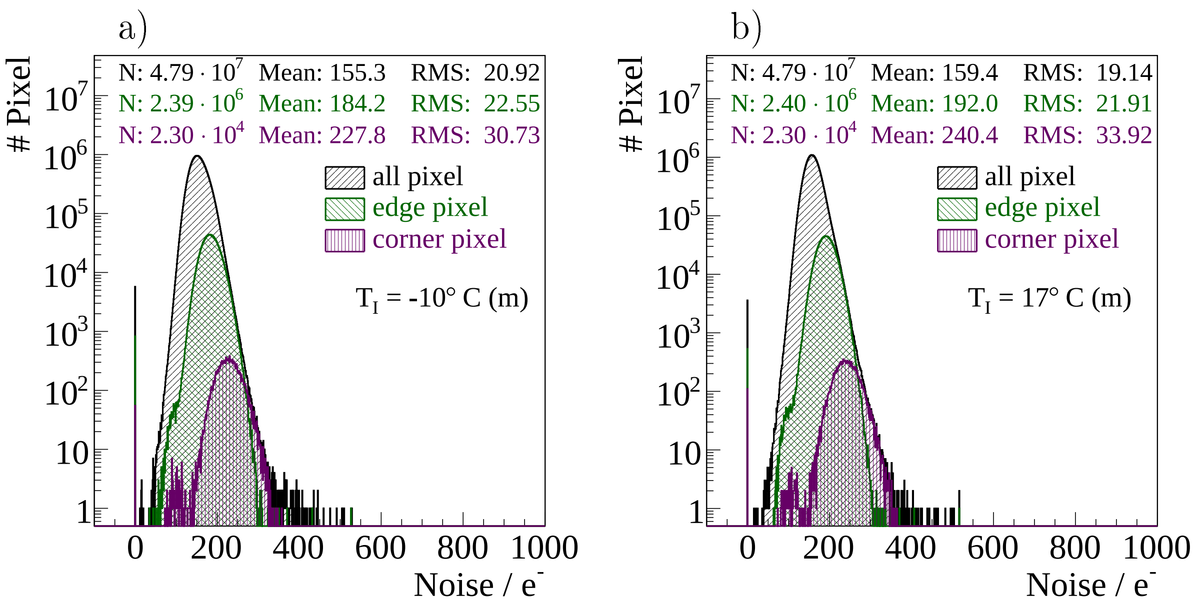

Figure 4.12: Distribution of the noise levels of the 47.9×10^6 pixels in the final system, measured during *test suite I* at a) −10 °C and b) 17 °C.

4.5.2 Sensor Leakage Current

The thermal excitation of electron-hole pairs in the silicon sensor gives rise to a leakage current. The dependency of the leakage current and the applied reverse bias voltage is shown in Figure 4.13 for a module with a flawless sensor. The curve that describes the current-voltage characteristics of a module is called the "IV-curve" and can be divided into three regions: 1) below the sensor depletion voltage the current increases with the square root of the voltage, 2) in the plateau region the current increment is very small and 3) at very high voltages a breakdown occurs. Beyond this point, a hard breakdown can occur which will destroy the device. Defects in the sensor lead to a deviation from the typical IV-curve. The IV characteristics thus provide a powerful tool to identify sensor imperfection and problems in the fabrication process (such as scratches and spikes).

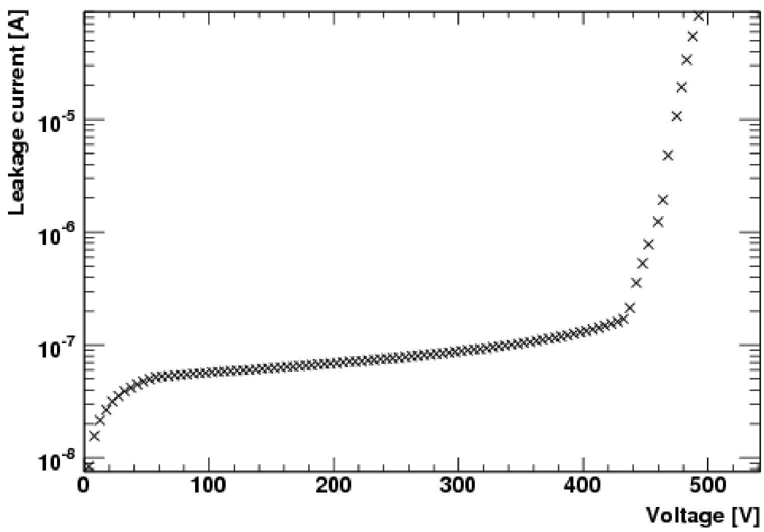

Figure 4.13: Sensor IV-curve.

The IV-curve measurement for a module starts at 0 V and is increased in steps of 5 V. The current is measured 5 s after incrementing the voltage. The procedure is stopped when the leakage current exceeds 100 μA or the voltage reaches 600 V. Repeated measurements for the same module yielded an accuracy of the measured current of $2.1 \cdot 10^{-3}\,\mu$A. Figures 4.14 and 4.15 illustrate the IV-curves of a few selected modules and half-modules, measured at $-10\,°$C and $17\,°$C respectively.

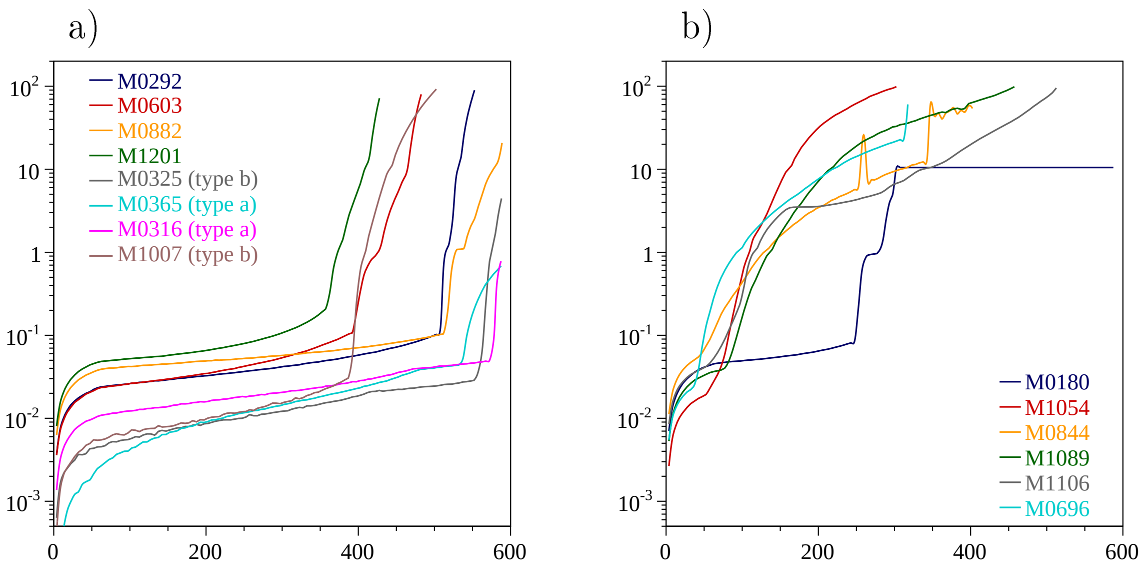

Figure 4.14: Sensor IV-curves measured at $-10\,^{\circ}$C for a) 'good' sensors and b) rejected sensors. The supplements "type a" and "type b" indicate the two different types of half-modules (right-handed or left-handed version).

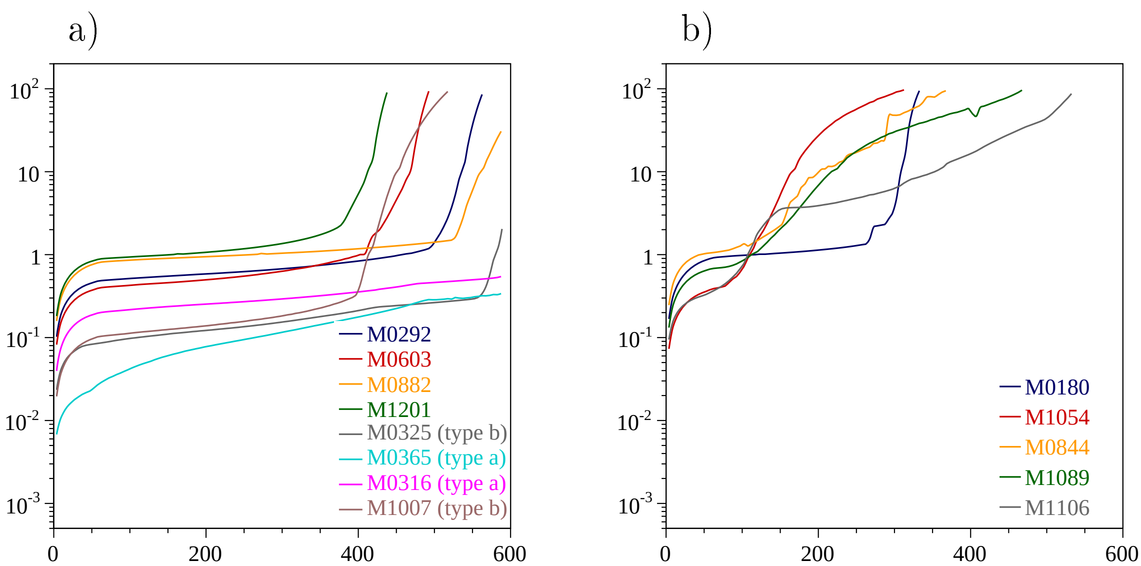

Figure 4.15: Sensor IV-curves measured at $17\,^{\circ}$C for a) 'good' sensors and b) rejected sensors. The supplements "type a" and "type b" indicate the two different types of half-modules (right-handed or left-handed version).

4.6 Calibration Tests

4.6.1 Vcal Calibration

The correlation between ionisation charge and the injected calibration signal for a given *Vcal* value was first investigated in a beam-test with 300 MeV pions at PSI in 2005 [41]. The ionisation charge in a single hit pixel can be represented by a Landau distribution. By varying the incident angle of the pions, the position of the Landau peak can be shifted and with it, the expected ionisation charge (known from [42]). From this, the dependence of the calibration voltage on the value of *Vcal* can be established. The beam-test results showed that the ionisation charge can be expressed as a linear function of the injected calibration signal. The results also indicated that the slope of this linear dependency varied between ROCs, ranging from $51 - 69$ electrons per *Vcal* unit. The average ionisation charge per *Vcal* unit was found to be $61.1\,e^-/$ *Vcal* DAC with an RMS of $5.5\,e^-/$ *Vcal* DAC [37].

A second study with a variable X-ray source embedded in the module testing procedure was carried out for 834 modules. The X-rays were produced by a primary Americium-241 source exciting the $K\alpha_1$ emission lines of a selectable target (Cu, Rb, Mo, Ag, Ba, Tb). Each module was calibrated with the Molybdenum and silver target, providing a photon energy of 17.48 keV and 22.16 keV respectively. This corresponds to an expected ionisation charge of approximately $5000\,e^-$ and $6000\,e^-$ respectively. As a first step, the *VthrComp*-threshold of each ROC was determined by randomly reading out the fully enabled module several thousand times and counting the hits on each ROC. For this measurement the clock sent to a module was stopped. This artificially stretches the bunch crossing and thus increases the probability of finding a hit in the corresponding bunch-crossing. The *VthrComp*-threshold value is extracted from the value of the error function fit of the threshold curve at 50%. The comparator thresholds of all ROCs are set to the resulting *VthrComp* and the *Vcal*-threshold is measured for each pixel. The average *Vcal* value then corresponds to the ionisation charge for the given X-ray energy. The two measurement points were fitted with a linear function, and the average slope was found to be $65.5\,e^-/$ *Vcal* DAC and a mean offset at zero *Vcal* of $-410\,e^-$. Taking into account measurement uncertainties, the RMS is of the order of $5\,e^-/$ *Vcal* DAC [37].

A subset of 69 modules was also tested in a more elaborate measurement that, besides the Molybdenum and silver, also included a Barium target. This provided an additional measurement point with a photon energy of 32.19 keV and an ionisation charge of 9000 e^-. The conclusion from this study was that the calibration constant varies less for ROCs from the same ROC wafer or within the same module. In those cases the RMS is 2.8 e^-/ *Vcal* DAC and 2.9 e^-/ *Vcal* DAC respectively. Therefore, the suggestion was to use a calibration constants averaged over modules (or wafers).

In the scope of this work, a *conversion factor* of 65 e^- per *Vcal* unit in the low range and 455 e^- per *Vcal* unit in the high range applies.

4.6.2 Pulse Height Calibration

The ionisation signal in a pixel is represented by a pulse height (PH) expressed in Analogue to Digital Converter (ADC) units. To associate the pulse height with the collected ionisation charge, the pulse height is recorded as a function of the injected calibration signal. With the results from the *Vcal* calibration, the signal amplitude in *Vcal* units can then be converted into electrons. Ten pulse height measurements at each of the following *Vcal* values are taken and averaged: 50, 100, 150, 200, 250 in the low range and 30, 50, 70, 90, 200 in the high range. Figure 4.16 shows the results of such a pulse height calibration measurement. Before saturating at approximately 120 *Vcal* units in the high range, the curve shows a linear behaviour (except in the very low range, see section 4.7.1).

The expected ionisation charges in a pixel will typically be below 30000 e^- [38], corresponding to *Vcal* values well below the saturated region. A linear fit in the respective range is therefore adequate to describe the dependency of the pulse height and the ionisation charge on a pixel. The two parameter extracted from this fit are the gain and pedestal: the *gain* is determined by the slope of the fit and the *pedestal* represents the *Vcal* offset corresponding to a pulse height of zero. The gain is expressed in units of ADC/DAC, whereas the pedestal is generally converted from *Vcal*-DAC units to electrons, using a conversion factor of 65 e^- per *Vcal* unit in the low range (see section 4.6.1).

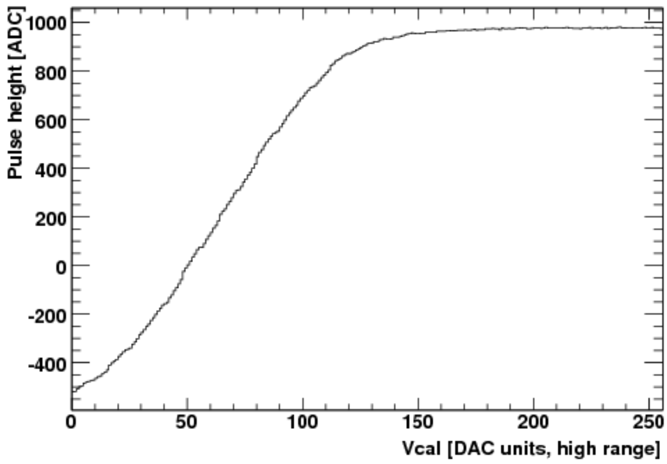

Figure 4.16: Analogue pulse height as a function of *Vcal*.

4.7 Performance Optimisation

4.7.1 Optimisation of the Pulse Height Calibration

The pulse height calibration curve over the entire *Vcal*-range is fitted with the hyperbolic tangent function given in equation 4.1. Towards the low *Vcal*-range, non-linear behaviour can occur. The degree of non-linearity is contained in the parameter p_1 of the function given in equation 4.1. A value of $p_1 \approx 1$ implies an almost linear behaviour down to the lowest *Vcal* values. With increasing p_1 the pulse height starts to saturate in the low *Vcal*-range. A method to optimise the linearity of the pulse height calibration curve was developed at PSI and is described in detail in [38].

$$y = p_3 + p_2 \ \tanh(p_0 \ x - p_1) \tag{4.1}$$

The linearity in the low range is optimised by increasing the value of the *Vsf*-DAC. Since the *Vsf*-DAC also affects the digital current, the optimisation terminates at the *Vsf*-value, for which $p_1 < 1.4$ or $I_{dig} > 5\,\mu$A. The distributions of the established *Vsf*-values are illustrated in Figure 4.17.

The linearity in the high range is optimised by adjusting the ADC range of the pulse heights. In a first step, the absolute ADC range is adjusted from -1000 to $+1000$ ADC units by setting the *VIbias_PH*-DAC to the appropriate value. Then *VoffsetOP*-DAC and *VOffsetR0*-DAC are used to shift the ADC range of the pulse height. As shown in [38], the range can always be moved to the required level by adjusting *VoffsetOP*, if *VOffsetR0* is set to above 100 DAC units. Taking into account temperature and pixel-to-pixel variations, *VOffsetR0* is set to 120. The pulse height curve is then shifted to the target ADC range by adjusting *VoffsetOP*. The distributions of the *VIbias_PH*-values and *VoffsetOP*-values established in this procedure are illustrated in Figures 4.18 and 4.19. The optimisation of the three DACs has no influence on the address levels, which have already been adjusted with *Ibias_DAC* in the start-up test (see section 4.3.2). In a similar way, *Ibias_DAC* has a negligible effect on the pulse height.

Figures 4.20—4.25 illustrate the effect of the optimisation on the gain, pedestal and parameter p_1. For each parameter, there are: i) two module maps—for before and after the optimisation—each showing the average value per pixel at $-10\,°$C; ii) two plots per parameter—for $-10\,°$C and $17\,°$C—showing a direct comparison of the parameter distributions before and after the optimisation. The parameter distributions are displayed using a linear scale. Plots showing gain, pedestal and parameter p_1 on a logarithmic scale can be found in section 5.2.3.2 on performance deficiencies.

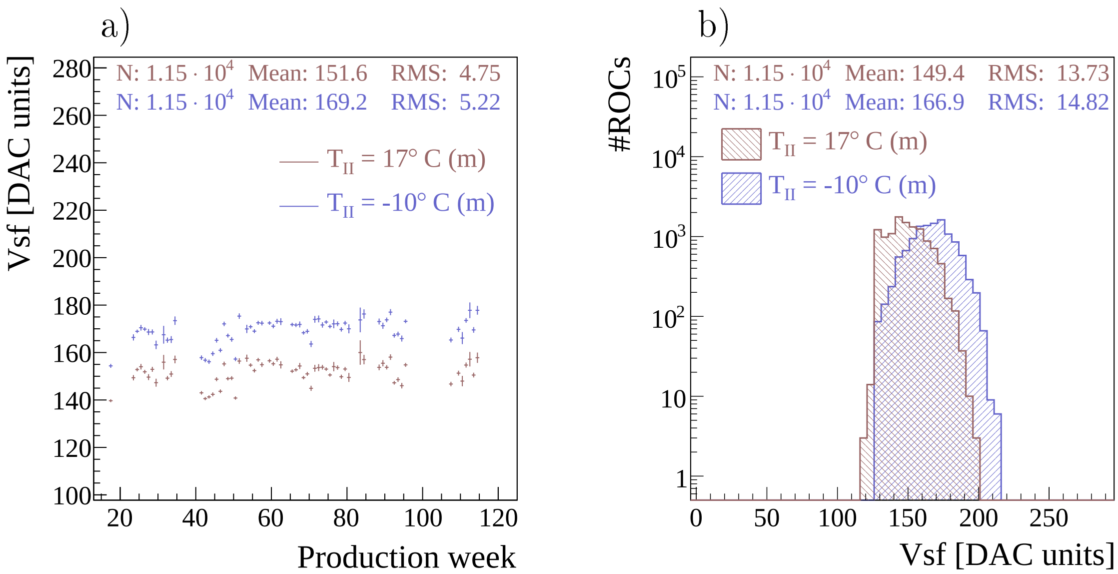

Figure 4.17: Distributions of the *Vsf*-values of the 1.15×10^4 ROCs in the final system: a) The average *Vsf*-value of modules produced in the same week as a function of time, determined at $-10\,°$C and $17\,°$C during *test suite II*. b) Distribution of the *Vsf*-values from *test suite II* at $-10\,°$C and $17\,°$C.

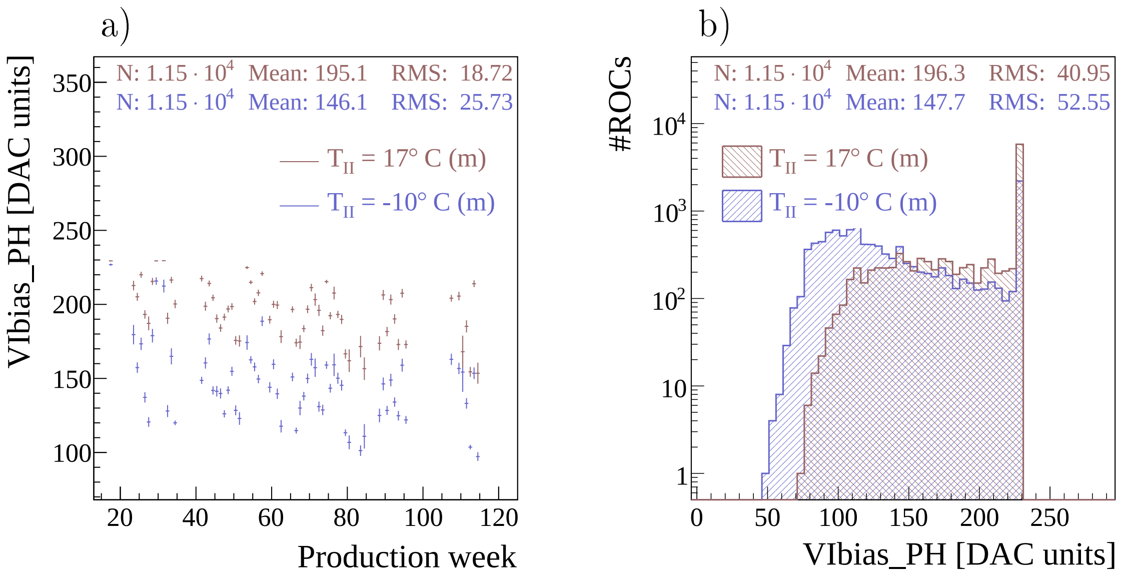

Figure 4.18: Distributions of the *VIbias_PH*-values of the 1.15×10^4 ROCs in the final system: a) The average *VIbias_PH*-value of modules produced in the same week as a function of time, determined at $-10\,°C$ and $17\,°C$ during *test suite II*. b) Distribution of the *VIbias_PH*-values from *test suite II* at $-10\,°C$ and $17\,°C$.

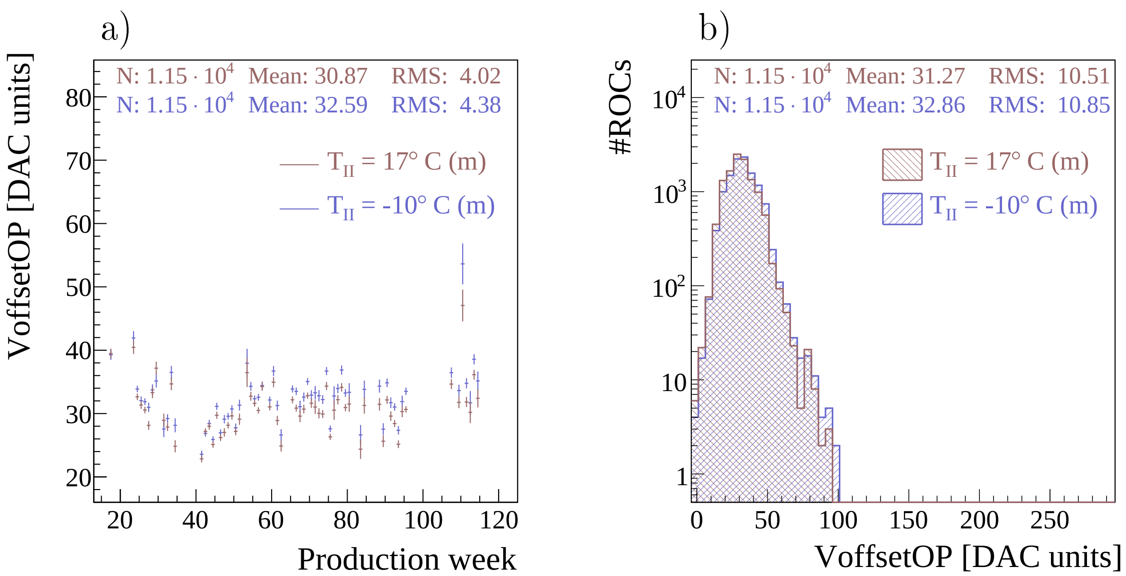

Figure 4.19: Distributions of the *VoffsetOP*-values of the 1.15×10^4 ROCs in the final system: a) The average *VoffsetOP*-value of modules produced in the same week as a function of time, determined at $-10\,°C$ and $17\,°C$ during *test suite II*. b) Distribution of the *VoffsetOP*-values from *test suite II* at $-10\,°C$ and $17\,°C$.

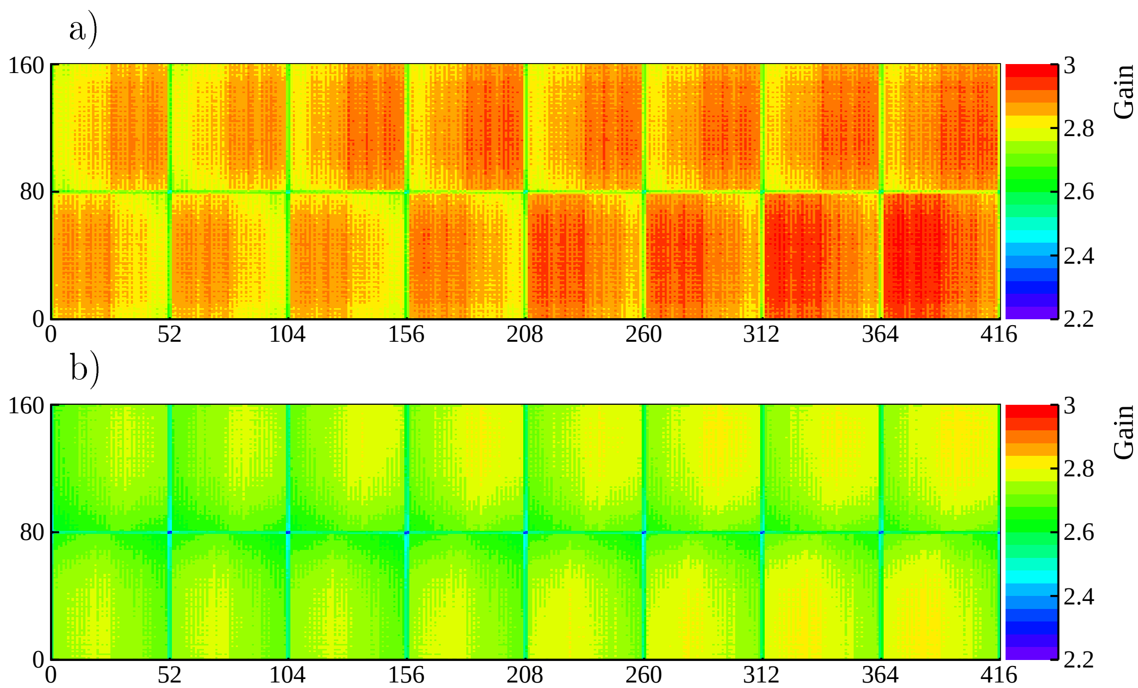

Figure 4.20: Average pixel gain in the modules of the final system measured at −10 °C, a) before and b) after the optimisation.

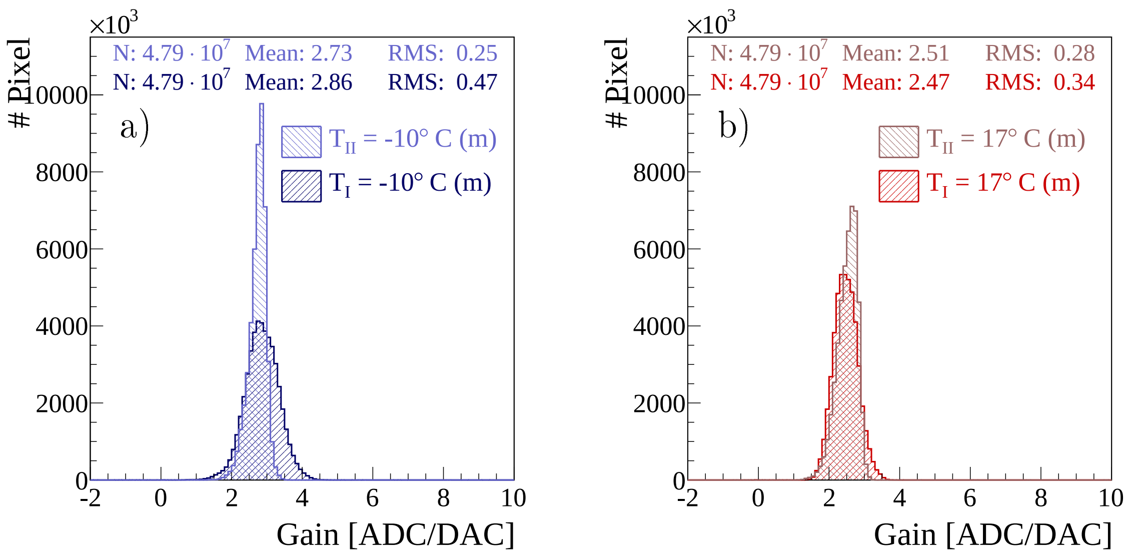

Figure 4.21: Comparison of the gain distributions of the 47.9×10^6 pixels in the final system before (T_I) and after (T_{II}) the optimisation, measured at a) −10 °C and b) 17 °C.

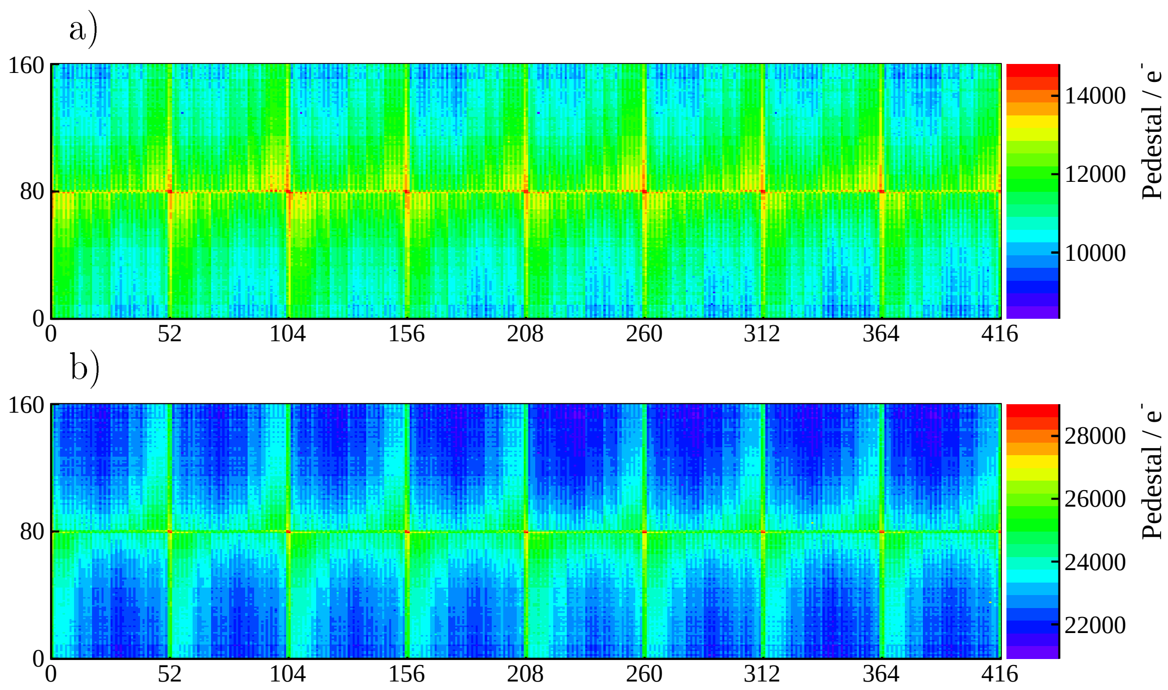

Figure 4.22: Average pixel pedestal in the modules of the final system measured at −10 °C, a) before and b) after the optimisation.

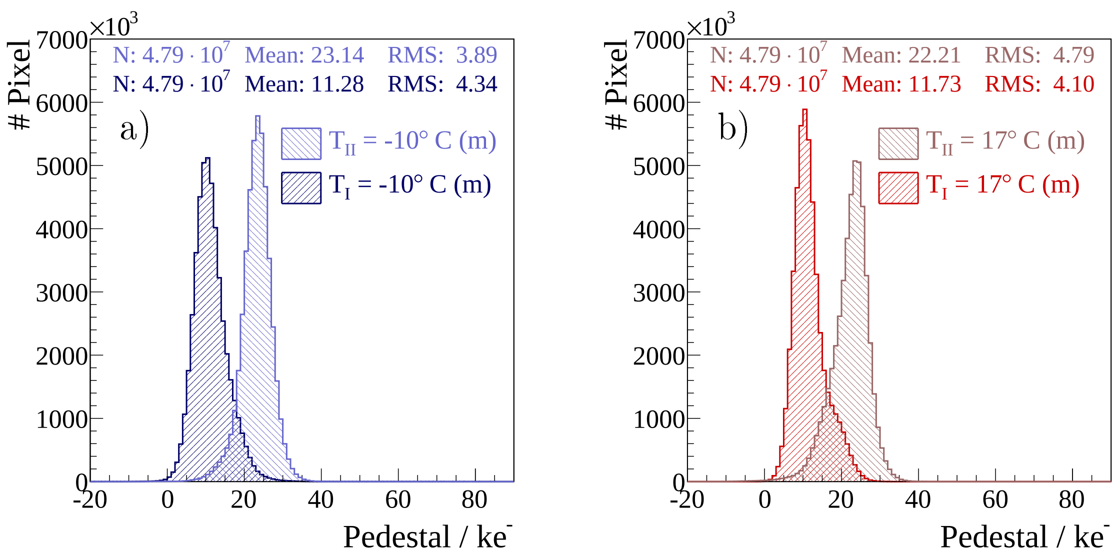

Figure 4.23: Comparison of the pedestal distributions of the 47.9×10^6 pixels in the final system before (T_I) and after (T_{II}) the optimisation, measured at a) −10 °C and b) 17 °C.

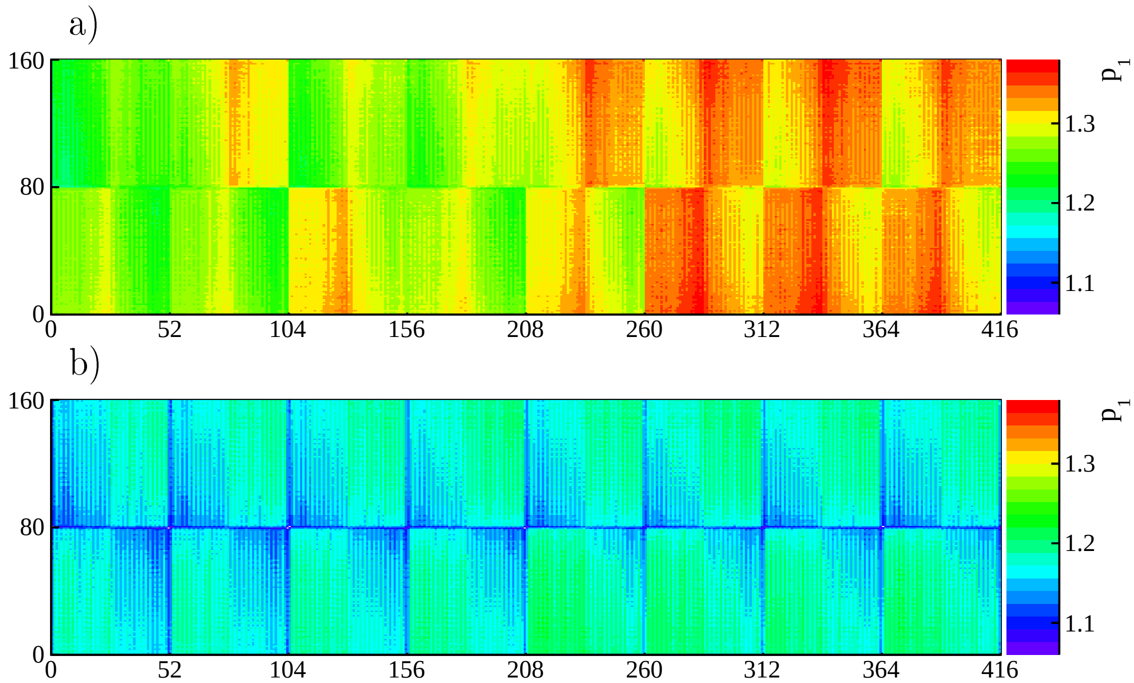

Figure 4.24: Average parameter p_1 in the modules of the final system measured at $-10\,°$C, a) before and b) after the optimisation.

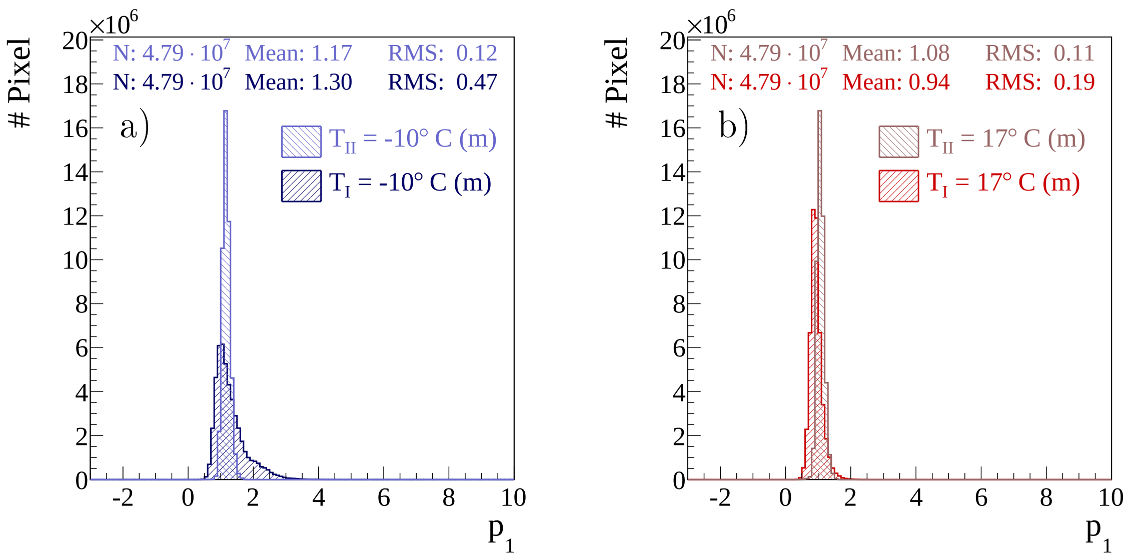

Figure 4.25: Comparison of the parameter p_1 distributions of the 47.9×10^6 pixels in the final system before (T_I) and after (T_{II}) the optimisation, measured at a) $-10\,°$C and b) $17\,°$C.

4.7.2 Threshold Unification (Trimming)

If the comparator thresholds are adjusted with *VthrComp* only, the mean spread of the physical thresholds per ROC is 309 e^- due to local transistor mismatches. By means of four trim bits and together with the *Vtrim* DAC, these physical thresholds can be unified. By default all four trim bits are enabled and the threshold of a pixel can be lowered by turning off one or more trim bits. The trim bits can take a value between 0 and 15 and the corresponding threshold difference can be scaled with *Vtrim*.

The goal of the trim algorithm, developed at PSI [37; 40], is to set all comparator thresholds in such a way, that they correspond to the same *Vcal*-threshold. By default the trimming was performed at a target threshold of 60 *Vcal* units. The procedure to unify the thresholds of all pixels on a ROC comprises the following three steps:

1. **VthrComp**: The global threshold is set to the *VthrComp* value of the pixel with the highest comparator threshold, i.e. the threshold for measuring a calibration signal injected with the selected *Vcal* value.

2. **Vtrim**: The *Vcal*-threshold of all pixels is measured. The trim voltage *Vtrim* is set by disabling all four trim bits of the pixel with the highest *Vcal*-threshold and increasing *Vtrim*, until the *Vcal*-threshold is lowered to the selected target threshold. The distribution of the *Vtrim*-values is illustrated in Figure 4.26.

3. **Trim Bits**: For each pixel the trim bit value is adjusted in a way, that the *Vcal*-threshold of the pixel differs least from the selected target threshold. The distribution of the value of the trim bits is illustrated in Figure 4.27.

All thresholds established by trim algorithm are *timing independent* thresholds. When determining *VthrComp* and *Vtrim*, outliers deviating more than five times the root mean square (RMS) from the mean are neglected. In case of *VthrComp* an upper limit applies, above which the ROC is no longer functional. A comparison of the threshold distributions before and after the trimming is illustrated in Figures 4.28 and 4.29. The threshold distributions are displayed using a linear scale. Plots showing the threshold distribution after trimming on a logarithmic scale can be found in section 5.2.3.2 on performance deficiencies.

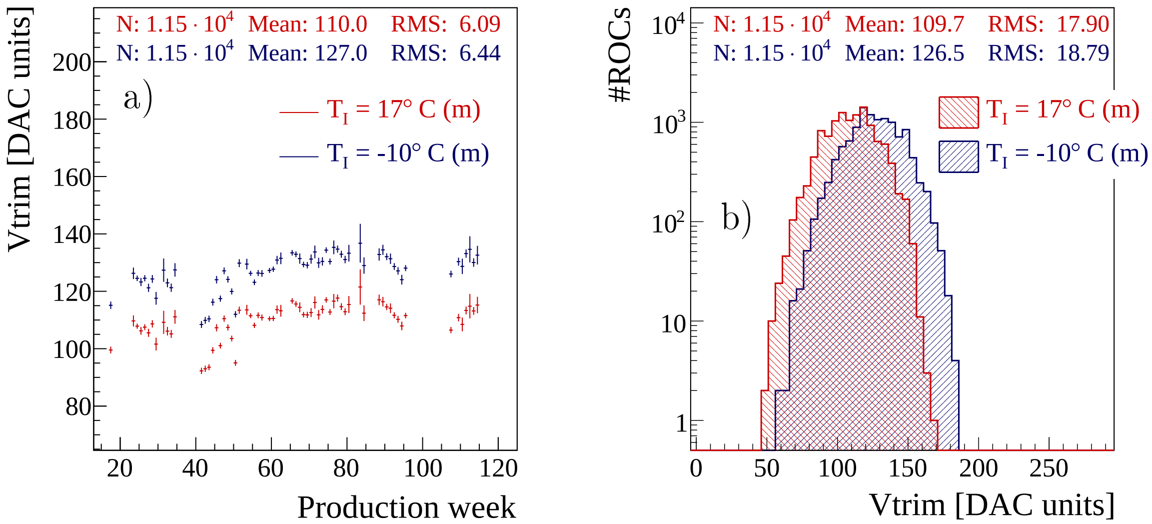

Figure 4.26: Distributions of the *Vtrim*-values of the 1.15×10^4 ROCs in the final system: a) The average *Vtrim*-value of modules produced in the same week as a function of time, determined at $-10\,°C$ during *test suite I* (T_I) and *test suite II* (T_{II}). b) Distribution of the *Vtrim*-values from *test suite II* at $-10\,°C$ and $17\,°C$.

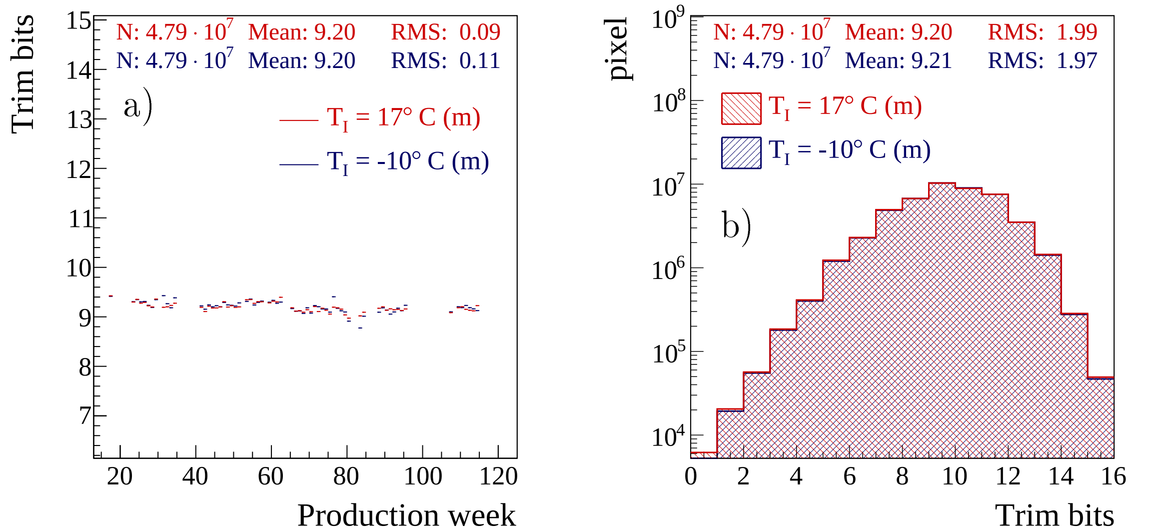

Figure 4.27: Distributions of the trim bits values of the 47.9×10^6 pixels in the final system: a) The average trim bits value of modules produced in the same week as a function of time, determined at $-10\,°C$ during *test suite I* (T_I) and *test suite II* (T_{II}). b) Distribution of the trim bits values from *test suite II* at $-10\,°C$ and $17\,°C$.

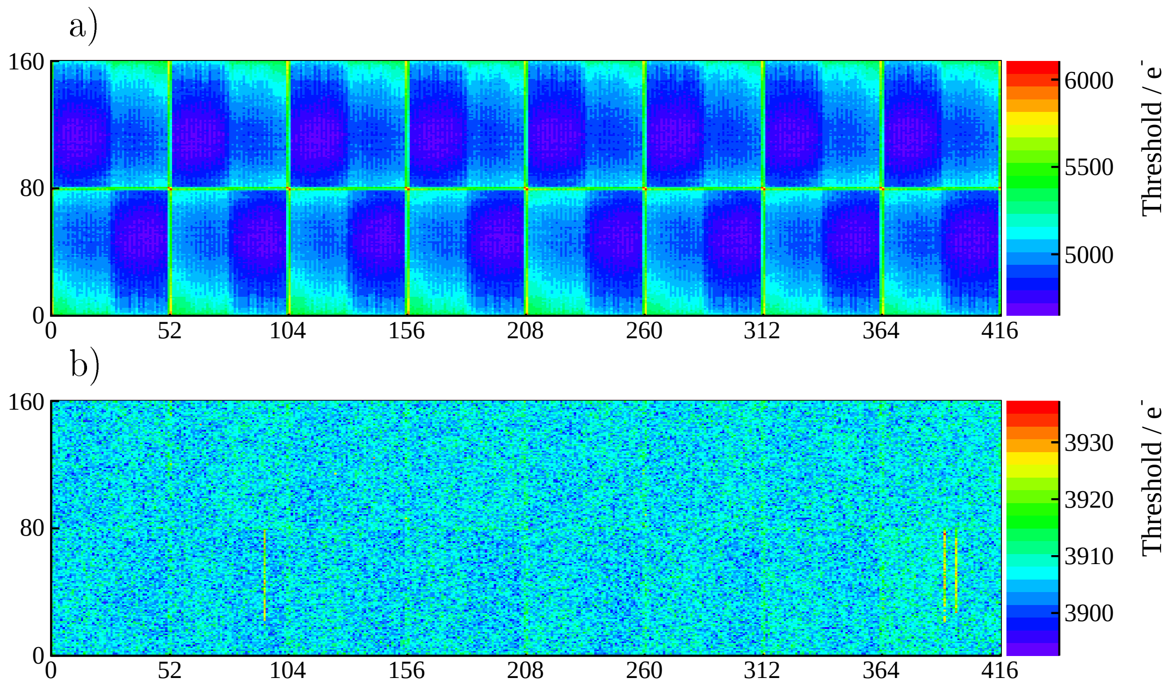

Figure 4.28: Average threshold in the modules of the final system measured at −10 °C, a) before and b) after threshold unification with the trim algorithm.

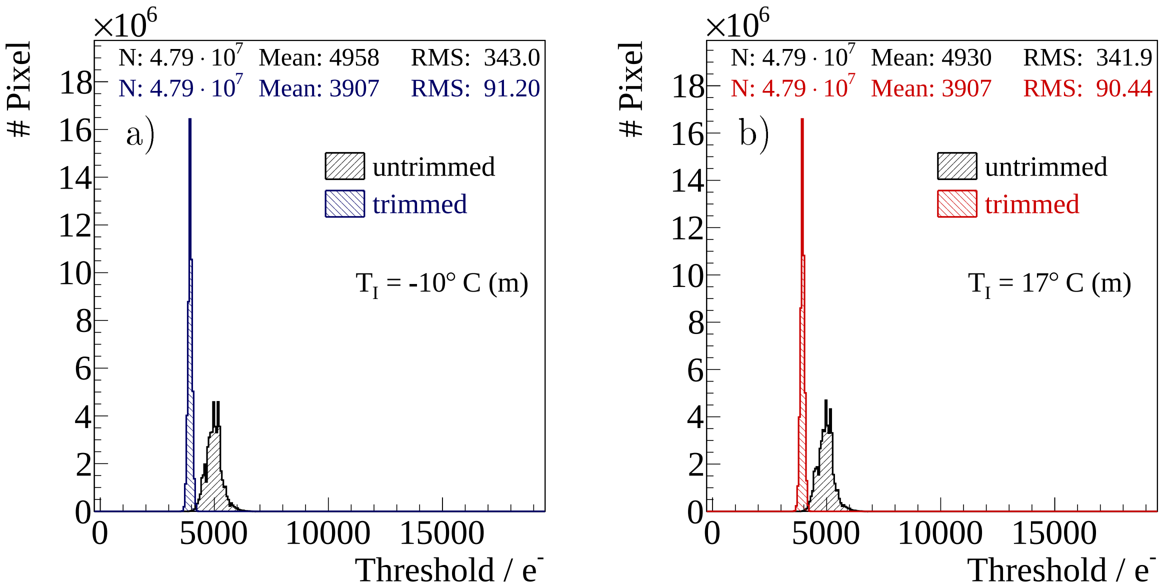

Figure 4.29: Comparison of the threshold distributions of the 47.9×10^6 pixels in the final system before and after trimming, measured at a) −10 °C and b) 17 °C during *test suite I*.

4.8 Results

This section summarises the results of the functionality, calibration and performance tests, and the performance optimisation described in the previous sections. In total, 971 fully assembled detector modules entered the process of module qualification (848 full and 123 half modules). Unless mentioned otherwise, the results presented in this section include only the 768 modules selected for implementation in the final detector system (672 full and 96 half modules), and therefore comprise of the results of the test and optimisation algorithms performed in 47.9×10^6 pixels and 1.15×10^4 ROCs.

Pixel defects: The only significant contributions to the number of pixel defects in the final detector system arise from defective bump bonds and dead readout pixels (see Table 4.2).

Table 4.2: Number and ratio of different types of defects in the 47.9×10^6 pixels in the final detector system, measured at −10 °C and 17 °C.

	−10 °C		17 °C	
	# Pixel	Ratio	# Pixel	Ratio
Dead pixel	1086	$2.3 \cdot 10^{-5}$	1141	$2.4 \cdot 10^{-5}$
Defective bump bond	6289	$1.3 \cdot 10^{-4}$	6289	$1.3 \cdot 10^{-4}$
Mask bit defect	0	0	0	0
Trim bit defect	30	$6.3 \cdot 10^{-7}$	26	$5.4 \cdot 10^{-7}$
Address decoding problem	5	$1.0 \cdot 10^{-7}$	185	$3.9 \cdot 10^{-6}$

On the modules implemented in the final detector system, the fraction of defective bump bonds is $1.3 \cdot 10^{-4}$ and the fraction of dead pixels amounts to $2.3 \cdot 10^{-5}$. In the entire module production, less than 23 pixels with a defective mask bit and 35 pixels with one or more defective trim bits were found, and pixels of which the address could not be decoded correctly, were restricted to a few modules and generally accumulated on the same ROCs in combination with other malfunctions. In the final detector system the fraction of pixels with defective trim bits and the fraction of pixels with address decoding errors are both of the order of 10^{-7}. Modules containing a pixel with a defective mask bit were not allowed in the detector system. The above fractions were obtained from measurements at −10 °C.

DAC settings: To achieve an optimum ROC performance, nine DAC parameters were adjusted separately for each ROC. The mean value and the RMS of the dynamically optimised DAC parameters of the ROCs in the final system are shown in Table 4.3. The table also shows the mean value and RMS of the trim bits values established by the trim algorithm in the pixels in the final system.

Table 4.3: Dynamically optimised DAC parameters of the 1.15×10^4 ROCs in the final detector system, established in *test suite I* and *test suite II* at −10 °C and 17 °C.

	Mean				RMS			
	−10 °C		17 °C		−10 °C		17 °C	
	I	II	I	II	I	II	I	II
Vana	156	154	141	141	14	14	13	12
Dacgain	164	215	154	206	40	34	36	33
Ibias_DAC	85	112	93	125	20	18	22	20
VthrComp	87	85	89	84	10	9	10	9
CalDel	87	87	79	80	10	10	9	10
Vsf	-	168	-	151	-	15	-	14
VIbiasPH	-	148	-	197	-	53	-	42
VoffsetOP	-	32	-	31	-	11	-	10
Vtrim	126	-	109	-	19	-	18	-
Trim Bits	9	-	9	-	2	-	2	-

ROC performance: A comparison of different performance characteristics averaged per ROC and per double-column (DC) is illustrated in Figures 4.30—4.35. The mean values and the mean spreads of the performance parameters are summarised in

Table 4.4. The following results have been obtained from measurements at −10 °C for the modules in the final detector system:

- The average threshold variation before trimming is 309 e^- (277 e^-) per ROC (DC). With the trim algorithm the mean threshold variation can be reduced to 87.6 e^- (86.9 e^-) per ROC (DC), see Figures 4.30 and 4.31.
- The mean noise on a ROC amounts to 155 e^-. The average spread of the noise is 18.5 e^- (17.4 e^-) per ROC (DC), see Figure 4.32.
- After optimising the DAC parameters to increase the linear range of the pulse height calibration curve, the mean parameter p_1 of the hyperbolic tangent fit is 1.2 and the average spread per ROC (DC) amounts to $4.2 \cdot 10^{-2}$ ($2.8 \cdot 10^{-2}$), see Figure 4.33. At −10 °C, the value of p_1 could not be lowered to below 2 in 6544 pixels. This relatively large number can mainly be attributed to a failed p_1 optimisation in two ROCs on two different modules (see section 5.2.3.2). At 17 °C, there were only 226 pixels with $p_1 > 2$ after the optimisation.
- Using the optimised DAC settings gives the following results for the linear fit parameters of the pulse height calibration curve (see Figures 4.34 and 4.35): the mean *relative* spread of the gain[1] amounts to $2.7 \cdot 10^{-2}$ ($1.9 \cdot 10^{-2}$) per ROC (DC); the average pedestal spread per ROC (DC) is 1.5 ke^- (0.91 ke^-); the mean gain is 2.7 ADC/DAC and the mean pedestal is 23 ke^-. Since the value of the pedestal is affected by the shift of the ADC range of the pulse height in the optimisation of the linear range, the mean pedestal is much lower before the optimisation and amounts to 11 ke^-.

[1]the spread of the gain distribution divided by its mean

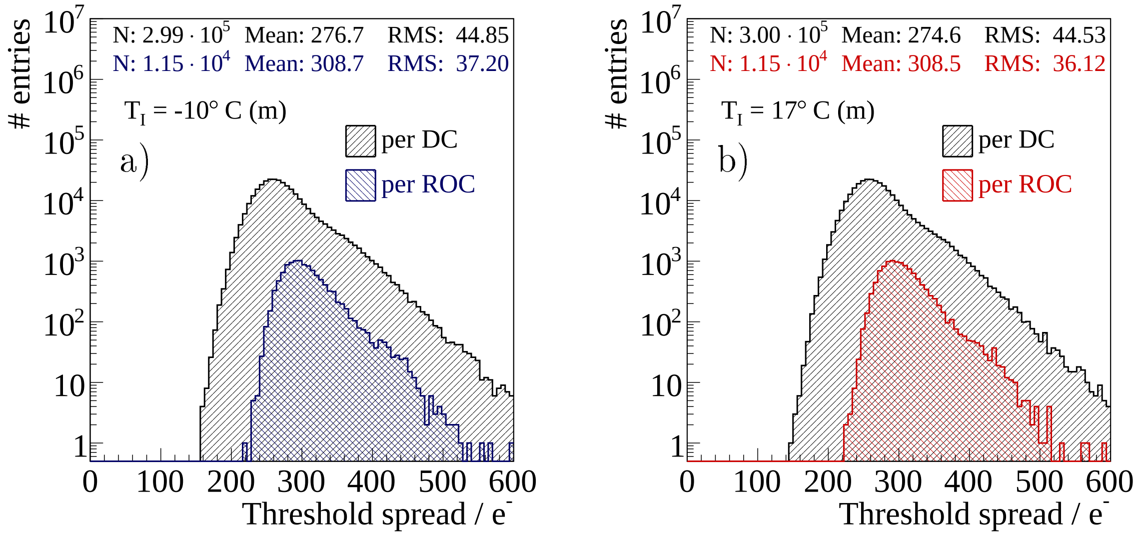

Figure 4.30: Comparison of the pixel threshold variation per ROC and per double-column (DC) before trimming, at a) −10 °C and b) 17 °C.

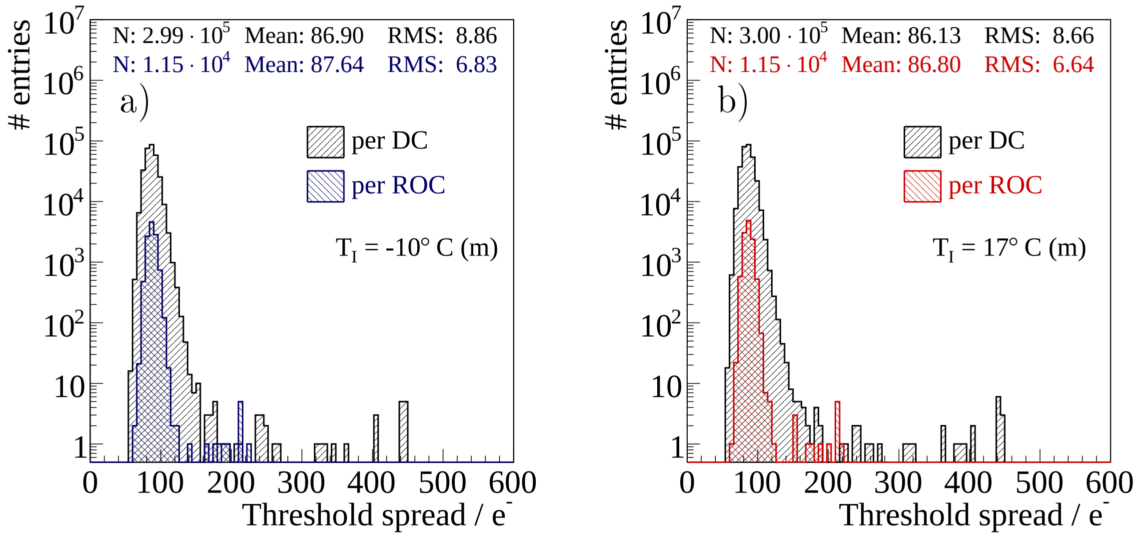

Figure 4.31: Comparison of the pixel threshold variation per ROC and per double-column (DC) after trimming, at a) −10 °C and b) 17 °C.

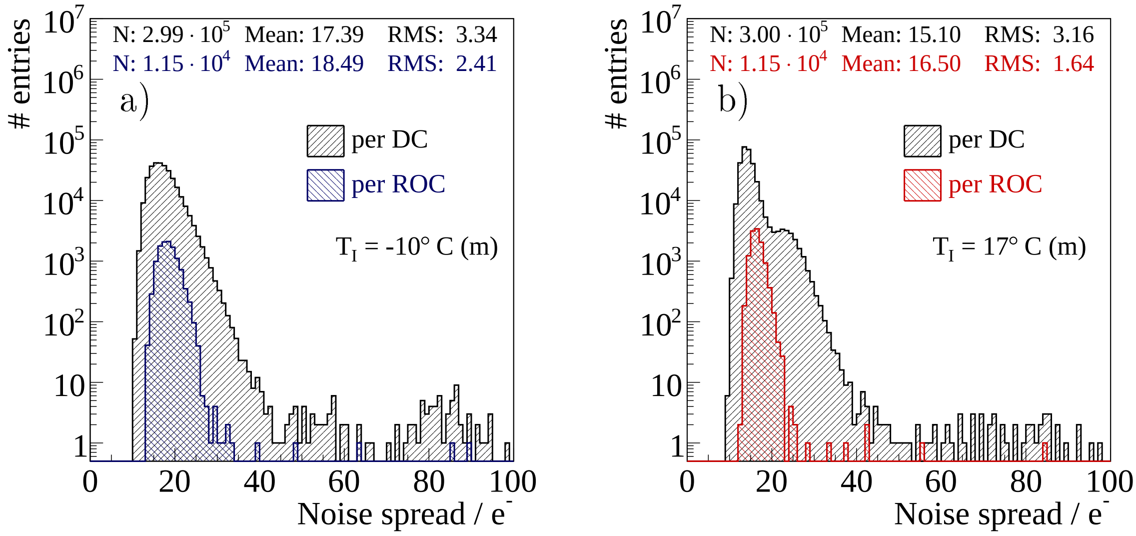

Figure 4.32: Comparison of the pixel noise spread per ROC and per double-column (DC), at a) $-10\,°\mathrm{C}$ and b) $17\,°\mathrm{C}$.

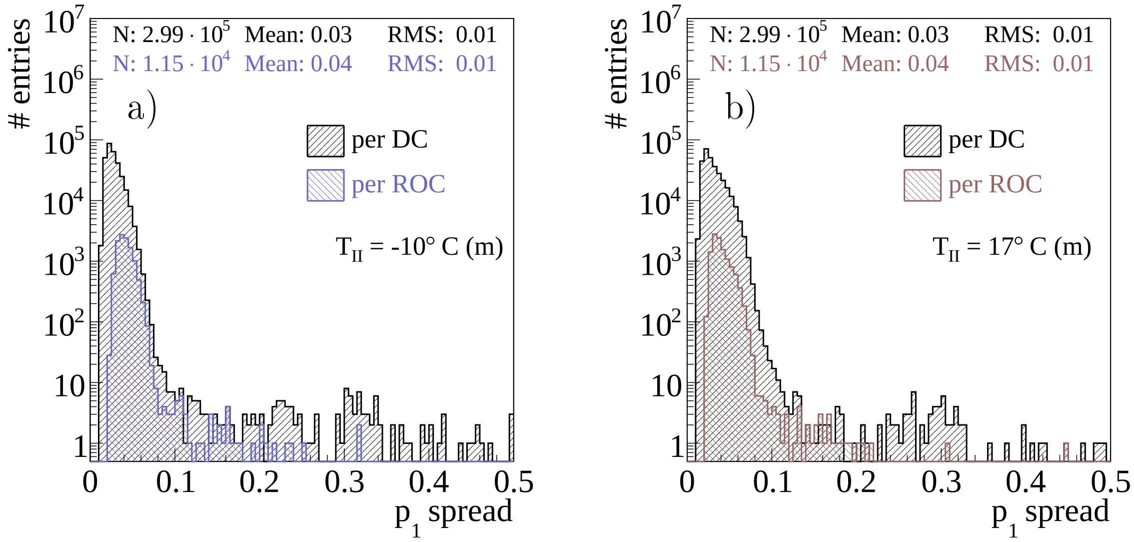

Figure 4.33: Comparison of the spread per ROC and per double-column (DC) of parameter p_1 of the hyperbolic tangent fit after the optimisation, at a) $-10\,°\mathrm{C}$ and b) $17\,°\mathrm{C}$.

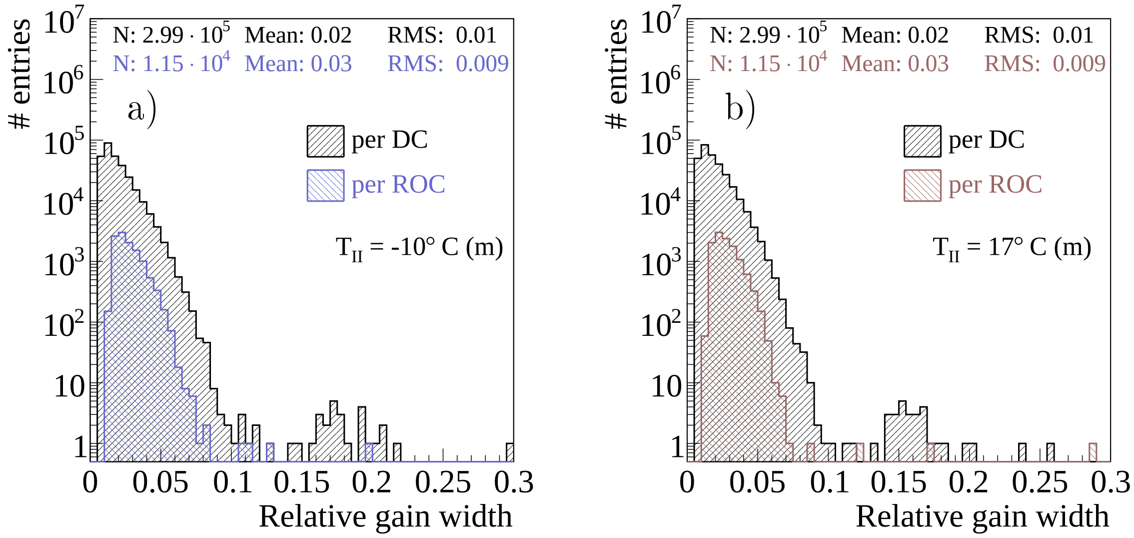

Figure 4.34: Comparison of the pixel relative gain width per ROC and per double-column (DC) after the optimisation, at a) −10 °C and b) 17 °C.

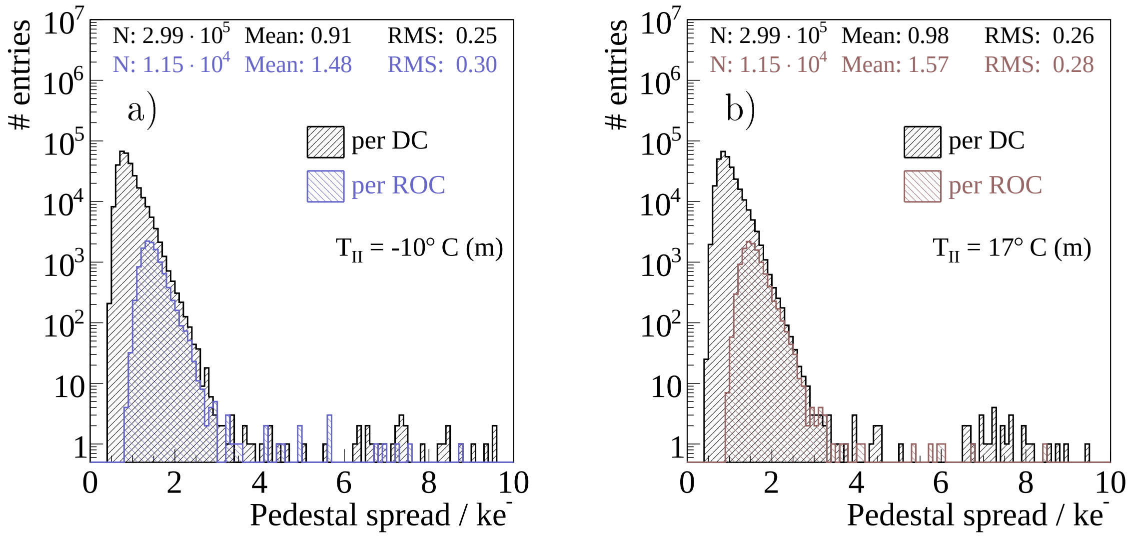

Figure 4.35: Comparison of the pedestal spread per ROC and per double-column (DC) after the optimisation, at a) −10 °C and b) 17 °C.

Table 4.4: ROC performance parameters of the 1.15×10^4 ROCs in the final detector system, established at $-10\,°C$ and $17\,°C$. The columns "RMS per ROC" and "RMS per DC" give the mean value of the performance parameter spreads per ROC and per DC respectively (see Figures 4.30—4.35).

		$-10\,°C$			$17\,°C$		
		Mean	RMS per ROC	RMS per DC	Mean	RMS per ROC	RMS per DC
Untrimmed Thr.	e^-	$4.96 \cdot 10^3$	309	277	$4.93 \cdot 10^3$	308	275
Trimmed Thr.	e^-	$3.91 \cdot 10^3$	87.6	86.9	$3.91 \cdot 10^3$	86.8	86.1
Noise	e^-	155	18.5	17.4	159	16.5	15.1
		Before optimisation of the PH calibration					
Gain	$\frac{ADC}{DAC}$	2.9	$9.7 \cdot 10^{-2}$	$6.1 \cdot 10^{-2}$	2.5	$6.6 \cdot 10^{-2}$	$3.8 \cdot 10^{-2}$
- relative spread	%	-	$3.5 \cdot 10^{-2}$	$2.2 \cdot 10^{-2}$	-	$2.7 \cdot 10^{-2}$	$1.5 \cdot 10^{-2}$
Pedestal	e^-	11	1.7	0.95	12	1.6	0.93
Parameter p_1		1.3	$4.8 \cdot 10^{-2}$	$3.7 \cdot 10^{-2}$	0.94	$3.8 \cdot 10^{-2}$	$2.9 \cdot 10^{-2}$
		After optimisation of the PH calibration					
Gain	$\frac{ADC}{DAC}$	2.7	$7.3 \cdot 10^{-2}$	$5.0 \cdot 10^{-2}$	2.5	$6.9 \cdot 10^{-2}$	$4.7 \cdot 10^{-2}$
- relative spread	%	-	$2.7 \cdot 10^{-2}$	$1.9 \cdot 10^{-2}$	-	$2.8 \cdot 10^{-2}$	$1.9 \cdot 10^{-2}$
Pedestal	e^-	23	1.5	0.91	22	1.6	0.98
Parameter p_1		1.2	$4.2 \cdot 10^{-2}$	$2.8 \cdot 10^{-2}$	1.1	$4.1 \cdot 10^{-2}$	$3.1 \cdot 10^{-2}$

Sensor leakage current: The sensor leakage current at the initial operation voltage ($V_{OP} = 150$ V) was measured several times during a test suite: either as a single point measurement at $V = 150$ V or as part of the IV-curve measurement (see section 4.5.2). The results of the latter are shown in Figure 4.36a) for the modules in the final system, giving a mean sensor leakage current of 0.729 μA at 17 °C and 0.118 μA at −10 °C. In 98 % of the cases the leakage current measured at 17 °C is below 3 μA and 95 % of the modules in the final system have a leakage current below 2 μA at the operation voltage.

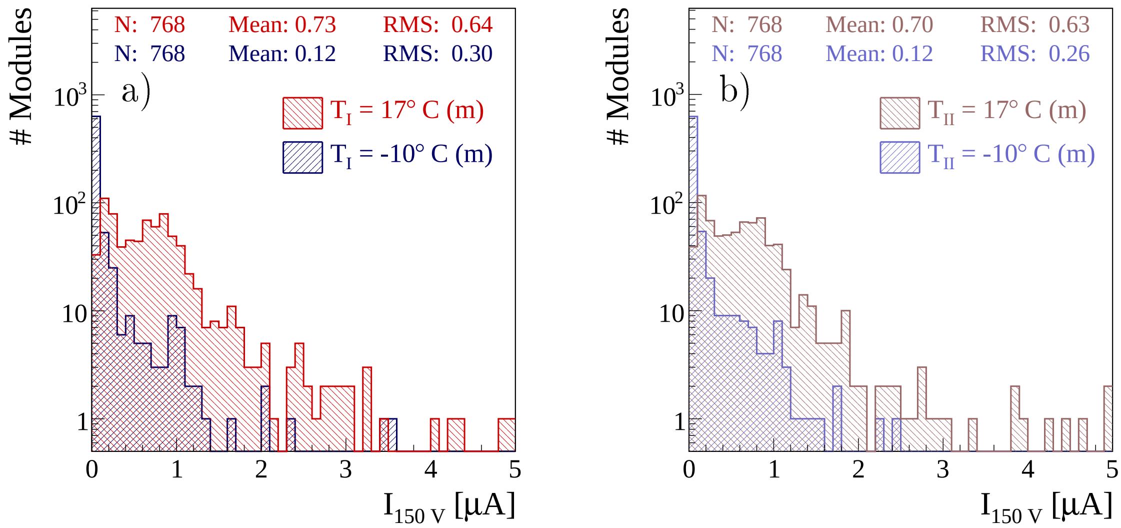

Figure 4.36: Sensor leakage current at the operation voltage $V_{OP} = 150$ V, measured at −10 °C and 17 °C in the 768 modules of the final system: a) the leakage current at V_{OP} extracted from the IV-curve measurements in *test suite I*; b) the leakage current determined in a single point measurement at V_{OP} in *test suite II*.

Chapter 5

Module Qualification

The quality of each module was assessed in a detailed procedure and included functionality, calibration and performance tests, as discussed in the previous chapter. Section 5.1 provides an overview of the general test setup and describes the two main test suites of the qualification procedure. The grading system and qualification criteria are devised in section 5.2. The results of the module qualification are presented in section 5.3.

5.1 Qualification Procedure

5.1.1 Module Assembly

A module consists of a sensor, 16 readout chips, a high density interconnect (HDI), a token bit manager (TBM), two basestrips, a signal and a power cable (see section 3). This section briefly describes the assembly process of a module [43] and the preselection that is applied at the different assembly steps. The sensor and ROC wafers are tested and pre-processed in several steps, before they are diced, picked and tested once more. Only devices with less than 1 % *noisy* or *dead pixel* and without *mask defects* are allowed at the next stage, where the 16 ROCs are connected to a sensor wafer by a dedicated bump bonding technique [36]. The emerging 'raw-module' is only further processed if each ROC passes the functionality tests and if the current-voltage characteristics of the sensor are acceptable. The pre-assembly of the HDI comprises of: gluing and wire bonding of the TBM to the HDI; soldering of the power cable to the HDI; gluing and wire bonding of the signal cable to the HDI. After verifying the functionality of the pre-assembled HDI, the base strips and the HDI are glued onto the raw-module and the module is completed by forming the electrical connection between ROCs and HDI with wire bonds. Only modules that reach the final assembly stage then enter an extensive qualification procedure.

The following sections will only focus on the modules that reached the final assembly stage. Modules that were rejected at an earlier stage, will not be considered. The implementation of additional tests at the module assembly stage (e.g. bump bonding test at the bare module stage) has led to a decrease of certain types of defects detected in the qualification of the fully assembled modules.

5.1.2 The Test Set-up

The fully assembled modules, which have successfully passed all assembly stages, enter an extensive test procedure in the module test station at the PSI laboratory. Up to four modules can be tested simultaneously at the test station. It is composed of the following elements:

- a cooling box to regulate the ambient temperature and humidity during module testing and thermal cycling
- four electronic testboards with a field-programmable gate array (FPGA)
- four module adapter boards to connect the modules to the test-board
- one Keithley high-voltage supply

The module testboards, the high-voltage power supply and the cooling box are connected to a desktop PC (Scientific Linux 4) from which they can be controlled remotely. The two different test suites of the module qualification procedure (see section 5.1.3) are executed by a supervisor script running on the PC. At the end of the test suite, the script initiates the automatic processing of the test results. Based on performance and functionality, a grade is assigned to each module (see section 5.2). Finally, a test summary comprising essential numbers and figures is uploaded to a web interface. Figure 5.1 depicts the general workflow of a module test suite.

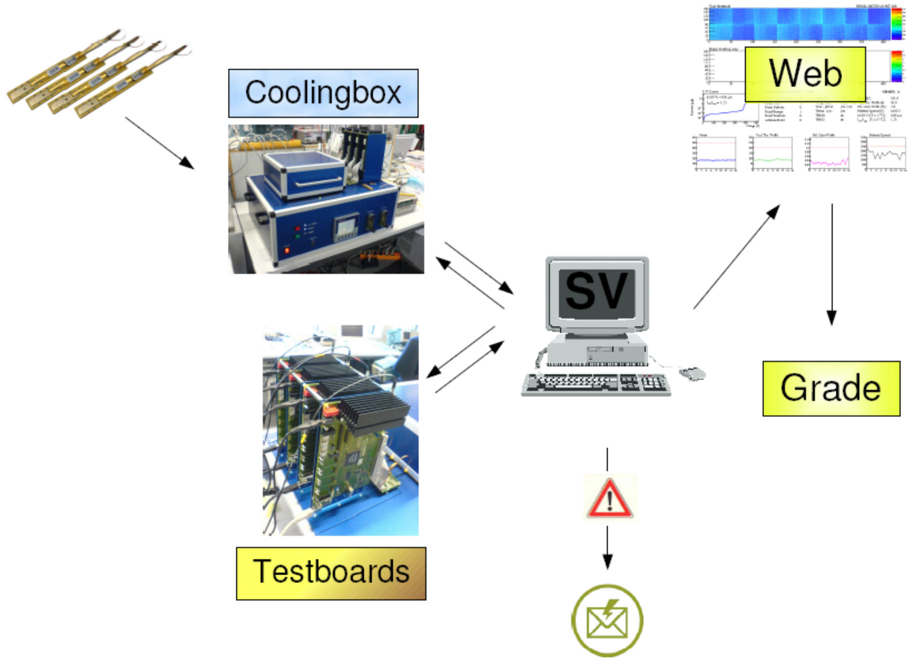

Figure 5.1: Test set-up for module qualification: Four modules can be tested simultaneously in the cooling box. Each module is connected to a testboard through a module adapter. The testboards are connected to the high-voltage supply (not shown) and to the PC (through a USB connection). The entire test suite is controlled by a supervisor script. The script initiates the automatic processing of the test results. A grade is assigned to each module and test summaries are uploaded to the web. After the test suite has finished or if a problem occurs, an email notification is sent to the tester.

5.1.2.1 The Cooling Box

The cooling box offers space for four modules. The temperature within the volume that contains the modules is adjusted by use of four water cooled, high-performance Peltier elements. To lower the humidity, Nitrogen is provided to the cooling box through two flow regulators - one with a high flow rate and one with a low flow rate. The first is only used at the beginning of the test while the second is constantly open during the test to maintain a low humidity in the cooling box. The Peltier elements and the N_2 flow regulators are connected to controller channels of a “JUMO Imago 500” process and program controller. The temperature is measured with a Platinum resistance thermometer (Pt-100) connected to the controller. Two program channels are allocated to regulate either heating or cooling. The communication between the controller and the PC is established with an RS422/485 serial interface (using Modbus protocol).

5.1.2.2 The Testboard

Control and readout signals are transmitted between module and PC by a testboard that was developed by engineers of ETH. It provides the digital and analogue voltages to the ROCs and the reverse bias voltage to the sensor. The core of the testboard is a Field-Programmable Gate Array (FPGA) with an implemented processor. The FPGA generates the electric signals such as calibrate, clock, trigger etc. The integrated processor allows test algorithms to be run directly on the testboard, thus reducing the data transfer between the PC and the testboard. Interactive test algorithms can therefore be performed much faster, e.g. the trim algorithm can be speeded up by a factor of three and the pulse height calibration by a factor of twelve by running parts of the algorithm directly on the testboard [37]. On the testboard, the two analogue output signals from the module are sampled with two 12-bit Analogue to Digital Converters (ADC) in the interval $[-2048, +2047]$, One ADC unit corresponds to 0.1275 mV. The data transfer between testboard and the PC takes place through a USB connection.

Data Trigger Level Scan: As explained in section 3.4, a series of ultra black levels—three in the TBM header and two in the TBM trailer—mark the beginning and the ending of the analogue readout of a module. For the testboard to detect the UBLs correctly, the data trigger level (DTL) has to be adapted to each module. Starting from zero, the data trigger level is decreased until a valid analogue empty readout is measured (i.e. a readout consisting of 64 clock-cycles) and the UBL can be determined. The data trigger level is then set to a value that lies 100 ADC above the UBL. Figure 5.2 shows the number of valid empty analogue readouts per module.

Sampling point adjustment: At the beginning of each test the pulse height of a random pixel is measured as a function of a delay that can be added to the module clock with respect to the ADC clock. The sampling point of the analogue signal is optimised by using a delay that corresponds to the centre of the range, in which the pulse height is less than 20 ADC units below the maximum value.

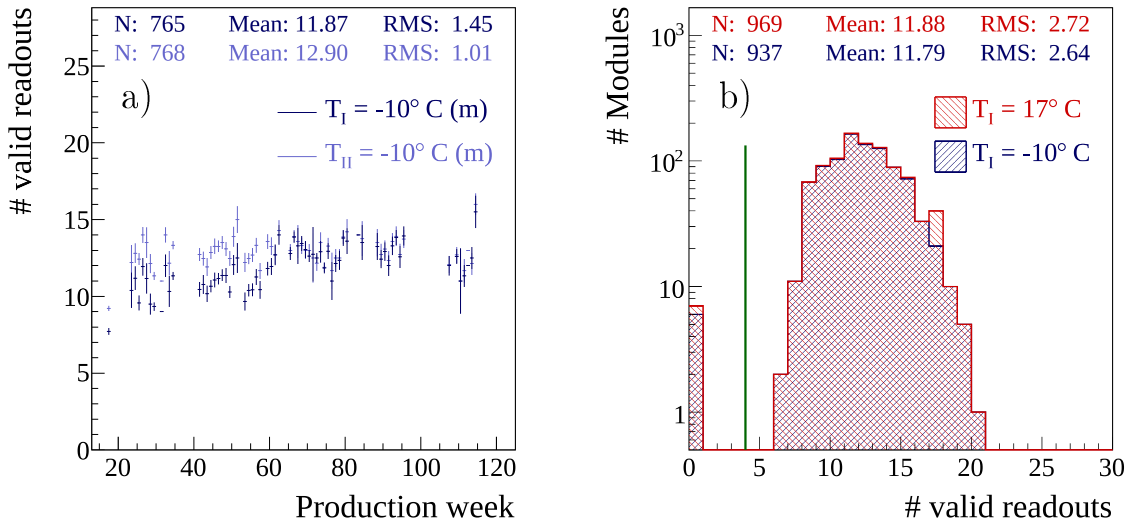

Figure 5.2: Number of valid empty analogue readouts (64 clock cycles) per module: a) The average number of valid readouts for modules produced in the same week as a function of time, determined at $-10\,°\mathrm{C}$ at the beginning of the first test suite (T_I) and second test suite (T_{II}). Only the modules in the final system are shown. b) Distribution of the number of valid readouts per module for all tested modules, determined at the beginning of the first test suite (first at $17\,°\mathrm{C}$ and then at $-10\,°\mathrm{C}$). Modules with less than four valid readouts were disabled.

5.1.2.3 Default Settings

Unless explicitly mentioned, the default settings listed below apply in both of the test suites described in section 5.1.3.

- A test sequence of a pixel always contains the following steps: enable the double-column, the calibration injection and the readout of the pixel, send a calibration signal and then a trigger to the module.
- Only the pixel being examined is enabled and thus the pixel address in the analogue readout does not need to be decoded (except in the address decoding test).
- The standard calibration signal is injected with a *Vcal* DAC of 200 in the low range.
- During testing a bias voltage of 150 V is applied to the sensor.
- The readout speed is 40 MHz.

5.1.3 The Test Suites

The module qualification was performed in two steps. A first, major series of tests—*test suite I*—comprising of all the functionality, calibration and performance tests described in section 4. Based on performance and functionality, a grade was assigned to each module (see section 5.2). The selection of modules that had qualified for the final system, underwent a second series of tests—*test suite II*—with the purpose of ensuring the functional integrity of the module before mounting it onto the detector half shells and also to optimise the ROC performance, as described in [38]. The second test suite also contained the *Vcal* calibration test of each module with two different X-ray sources, as described in section 4.6.

A colour coding scheme for the different test suites and different temperatures has already been introduced in section 4.1 and is illustrated in Figure 4.1.

5.1.3.1 Test Suite I (after assembly)

The first test suite consists of three test sequences, two IV-curve measurements and one thermal cycling part. At the beginning of a test suite, a data trigger level scan is performed (see section 5.1.2.2). Modules with less than four valid readouts are disabled. Each test sequence can be divided into three main parts: At the beginning of every test sequence all 26 DACs and the three registers are initialised to their default values to set the ROCs into the operating regime (see section 4.3). In the next step, the pixel readout circuit and electrical connections to sensor pixels of each pixel are evaluated, comprising the test algorithms for the pixel response, bump bonding quality, trim bit test and address decoding (see section 4.4). The third part includes the performance and calibration tests, such as the noise measurement and the pulse height calibration (without optimisation) for each pixel as well as the threshold unification with the trim algorithm (see sections 4.5—4.7). During the thermal cycling process, the modules are cycled ten times from $-10\,°\mathrm{C}$ to $17\,°\mathrm{C}$. The order of the different testing steps is the following:

- Test sequence at $-10\,°\mathrm{C}$
- Thermal cycling between $-10\,°\mathrm{C}$ and $+17\,°\mathrm{C}$
- Test sequence at $-10\,°\mathrm{C}$, IV-curve measurement
- Test sequence at $+17\,°\mathrm{C}$, IV-curve measurement

To avoid running into the compliance of the power supply, the leakage current of every module is checked each time after the temperature has been adjusted to a new value. Modules with a leakage current above 25 μA at the operating voltage of 150 V are disabled. The test sequences are performed simultaneously in all four modules and generally last about one and a half hours. The IV-curve has to be measured consecutively for each module and takes about ten minutes per module. The thermal cycling process lasts about one hour. Figure 5.3 shows the temperature profile and the different steps of *test suite I*. As shown in Figure 5.4a), the initial test duration of about ten hours was reduced to about six hours after optimising the thermal cycling and the time consuming IV-curve measurements. The results of *test suite I* are denoted with the subscript I (e.g. $T_I = 17\,°C$) and are displayed in bright colours: red represents the results of measurements at 17 °C and blue the results of measurements at −10 °C *after* thermal cycling (see Figure 4.1).

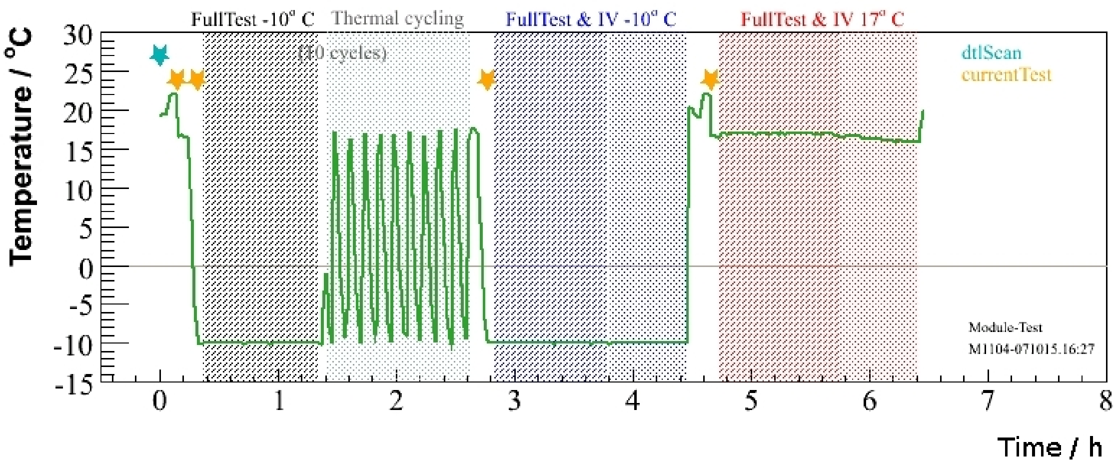

Figure 5.3: The temperature profile shows the first test sequence (without IV-curve measurement) at −10 °C, followed by ten cycles, in which the temperature is continuously alternated between −10 °C and 17 °C. After the thermal cycling, the test sequence is repeated once more at −10 °C followed by a separate IV-curve measurement in each module. The same two steps, a test sequence and an IV-curve measurement, are then repeated once more at 17 °C. At the beginning of the test suite, before cooling down to −10 °C, the leakage current is checked and a DTL scan is performed, checking for at least four valid empty analogue readouts. In addition, the current is checked every time before starting a new test sequence.

5.1.3.2 Test Suite II (before mounting)

Before mounting a module that has successfully passed the first test stage onto the detector half shells, the module is examined once more in a series of two reduced test sequences, featuring some basic functionality tests as well as the linearity optimisation of the pulse height calibration curve. A data trigger level scan is again performed at the beginning of the test suite and modules with less than four valid readouts are disabled. At the beginning of each test sequence, the DACs of each ROC are either dynamically optimised or initialised to the updated set of nonrecurringly optimised DAC parameters (see Table B.1). In addition, the dynamic DAC optimisation in *test suite II* also includes the algorithm to maximise the linear range of the pulse height calibration curve (see section 4.7). The pixel functionality tests of the reduced test sequence contains only the pixel readout test, that is repeated in each pixel. Again, the leakage current of every module is checked each time after the temperature has been adjusted to a new value and modules with a leakage current above 25 μA at a voltage of 150 V are disabled. Furthermore, the correlation of the injected calibration signal at a given $Vcal$ and the ionisation charge is established in the X-ray test(see section 4.6). Each module was irradiated separately at a test station outside the cooling box, using a Molybdenum and a Silver X-ray source. The order of the different testing steps is the following:

- Test sequence at −10 °C
- Test sequence at 17 °C
- X-ray test

Figure 5.4b) shows the duration of *test suite II* not including the X-ray test. The latter takes about 20 minutes per module. Altogether the second test suite lasted about three hours. The results of *test suite II* are denoted with the subscript *II* (e.g. $T_{II} = 17\,°\mathrm{C}$) and are displayed in muted colours: red represents the results of measurements at 17 °C and blue the results of measurements at −10 °C (see Figure 4.1).

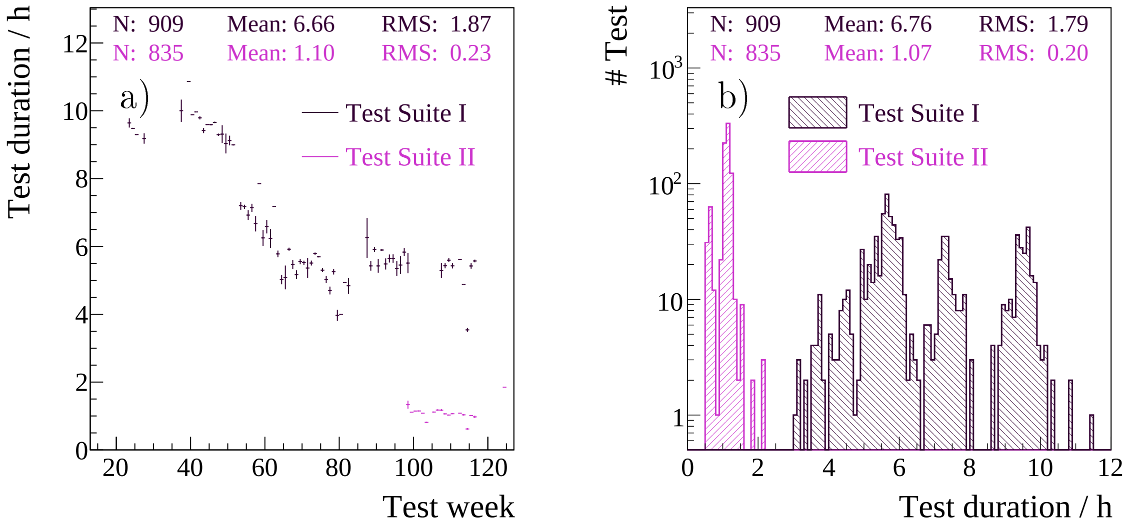

Figure 5.4: Test durations of *test suites I* and *II*: a) The average duration of a test suite as a function of time. b) Distribution of the test duration for *test suite I* (T_I) and *test suite II*

5.2 Qualification Criteria

A grading system has been established and consists of three categories: A, B and C. Modules with grade A have no or only minor defects and qualify for use anywhere in the final detector. Modules with grade B are of lesser quality than modules with grade A, but are still working acceptably well to be used in the experiment. The type of defect(s) of a grade B module should be considered before assigning it to a specific detector layer. Modules with grade C are seriously flawed or not working at all. If all attempts to recuperate such a module failed, the final grade was left at C and the module was considered to be waste.

The grades are assigned according to qualification criteria derived from performance and lifetime requirements of the pixel detector in the experimental conditions of CMS. These criteria can be divided into three categories—*sensor performance, chip performance and pixel defects*—and will be explained in detail in the following sections. The grading criteria are summarised in Table 5.3 at the end of section 5.2.

5.2.1 Module sensor quality

Being the innermost measurement device of CMS, the pixel detector has to sustain the harsh radiation environment close to the interaction region—with charged particle fluxes up to 10^8 $\text{cm}^{-2}\,\text{s}^{-1}$ in the first layer. The consequences of radiation-induced defects in the sensor are charge trapping, rising leakage current and—subsequent to a space charge sign inversion—an increasing full depletion voltage. To ensure reliable operation at nominal luminosity throughout the expected lifetimes (2 years for the innermost and more than 10 years for the outermost layer), a stable current voltage characteristic of unirradiated sensors up to high voltages is imperative. Problems in the sensor production process or damage inflicted upon the sensor during module assembly alter the expected current-voltage dependence explained in section 4.5 and can cause high leakage currents.

To identify damaged or malfunctioning sensors, the IV-characteristics of each module were recorded at room temperature (17 °C) and at −10 °C. At the initial operational voltage of $V_{OP} = 150$ V, the leakage current measured at room temperature should be below 2 μA for a module to be of grade A, and to be of grade B the leakage current should not exceed 10 μA. Since these reference values are only valid at room temperature, the measurements performed at −10 °C were recalculated, using the following correlation between the sensor leakage current and the ambient temperature

$$I \propto T\ \mathrm{e}^{-E_g/2kT}, \tag{5.1}$$

where k is the Boltzmann constant and E_g is the energy gap, defined as the difference between the lower edge of the conduction band and the upper edge of the valence band ($E_g = 1.12\,\text{eV}$ in Silicium). Figure 5.5 shows that, the average ratio of the value recalculated from the current measured at −10 °C and the actual current measurement at room temperature, tends to be around 1.5 rather than one. Therefore, the reference values for grading based on measurements at −10 °C were multiplied by a factor of 1.5. The adapted limits for the values recalculated for room temperature are 3 μA and 15 μA for a module to qualify for grade A or grade B respectively. The distributions of the leakage currents, illustrated in Figure 5.6, have been extracted from the IV-curves measured at 17 °C and −10 °C in *test suite I* and correspond to the sensor leakage currents at 150 V at 17 °C.

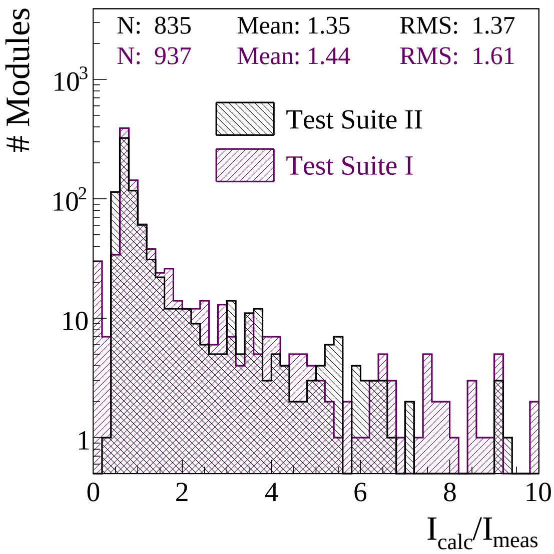

Figure 5.5: Ratio of the value recalculated from the current measured at −10 °C and the actual current measurement at room temperature in *test suite I* and *II*. (In *test suite I*, the leakage current at 150 V is extracted from the *IV*-curve measurement and in *test suite II*, it is determined in a single point measurement at 150 V.)

Accumulating radiation damage will increase the full depletion voltage and require higher operation voltages. Whenever possible the sensor should be operated within the plateau region and always below the breakdown voltage. Although the slope of the *IV*-curve, defined as the ratio between the two current measurements at 150 V and at 100 V, allows to detect early sensor breakdowns, it primarily serves as a measure of the current increase in the plateau region towards higher voltages. The *IV*-slope criterion in equation 5.2 therefore merely classifies the modules into category A and B.

$$I(V_{OP})/I(V_{OP} - 50\ \mathrm{V}) \leq 2. \tag{5.2}$$

Figure 5.7 shows the values of *IV*-slope extracted from the *IV*-curve, measured at 17 °C and −10 °C in *test suite I*. Modules with early sensor breakdowns and with *IV*-characteristics, that deviate significantly from the *IV*-curve in Figure 4.13 were rejected.

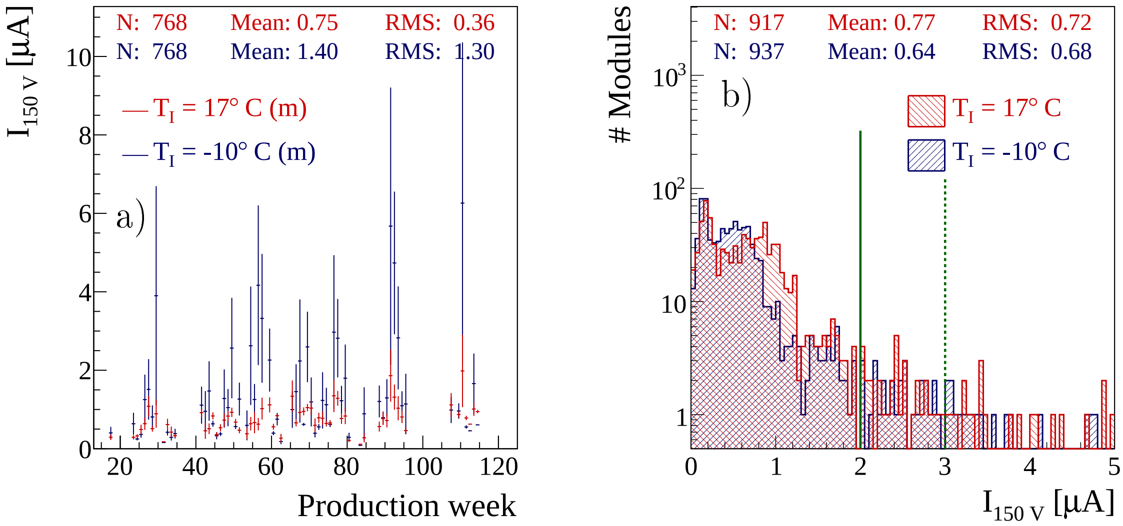

Figure 5.6: Distributions of the leakage current at 17 °C, determined at −10 °C and 17 °C during *test suite I* (current measurements at −10 °C were recalculated to the corresponding current at 17 °C, using formula 5.1): a) The average leakage current of modules produced in the same week as a function of time. b) The solid green line indicates the criteria for grade A for measurements at 17 °C and the dashed green line shows the adapted value for measurements at −10 °C.

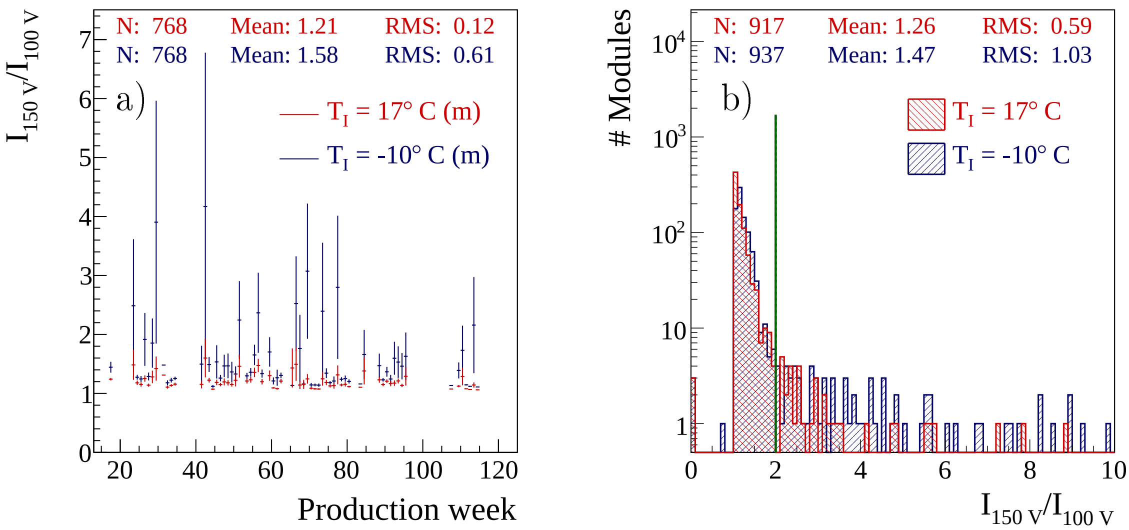

Figure 5.7: Distributions of the slope of the *IV*-curve, determined at −10 °C and 17 °C in *test suite I*: a) The average *IV*-slope of modules produced in the same week as a function of time. b) The green line separates modules into grade A and grade B .

5.2.2 Chip performance

Performance tests and performance optimisation of the ROCs are important elements in the module qualification since the chip performance affects the efficiency and precision of the hit reconstruction. Therefore, performance based grading criteria were introduced. They were derived from the study on position resolution and reconstruction efficiency in [44]. The angle of the trajectory and the Lorentz drift of charged particle in the magnetic field lead to charge sharing among pixels in the detector. Therefore, the hit reconstruction algorithm is based on clusters. A cluster is defined as a set of adjacent pixels above a certain threshold in units of noise and generally contains more than one pixel. The hit position is evaluated from the track angle and the charge distribution in the cluster pixels. The position resolution is affected by various factors, some of which can be controlled on the level of the PUC with the appropriate parameters. The *threshold*, for example, deteriorates the resolution as it increases, since the detector is operated in a zero suppressed mode and only pixels with a signal above a certain threshold are read out. The binary readout resolution is reached at a threshold of approximately $5500\,e^-$ in the z-direction and $9200\,e^-$ in the $r-\phi-$direction. Both of these values are far off the nominal threshold of $2500\,e^-$ for unirradiated modules. The position resolution is indirectly affected by the electronic *noise*, since the amount of noise in a pixel influences the threshold level. To be at least 5σ below the threshold requirement, the average noise should not exceed $400-500\,e^-$. Below $1000\,e^-$ the value of the noise itself only has little direct effect on the resolution. A more recent study [38] has shown, that the typical *threshold variations* of a few hundred electrons before threshold unification with the trim algorithm, also have no influence on the position resolution. However, if the same *pulse height calibration constants* are applied for all pixels on a ROC, the non-uniformity of the pixel response leads to a degradation of the position resolution. In this case a compromise has to be found to keep the number of calibration constants low while maintaining an acceptable position resolution. Applying an approximate pixel calibration is acceptable as long as the impact of the miscalibration on track and vertex reconstruction is inconsiderable compared to multiple scattering and misalignment effects. Relative *gain variations*, defined as the spread of the gain distribution divided by its mean, up to $20-40\,\%$ and *pedestal spreads* as large as $1000-2000\,e^-$ are tolerable according to [44].

Based on these consideration, the following grading scheme with respect to the chip performance was established. A module will be graded A (B), if

- The mean of the *noise* distribution is below 500 e^- (1000 e^-).
- The spread of the *threshold* after trimming does not exceed 200 e^- (400 e^-)-
- The relative *gain* width is less than 10 % (20 %).
- The *pedestal* spread lies below 2500 e^- (5000 e^-).

Figures 5.8—5.17 illustrate the ROC mean and the ROC spread distributions of the different performance parameters. The criteria for grade A (B) listed above are indicated in the corresponding histograms with solid (dashed) lines.

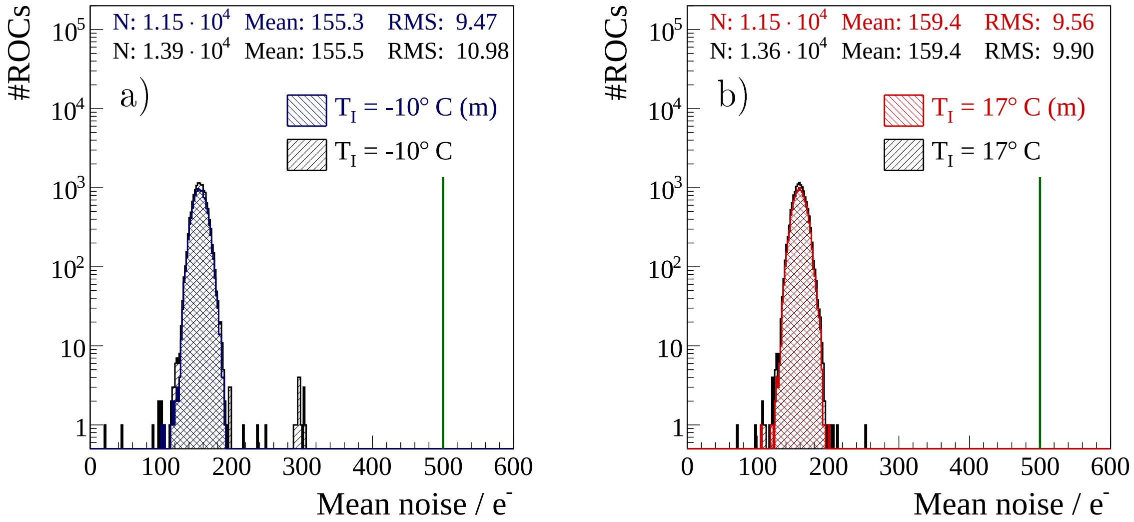

Figure 5.8: Mean noise per ROC: for the 1.15×10^4 ROCs of the modules mounted in the final system (m) and a) the $1.39 \cdot 10^4$ ROCs of all modules tested at $-10\,°\mathrm{C}$, b) the $1.36 \cdot 10^4$ ROCs of all modules tested at $17\,°\mathrm{C}$. The solid green line indicates the upper limit for the mean noise per ROC on a grade A module.

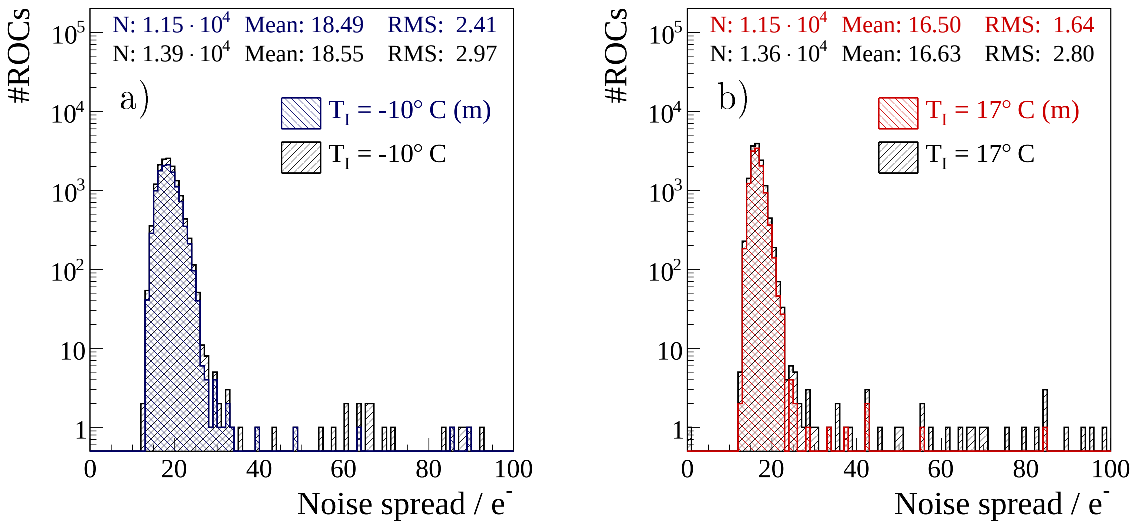

Figure 5.9: Noise spread per ROC: for the 1.15×10^4 ROCs of the modules mounted in the final system (m) and a) the $1.39 \cdot 10^4$ ROCs of all modules tested at $-10\,°\mathrm{C}$, b) the $1.36 \cdot 10^4$ ROCs of all modules tested at $17\,°\mathrm{C}$.

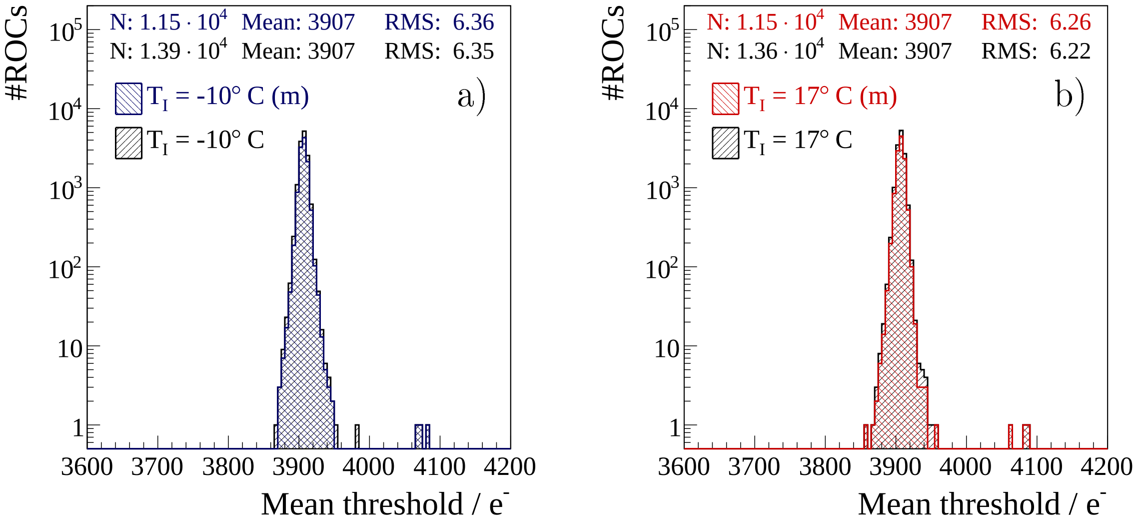

Figure 5.10: Mean threshold per ROC: for the 1.15×10^4 ROCs of the modules mounted in the final system (m) and a) the $1.39 \cdot 10^4$ ROCs of all modules tested at $-10\,°C$, b) the $1.36 \cdot 10^4$ ROCs of all modules tested at $17\,°C$.

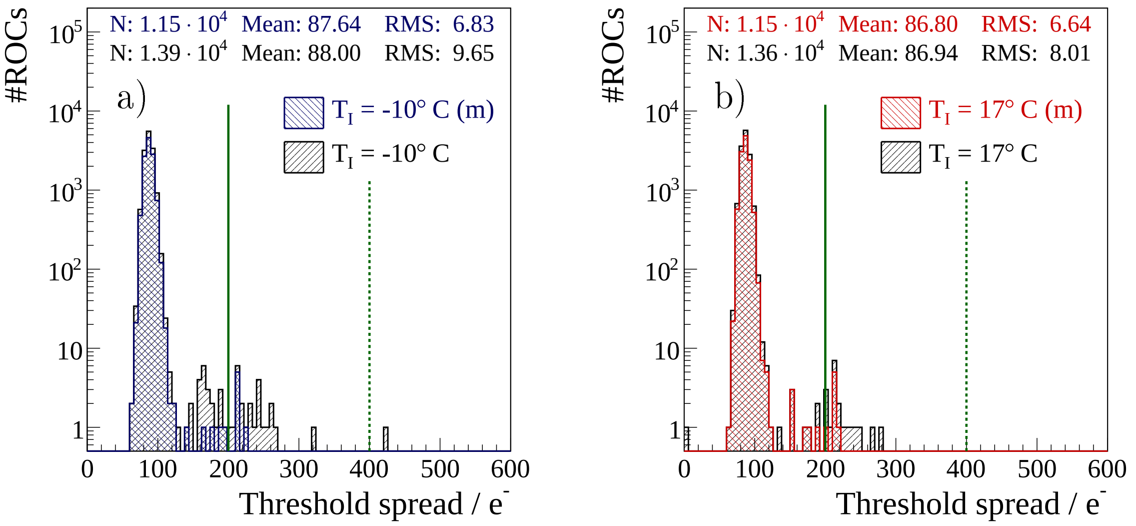

Figure 5.11: Threshold variation per ROC: for the 1.15×10^4 ROCs of the modules mounted in the final system (m) and a) the $1.39 \cdot 10^4$ ROCs of all modules tested at $-10\,°C$, b) the $1.36 \cdot 10^4$ ROCs of all modules tested at $17\,°C$. The solid (dashed) green line indicates the upper limit for the threshold variation per ROC on a grade A (B) module.

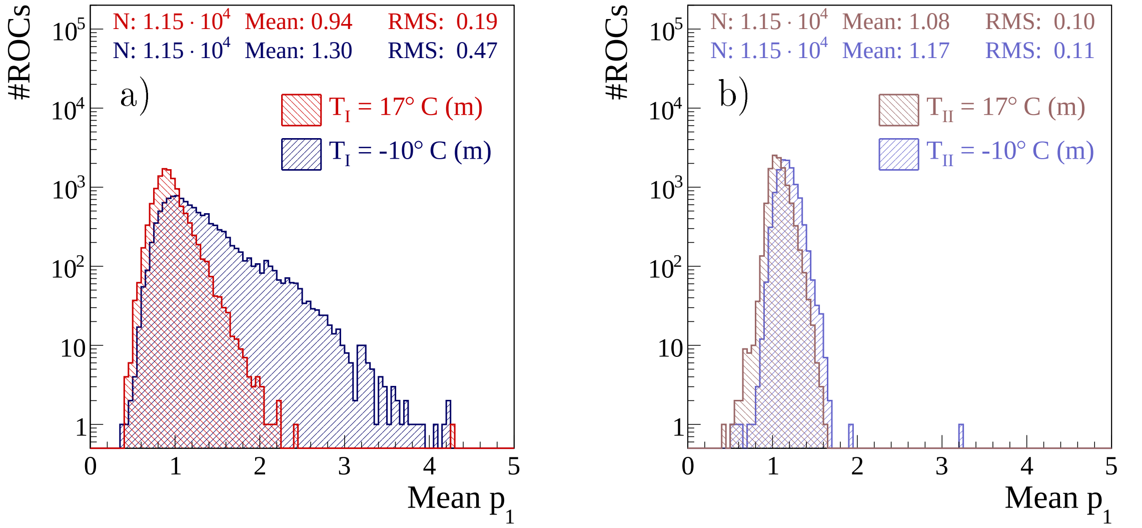

Figure 5.12: Mean parameter p_1 per ROC at $-10\,°\mathrm{C}$ and $17\,°\mathrm{C}$ for the 1.15×10^4 ROCs of the modules mounted in the final system: a) before the optimisation, b) after the optimisation of the pulse height calibration curve.

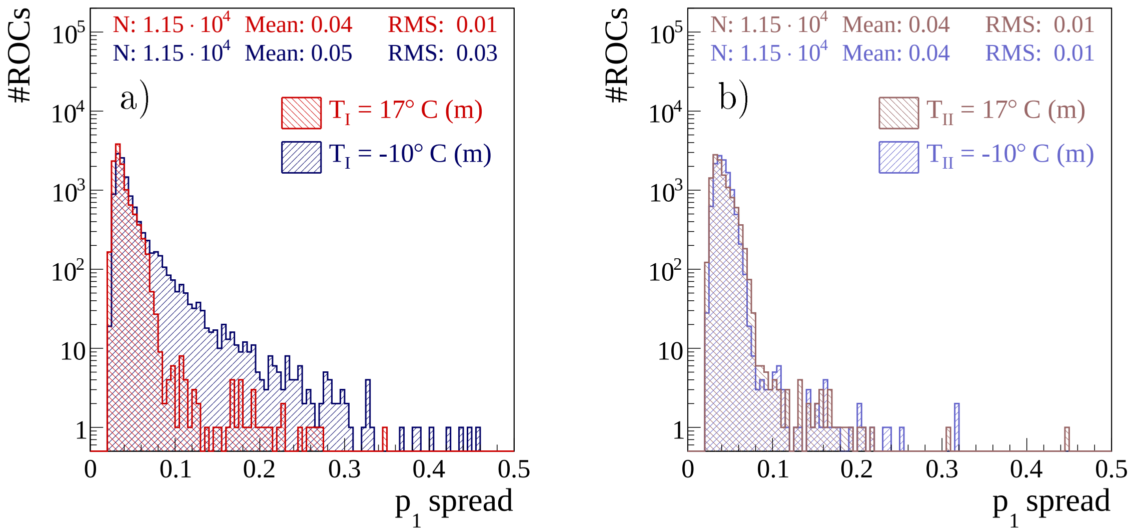

Figure 5.13: Parameter p_1 spread per ROC at $-10\,°\mathrm{C}$ and $17\,°\mathrm{C}$ for the 1.15×10^4 ROCs of the modules mounted in the final system: a) before the optimisation, b) after the optimisation of the pulse height calibration curve.

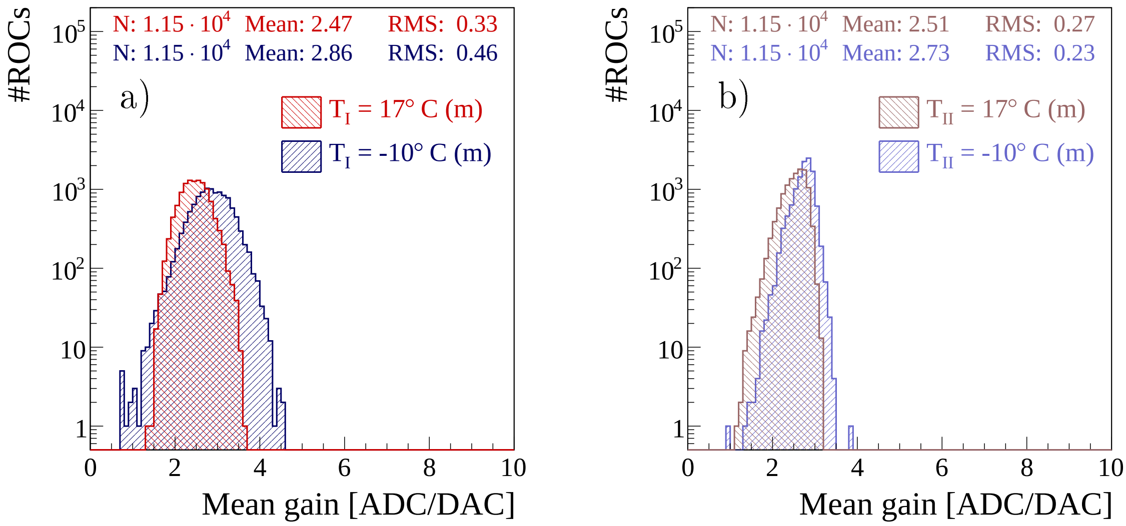

Figure 5.14: Mean gain per ROC at $-10\,°\mathrm{C}$ and $17\,°\mathrm{C}$ for the 1.15×10^4 ROCs of the modules mounted in the final system: a) before the optimisation, b) after the optimisation of the pulse height calibration curve.

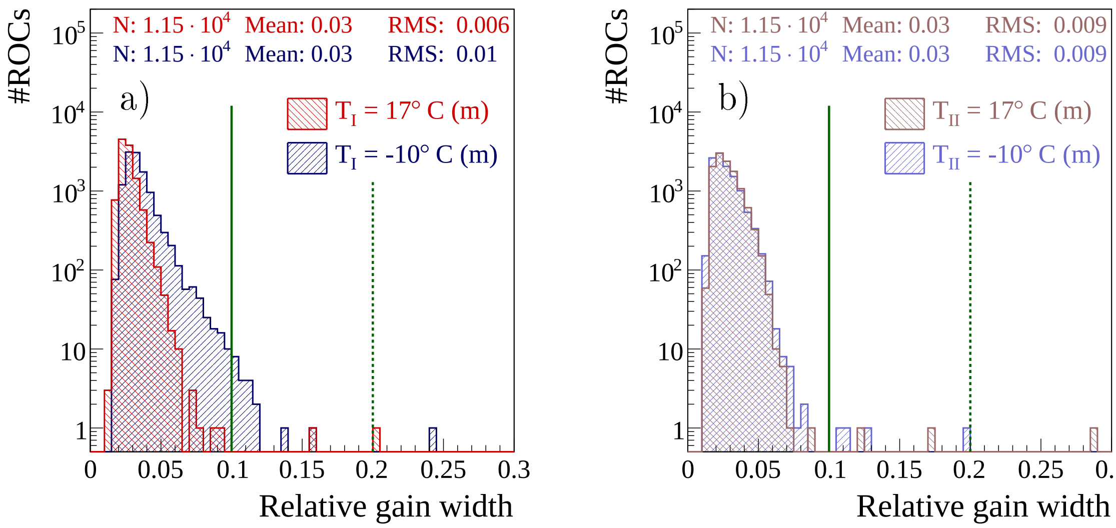

Figure 5.15: Relative gain width per ROC at $-10\,°\mathrm{C}$ and $17\,°\mathrm{C}$ for the 1.15×10^4 ROCs of the modules mounted in the final system: a) before the optimisation, b) after the optimisation of the pulse height calibration curve. The solid (dashed) green line indicates the upper limit for the relative gain width per ROC on a grade A (B) module.

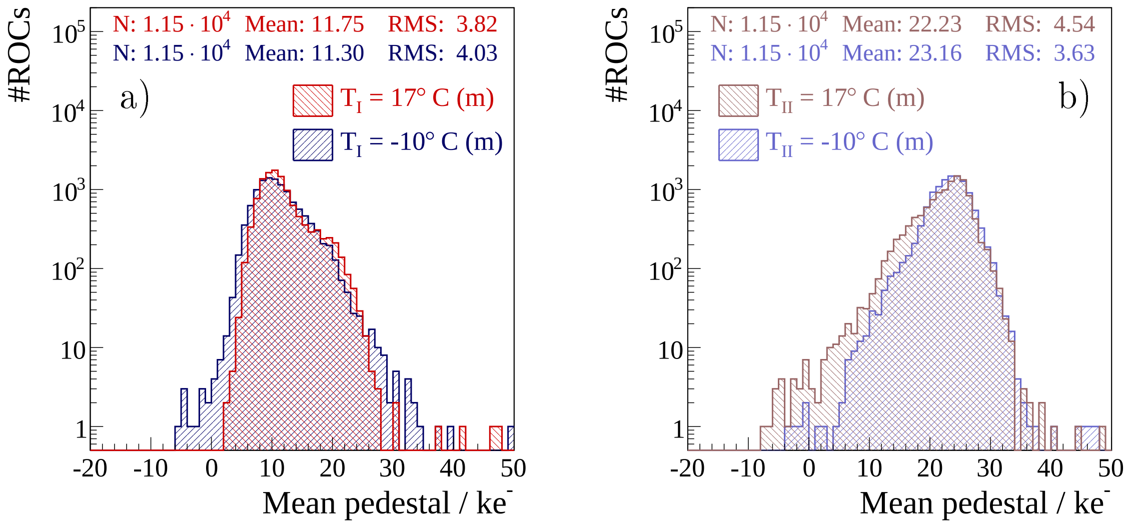

Figure 5.16: Mean pedestal per ROC at $-10\,°C$ and $17\,°C$ for the 1.15×10^4 ROCs of the modules mounted in the final system: a) before the optimisation, b) after the optimisation of the pulse height calibration curve.

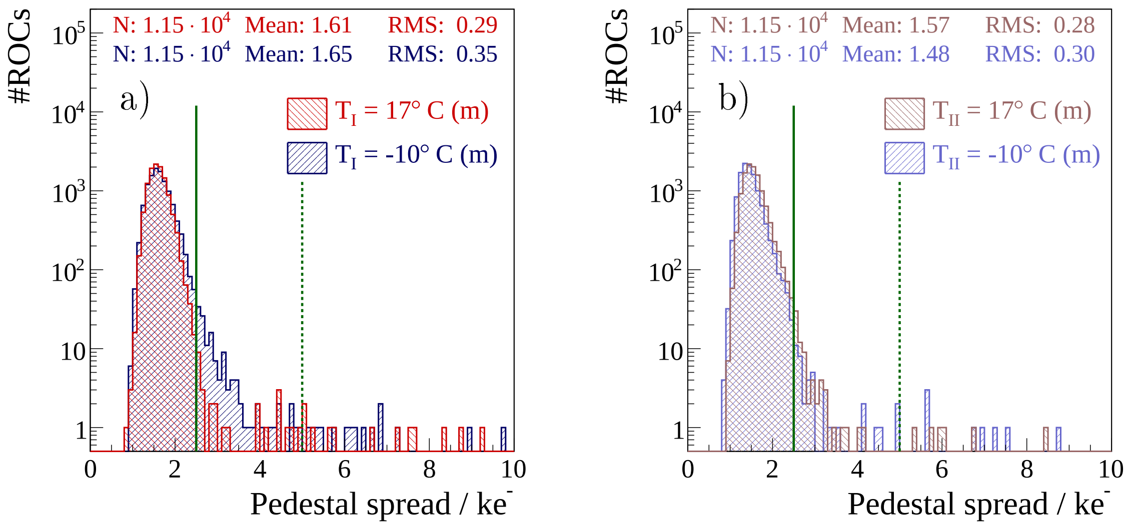

Figure 5.17: Pedestal spread per ROC at $-10\,°C$ and $17\,°C$ for the 1.15×10^4 ROCs of the modules mounted in the final system: a) before the optimisation, b) after the optimisation of the pulse height calibration curve. The solid (dashed) green line indicates the upper limit of the relative gain width per ROC on a grade A (B) module.

5.2.3 Pixel defects

5.2.3.1 Functionality defects

In a similar manner to chip performance, inefficient or broken pixels will degrade the reconstruction efficiencies. Missing charge will lead to inaccurate hit position due to incorrect charge interpolation. In particular hits at high rapidities are affected, where a long pattern can be misidentified as two separate clusters due to a lost pixel. The readout circuit and the electrical connection between the PUC and the sensor pixel are tested for each pixel as part of the qualification procedure. The pixel functionality is evaluated based on the test algorithms explained in section 4.4. A pixel is counted as defective, if one or several of the following tests failed:

- Pixel readout test
- Bump bonding test
- Trim bits test
- Address decoding test

Each module is qualified based on the amount of pixel defects per ROC: Less than 1 % of defective pixel are allowed on the ROCs of a module of grade A and at most 4 % on the ROCs of a modules of grade B.
Being able to mask a noisy pixel is crucial, since such a pixel may jam the buffers of the readout system. Therefore, a module with as much as one mask bit defect was graded as C. Table 5.1 summarises the number and fraction of the different pixel functionality defects.

5.2.3.2 Performance deficiencies

In addition to pixel defects, performance deficiencies were defined based on the results of the performance and calibration tests. The four performance parameters to consider are the *noise*, the *trimmed threshold*, the *gain* and *parameter* p_1 of the hyperbolic tangent fit. A performance deficiency does not imply that the pixel per se is defective. Many pixels with a performance deficiency may still function acceptably. For example, a pulse height calibration curve that is non-linear in the low range, does not render a pixel completely inoperative, nor does a pixel, that could not be trimmed at the specific *Vcal* value of 60, mean that it cannot be trimmed at all. In some cases the noise

Table 5.1: Number and fraction of pixel with functionality defects. The table shows the number of defects measured at $-10\,°C$ and $17\,°C$, on the left for the entire production and on the right for the modules in the final system. The 'ROCs column' shows the number of ROCs containing one or more pixel with functionality defects. A dead pixel is not attributed any other defects, since the pixel being 'alive' is the premise for all other pixel functionality tests.

		Entire Production			Modules in final system		
	Temp.	#ROCs	#Pixels	Fraction	#ROCs	#Pixels	Fraction
Tested	-10	$1.39 \cdot 10^4$	$5.76 \cdot 10^7$	-	$1.15 \cdot 10^4$	$4.79 \cdot 10^7$	-
	17	$1.36 \cdot 10^4$	$5.64 \cdot 10^7$	-	$1.15 \cdot 10^4$	$4.79 \cdot 10^7$	-
Readout (dead)	-10	744	1188	$2.1 \cdot 10^{-5}$	732	1086	$2.3 \cdot 10^{-5}$
	17	912	1480	$2.6 \cdot 10^{-5}$	760	1141	$2.4 \cdot 10^{-5}$
Noisy Readout	-10	2	2	$3.5 \cdot 10^{-8}$	0	0	0
	17	2	2	$3.6 \cdot 10^{-8}$	0	0	0
Mask Bit	-10	4	< 23	$< 4.0 \cdot 10^{-7}$	0	0	0
	17	5	< 24	$< 4.3 \cdot 10^{-7}$	0	0	0
Bump Bond	-10	2149	$4.04 \cdot 10^4$	$7.0 \cdot 10^{-4}$	1668	6289	$1.3 \cdot 10^{-4}$
	17	2115	$4.08 \cdot 10^4$	$7.3 \cdot 10^{-4}$	1671	6289	$1.3 \cdot 10^{-4}$
Trim Bit	-10	33	35	$6.1 \cdot 10^{-7}$	28	30	$6.3 \cdot 10^{-7}$
	17	30	30	$5.3 \cdot 10^{-7}$	26	26	$5.4 \cdot 10^{-7}$
Address Decoding	-10	6	85	$1.5 \cdot 10^{-6}$	5	5	$1.0 \cdot 10^{-7}$
	17	10	267	$4.7 \cdot 10^{-6}$	7	185	$3.9 \cdot 10^{-6}$

of a pixel could not be determined due to a non-converging fit in the test algorithm. Therefore, contrary to the pixel defects listed in section 5.2.3.1, performance deficiencies will not be included in the total number of pixel defects on a ROC, that is used to grade the module, but merely serve as a figure of merit to assess the individual pixel performance. The number and fraction of the different pixel performance deficiencies are summarised in Table 5.2.

Noise deficiencies: The distributions of the pixel noise measured at $-10\,°C$ and $17\,°C$ are illustrated in Figures 5.18 and 5.19. If the noise of a pixel exceeds $400\,e^-$, the pixel is considered to be noisy. A noise below $50\,e^-$ indicates abnormal behaviour or, for entries at -1, a failed S-curve fit. A pixel with a noise of less than $50\,e^-$ or above $400\,e^-$ is considered to be deficient. Figure 5.20 shows the location of the 5811 pixels with noise deficiencies at $-10\,°C$ on the modules of the final system. These noise deficiencies can largely be attributed to failed S-curve fits in two ROCs, each on different modules (M0459, M0482).

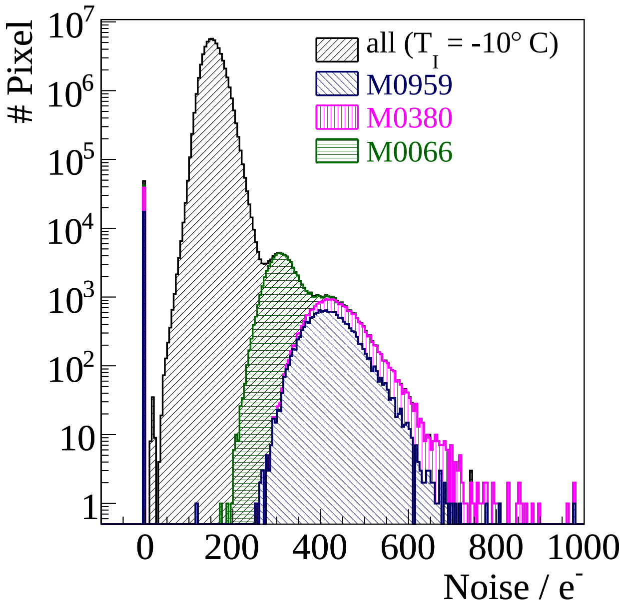

Figure 5.18: Stack of the noise distributions of modules M0959, M0380 and M0066 (all of grade C) causing several peaks in the tail of the noise distribution at −10 °C.

Deficient trimmed threshold: The distributions of the trimmed pixel thresholds measured at −10 °C and 17 °C are illustrated in Figure 5.21. If the *Vcal*-threshold of a pixel after trimming deviates more than 10 *Vcal* units from the target threshold, the trim algorithm is considered to have failed and the pixel is considered to have a trimming deficiency. Figure 5.22 shows the location of the 2492 pixels with trimming deficiencies at −10 °C on the modules of the final system. In 80 % of the cases, the trim algorithm failed to sufficiently lower the threshold to be within 10 *Vcal* units from the target threshold. The occurrence of these deficiencies is concentrated on the module edges, where the untrimmed *Vcal*-threshold per pixel tends to be higher on average (see Figure 4.28a)).

Deficient pulse height calibration: The distributions of the gain and pedestal of the linear fit and of parameter p_1 of the hyperbolic tangent fit of the pulse height calibration curve are illustrated in Figure 5.23—5.25, before and after the optimisation of the linear range. A very low gain indicates an inconsistent pulse height calibration curve or a failed linear fit. The following criteria were established based on Figure 5.24a) before the optimisation of the linear range. The gain of a pixel is considered to be deficient, if it is below 1 or above 4.5 ($T = 17\,°\text{C}$), or if it is below 0.5 or above 6.0 ($T = -10\,°\text{C}$). The few outliers in the pixel distribution on the modules in the final system disappear after the optimisation of the linear range. The upper boundary becomes redundant and only a few pixel on the modules in the final system fail to be above the lower boundary, see Figure 5.24b). Since the optimisation of the linearity in the low range aims at a target value of 1.4, pixels with p_1 above 2 or a negative p_1 value after the optimisation are considered to be deficient, see Figure 5.25b). Figure 5.26 shows the location of the 6544 pixels with a deficient parameter p_1 at $-10\,°\text{C}$ on the modules of the final system. The number of pixels with a deficient parameter p_1 is significantly lower at $17\,°\text{C}$ and amounts to 226. The increase of p_1 deficiencies at $-10\,°\text{C}$ can largely be attributed to a failed p_1 optimisation in two ROCs, each on different modules (M0562, M0657).

Table 5.2: Number and fraction of pixels with deficient performance. The table shows the number of defects measured at $-10\,°\text{C}$ and $17\,°\text{C}$ for the modules in the final system. The "ROCs column" shows the number of ROCs containing one or more pixel with a performance deficiency. A dead pixel is not attributed a performance deficiency.

		Modules in final system		
	Temp.	#ROCs	#Pixels	Fraction
Tested	-10	$1.15 \cdot 10^4$	$4.79 \cdot 10^7$	-
	17	$1.15 \cdot 10^4$	$4.79 \cdot 10^7$	-
Noise	-10	583	5811	$1.2 \cdot 10^{-4}$
	17	560	3649	$7.6 \cdot 10^{-5}$
Trimmed Thr.	-10	1737	2492	$5.2 \cdot 10^{-5}$
	17	1305	1981	$4.1 \cdot 10^{-5}$
Gain	-10	114	179	$3.7 \cdot 10^{-6}$
	17	104	174	$3.6 \cdot 10^{-6}$
Parameter p_1	-10	114	6544	$1.4 \cdot 10^{-4}$
	17	94	226	$4.7 \cdot 10^{-6}$

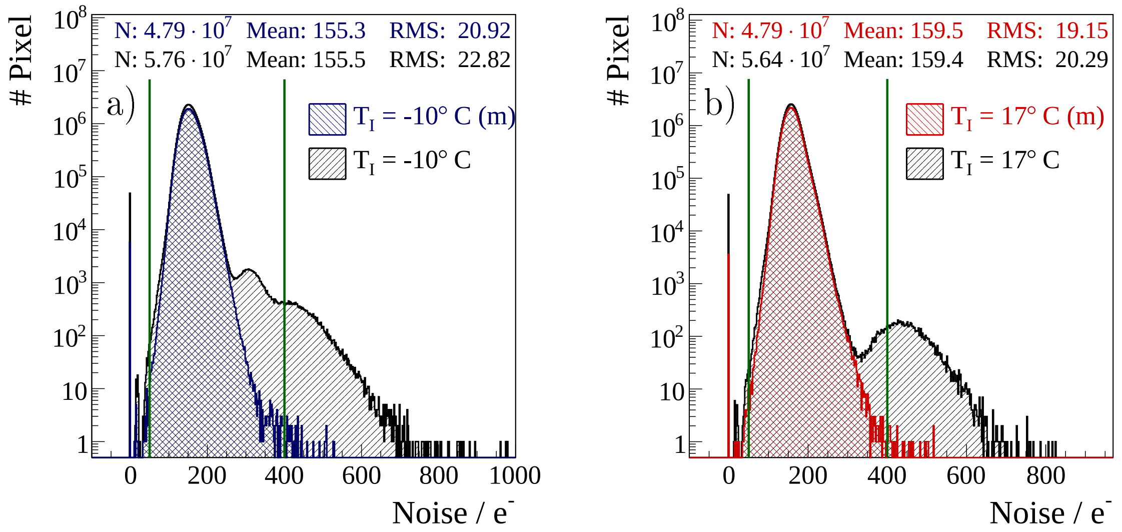

Figure 5.19: Pixel noise distributions: for the 47.9×10^6 pixels in the modules mounted in the final system (m) and a) the 57.6×10^6 pixels in all modules tested at $-10\,°\mathrm{C}$, b) the 56.4×10^6 pixels in all modules tested at $17\,°\mathrm{C}$. A pixel with a noise of less than $50\,\mathrm{e}^-$ or above $400\,\mathrm{e}^-$ is considered to be deficient. These limits are indicated by the solid green lines. The peaks in the tail of the noise distributions of all tested modules are caused by a few single modules (see Figure 5.18).

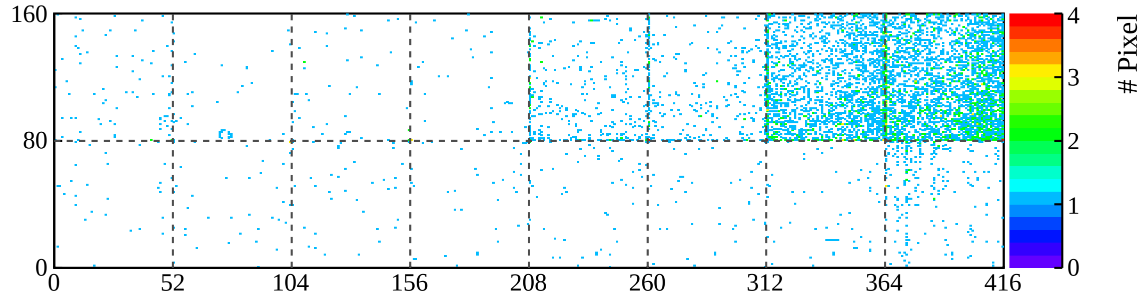

Figure 5.20: Superposition of the 768 modules in the final system (47.9×10^6 tested pixels) showing the location of the 5811 pixels with noise deficiencies, measured at $-10\,°\mathrm{C}$. The resulting fraction in the final system is $1.2 \cdot 10^{-4}$.

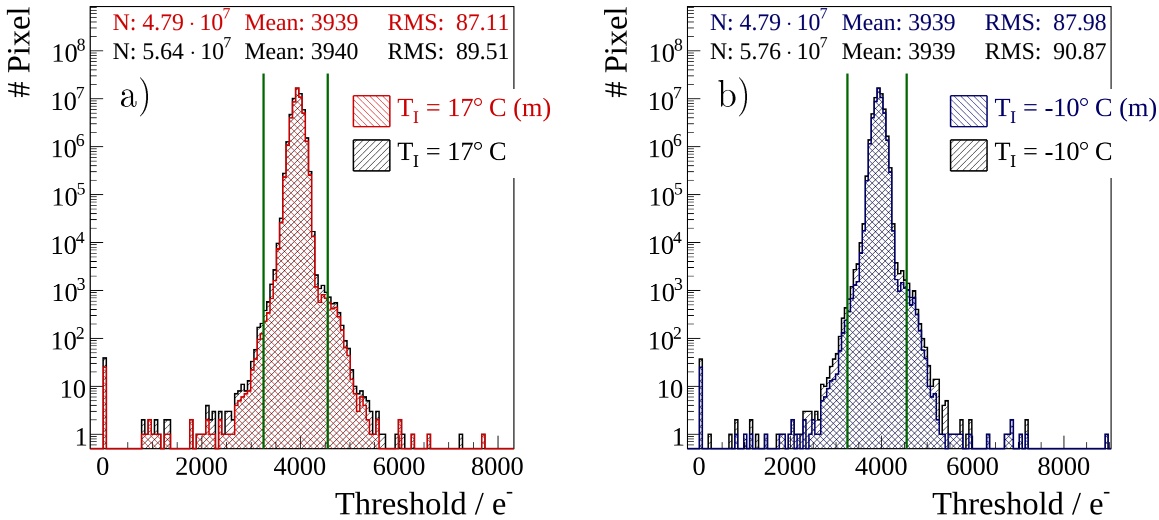

Figure 5.21: Pixel threshold distributions after trimming: for the 47.9×10^6 pixels in the modules mounted in the final system (m) and a) the 57.6×10^6 pixels in all modules tested at $-10\,°\mathrm{C}$, b) the 56.4×10^6 pixels in all modules tested at $17\,°\mathrm{C}$. A pixel with a threshold deviating more than than 10 *Vcal* units ($650\,\mathrm{e}^-$) from the target threshold after trimming—here at *Vcal* 60 ($3900\,\mathrm{e}^-$)—is considered to have a trimming deficiency. These limits are indicated by the solid green lines.

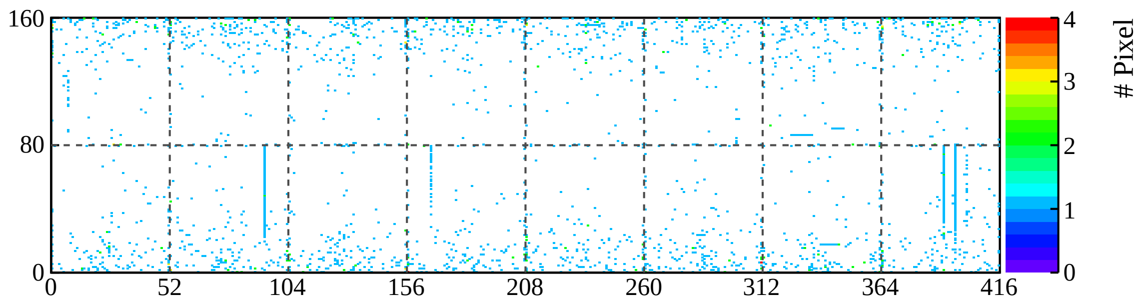

Figure 5.22: Superposition of the 768 modules in the final system (47.9×10^6 tested pixels) showing the location of the 2492 pixels with a deficient trimmed threshold, measured at $-10\,°\mathrm{C}$. The resulting fraction in the final system is $5.2 \cdot 10^{-5}$.

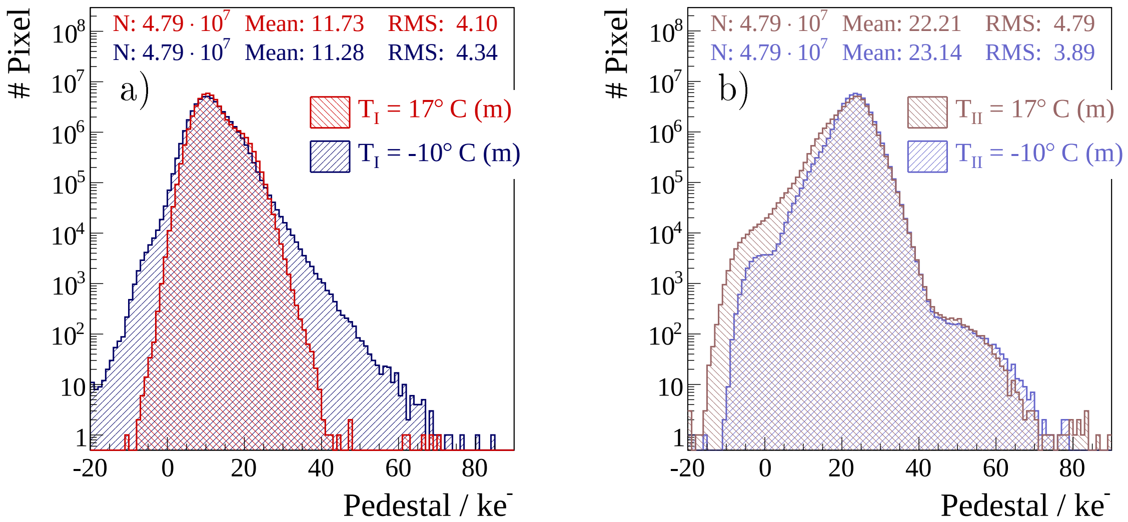

Figure 5.23: Pixel pedestal distributions at $-10\,°\mathrm{C}$ and $17\,°\mathrm{C}$ for the 47.9×10^6 pixels in the modules mounted in the final system: a) before the optimisation, b) after the optimisation of the pulse height calibration curve.

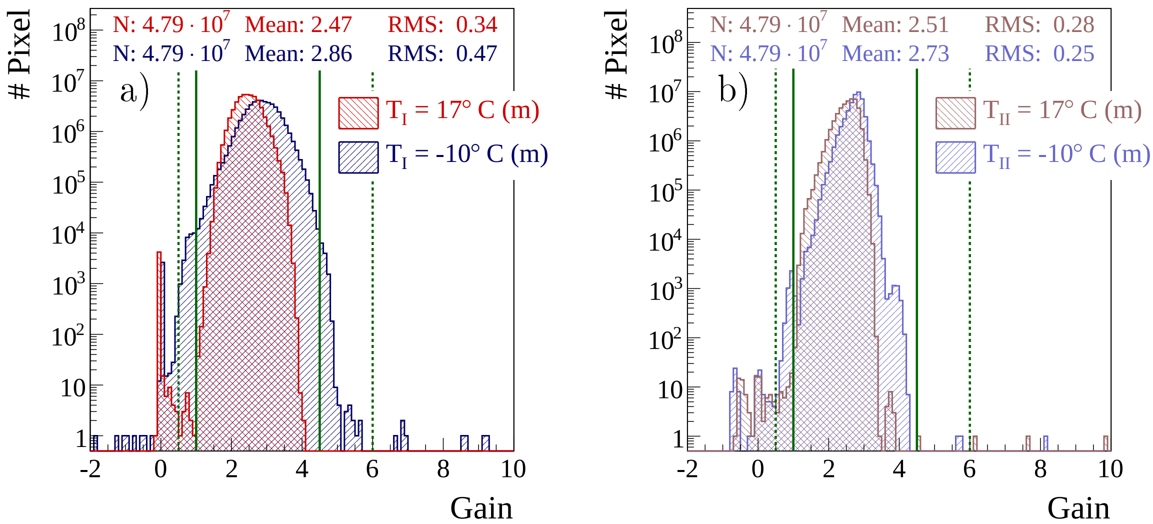

Figure 5.24: Pixel gain at $-10\,°\mathrm{C}$ and $17\,°\mathrm{C}$ for the 47.9×10^6 pixels in the modules mounted in the final system: a) before the optimisation, b) after the optimisation of the pulse height calibration curve. The gain of a pixel is considered to be deficient, if it is below 1 or above 4.5 ($T = 17\,°\mathrm{C}$), or if it is below 0.5 or above 6.0 ($T = -10\,°\mathrm{C}$). These limits are indicated by the solid and dashed green lines respectively.

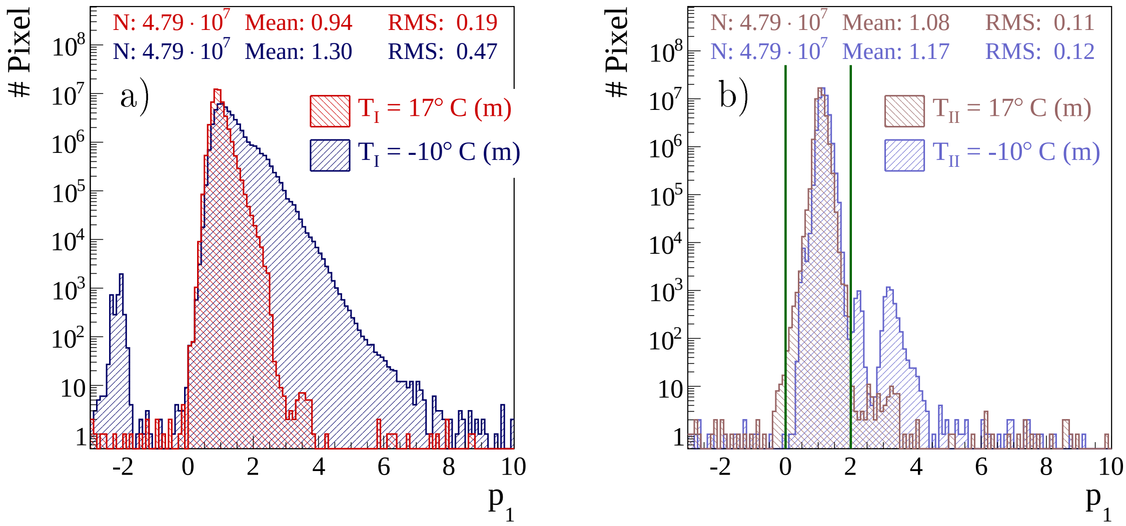

Figure 5.25: Pixel p_1 distributions at $-10\,°$C and $17\,°$C for the 47.9×10^6 pixels in the modules mounted in the final system: a) before the optimisation, b) after the optimisation of the pulse height calibration curve. Pixels with p_1 above 2 or a negative p_1 value after the optimisation are considered to be deficient, as indicated by the solid green lines.

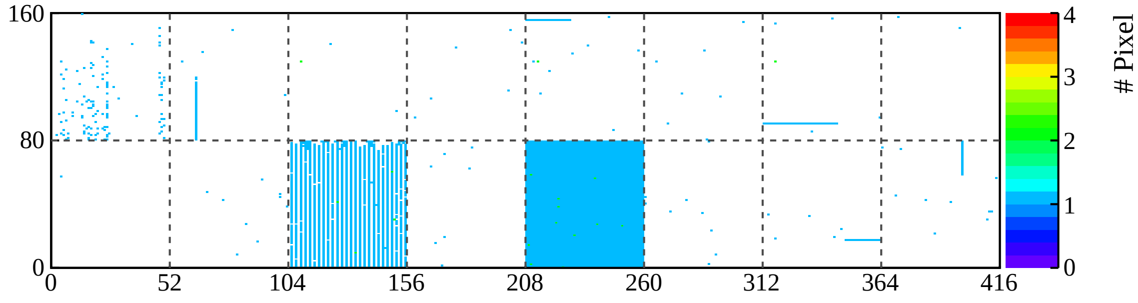

Figure 5.26: Superposition of the 768 modules in the final system (47.9×10^6 tested pixels) showing the location of the 6544 pixels with a deficient parameter p_1 on the modules of the final system, measured at $-10\,°$C. The resulting fractions in the final system is $1.4 \cdot 10^{-4}$.

5.2.4 Grading scheme

Table 5.3 summarises the grading criteria explained in sections 5.2.1— 5.2.3.

Table 5.3: Summary of qualification criteria.

	A	B	C
Defects / ROC	$\leq 1\,\%$	$\leq 4\,\%$	$> 4\,\%$
Mask defects	-	-	≥ 1
Mean Noise in e^-	< 500	< 1000	> 1000
Relative Gain Width	$< 10\,\%$	$< 20\,\%$	$> 20\,\%$
Pedestal Spread in e^-	< 2500	< 5000	> 5000
Vcal Thr. Width in e^-	< 200	< 400	> 400
$I^{meas}_{+17^o}(150V)$	$< 2\,\mu$A	$< 10\,\mu$A	$> 10\,\mu$A
$I^{recalc}_{-10^o}(150V)$	$< 3\,\mu$A	$< 15\,\mu$A	$> 15\,\mu$A
$\frac{I(150V)}{I(100V)}$	< 2	> 2	-

Only pixels with functionality defects contribute to the number of “Defects /ROC”. Pixels with performance deficiencies are not included in this number (see section 5.2.3.2). The “Relative Gain Width” is defined as the spread of the gain distribution divided by its mean. To grade a module with respect to its IV-characteristics, the current measurements at $-10\,°$C were recalculated to the corresponding current at $17\,°$C, using formula 5.1. Since the average ratio of the value recalculated from the current measured at $-10\,°$C and the actual current measurement at room temperature, tended to be around 1.5 rather than one, the reference values for grading based on measurements at $-10\,°$C were increased by a factor of 1.5.

5.3 Results

A summary plot of the module test carried out between April 2006 and March 2008 is shown in Figure 5.27. In total, 971 modules (848 full, 123 half modules) had successfully been assembled at PSI and entered the process of module qualification. Of these, 817 modules passed the first stage of module testing (*test suite I*). After the second test series (*test suite II*) carried out before the modules were mounted onto the final system, the total number of modules suitable for use in the CMS pixel barrel detector amounted to 813. Of these, 613 were of excellent quality (grade A) and 200 of grade B quality. 158 modules in total did not qualify for use in the detector system (grade C). This gives a final yield of 63 % A, 21 % B and 16 % C modules. The outcome of each testing part of the test suites is shown separately for full and half modules in Table 5.4. Figures 5.28 and 5.29 illustrate the module failure statistics in *test suite I* and *test suite II* respectively.

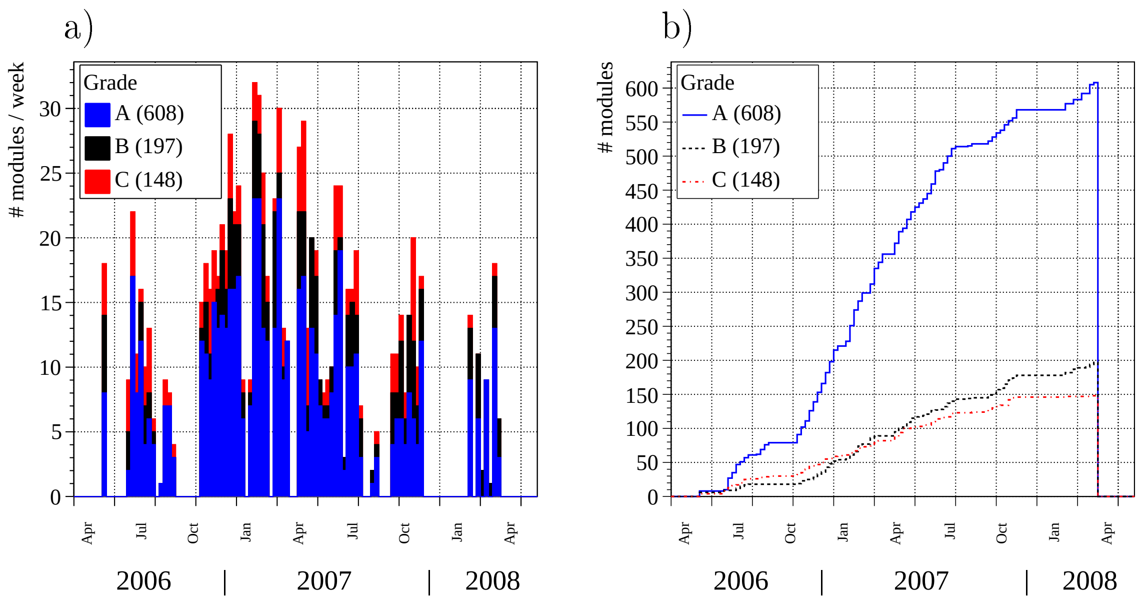

Figure 5.27: The number of produced modules between April 2006 and March 2008 as a function of time: a) number of modules produced per week, b) integrated number of produced modules as a function of time. The first 18 modules assembled as part of a pre-production in December 2005 are not shown. Adding the modules from the pre-production, 613 qualified as A, 200 as B and 158 as C—giving the following yields A:63 %, B:21 % and C: 16 %.

Table 5.4: The resulting grades from the different parts of *test suites I* and *II* are shown for full and half modules separately. The tests are listed in chronological order. The grades of the first test sequence at −10 °C do not include the sensor leakage current criteria, since the IV-characteristics of a module was only measured *after* the thermal cycling (once at −10 °C and 17 °C). Of 971 tested modules, 613 qualified as A, 200 as B and 158 as C. The combined final yields of modules for the three given grades are A:63 %, B:21 % and C:16 %.

		Full Modules			Half Modules		
		A	B	C	A	B	C
I	$T = -10\,°C$	702	76	70	111	5	7
I	Thermal cycling						
I	$T = -10\,°C$	597	148	103	97	18	8
I	$T = +17\,°C$	664	97	87	100	12	11
Overall Grade I		563	149	136	85	20	18
II	$T = -10\,°C$	634	81	16	97	6	3
II	$T = +17\,°C$	683	40	8	103	3	0
Overall Grade II		625	88	18	96	7	3
Final Grade*		533	177	138	80	23	20
Yield		63 %	21 %	16 %	65 %	19 %	16 %

*after regrading (in specific cases). See main text.

Of the 837 modules entering *test suite II*, 20 modules that had only narrowly failed to meet grade B criteria in *test suite I* were included. Following review of the results of both test suites, nine of the 20 modules were considered to be within the requirements for use in the detector system and were regraded to category B. Of the 817 modules that had passed the first test stage, 17 modules failed in *test suite II*, generally due to a significant increase of the leakage current at −10 °C. Their leakage current measured at 17 °C in *test suite II*, however, did not exceed the grade B requirement and they had also exhibited excellent IV-characteristics at both temperatures in the first test suite.[1]

[1] The accuracy of the IV-curve measurement in *test suite I* is $2.1 \cdot 10^{-3}\,\mu A$. The current established in a single point measurement in *test suite II* is only accurate within 1 % of the measured value.

Of the 17 modules, four were therefore graded as B. Taking these special cases into account, resulted in a total of 813 modules qualified for use in the detector system.

Grade B failures in *test suite I* were generally attributed to sensor IV-characteristics, in particular the IV-slope (see Figures 5.28a)). Second most relevant were functionality defects on one or several ROCs, caused in particular by bump bonding defects (with a fraction of $7.0 \cdot 10^{-4}$ in the entire production, defective bump bonds constituted the main contribution to pixel functionality defects, followed by dead pixels with a fraction of $2.1 \cdot 10^{-5}$). Note that failures attributed to large pedestal and gain spreads at *test suite I* decreased significantly following optimisation of the pulse height calibration curve in *test suite II*. Grade C failures in *test suite I* comprised of either high leakage currents or large amounts of defective pixels on a ROC (see Figures 5.28b)). A small contribution also arose from failed DTL scans. Contributions from the spread of the trimmed threshold and the mean noise were negligible or non-existent.

As in *test suite I*, the number of grade B failures in *test suite II* was generally attributed to a high sensor leakage current (see Figure 5.29a)). ROC failures attributed to pixel functionality defects are not present in the grade B failure statistics of *test suite II*, since only the pixel readout test was repeated at that stage and the number of dead pixel per ROC was below 1 %. As mentioned before, performance related problems decreased after the optimisation of the pulse height calibration. A significant increase in leakage current at $-10\,^{\circ}\mathrm{C}$ from *test suite I* to *test suite II*, resulted in several modules failing the grade B criteria (see Figure 5.29b)) and being categorised as grade C (with the exceptions mentioned above).

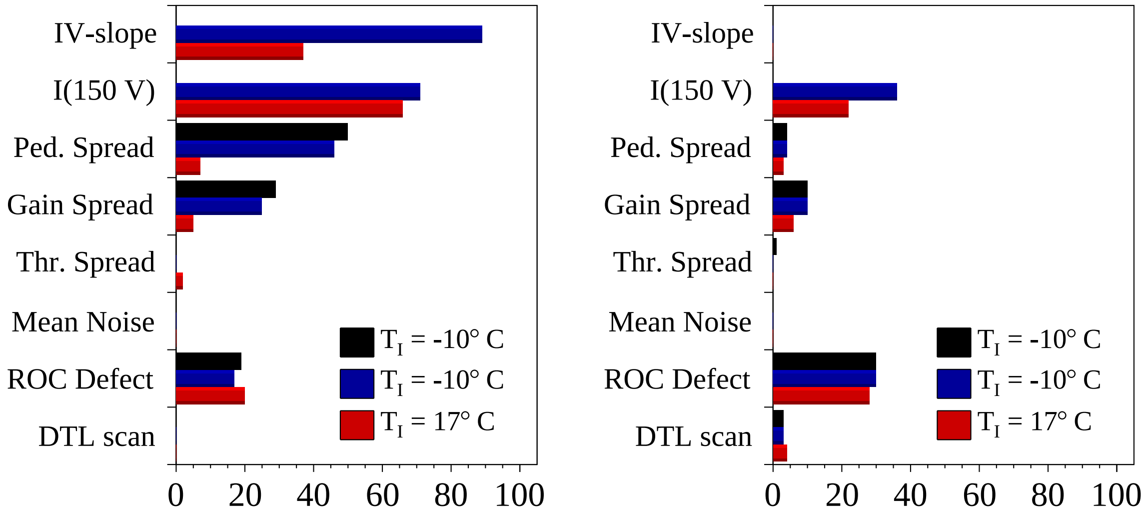

Figure 5.28: Module failure statistics in *test suite I* at −10 °C and 17 °C: failure reason for grade B (left) and grade C (right). The black bars show the failure statistics at −10 °C *before* thermal cycling and do not include the sensor leakage current criteria, since the *IV*-characteristics of a module was only measured *after* the thermal cycling, once at −10 °C and 17 °C.

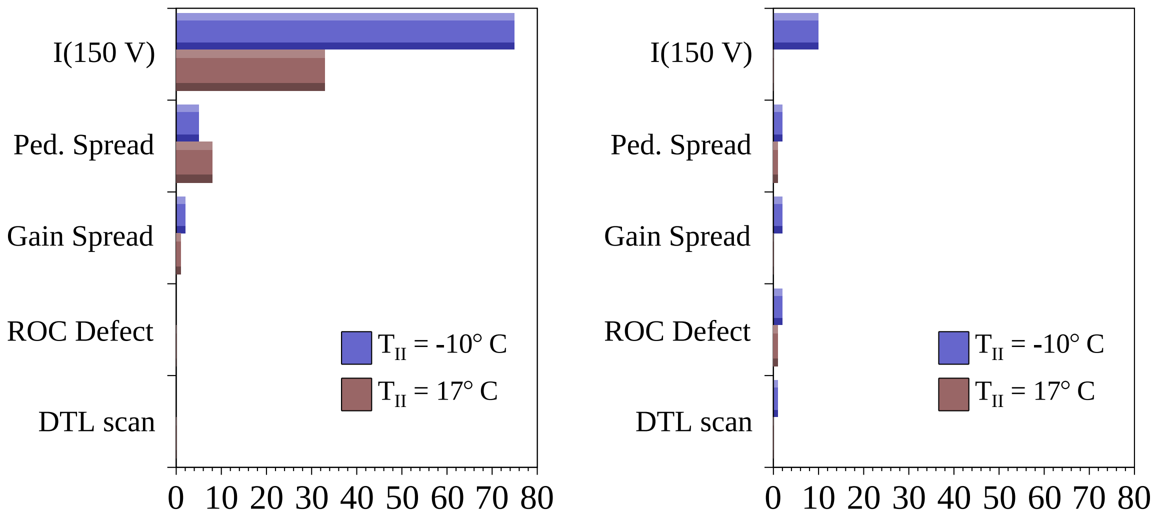

Figure 5.29: Module failure statistics in *test suite II* at −10 °C and 17 °C: failure reason for grade B (left) and grade C (right).

Part II

Search for $B_s^0 \to \mu^+\mu^-$

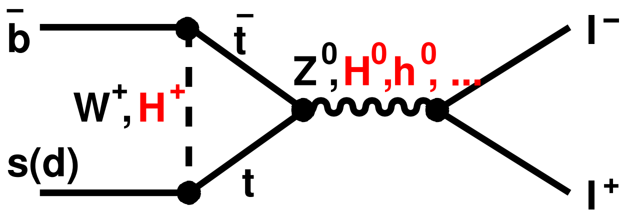

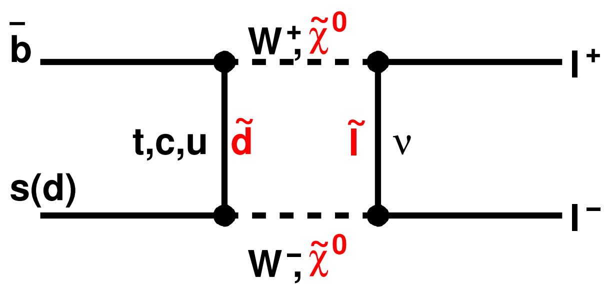

Chapter 6

B-physics

In the Standard Model of particle physics, matter is made up of two elementary particles: quarks and leptons. Both of them come in six flavours and can be arranged into three families (generations) differing only by their masses. The lightest elementary particles of the first family are the constituents of the stable matter in the universe. The quarks of the first family—up and down quarks—form the building block for atomic nuclei. Together with the electron, a lepton of the first family, they constitute the atoms. In nature, quarks are never found separated from each other. Five of the six known quarks always form composite particles called hadrons. The most massive known elementary particle—the top quark—however, decays weakly on a time scale that is too short for the top quark to form a hadron.[1] The third-generation partner of the top quark is the bottom quark. Hadrons containing a bottom quark are thus the heaviest that are experimentally accessible. Offering a multitude of diverse physics opportunities, b-hadron systems allow us to test the consistency of the Standard Model as well as to study new physics effects. The research in B-physics focuses on two main goals, which are to study the structure of quark (flavour) mixing and to explore the phenomenon of CP violation.

A short historical introduction to the discovery of the bottom quark in the context of the rise of the Standard Model is given in section 6.1. The basic concept of the Standard model are summarised in section 6.2, with the emphasis on flavour mixing and CP violation. Section 6.3 outlines the major goals of B-physics and the prospects at LHC. Finally, the production mechanism of B mesons will be explained in section 6.3.

[1] A top quark decays almost exclusively into a bottom quark and a W boson ($t \to bW$).

6.1 The Discovery of the Bottom Quark

The bottom quark, a quark of the third quark generation with charge $-\frac{1}{3}$, was discovered in the 1970s—a decade of remarkable experimental and theoretical progress in particle physics. By the end of that decade, a single model had been established in which the three fundamental forces in nature—the electromagnetic, the weak and the strong forces—could be described by three closely related gauge theories.

In the 1960s, the known elementary constituents comprised of the two known lepton pairs—the electron e and its neutrino ν_e, the muon μ and its neutrino ν_μ—and the three known quarks—u, d and s. The series of new particles discovered in the 1970s had already been hinted at in earlier theoretical papers: a mechanism, proposed by Glashow, Iliopoulos and Maiani in 1969 to explain the absence of strangeness-changing weak currents, required the existence of a fourth quark. Kobayashi and Maskawa concluded in their paper 1973, that a model with only two quark families could not account for the violation of CP invariance, that had been measured in decays of the K^0_L-meson almost a decade before. The most natural explanation implied the existence of a third quark family. A revolution of sorts started with the observation of the J/ψ meson, a bound state of the charm quark and its antiquark, in 1974. The same year evidence occurred for the existence of a heavy lepton, the τ-lepton. After several years of confusion and controversy about a third lepton family, the τ-lepton and its neutrino ν_τ were confirmed around 1978. In 1977 Leon Lederman[1] finally discovered the Υ meson at Fermilab. The Υ was immediately interpreted as a bound state of a new type of quark and antiquark—the bottom quark. The picture of particle physics finally settled on three generation of quarks and leptons, and culminated in the emergence of Standard Model (SM).

After the discovery of the Υ, the existence of B mesons was a logical consequence and was soon confirmed by measurements. The bound states of hadrons containing one b quark that have been confirmed up to date are listed in Table 6.1 together with their masses and lifetimes.

[1] who had missed the discovery of the J/ψ at Brookhaven in 1968 due to an insufficient mass resolution

Table 6.1: b-hadrons [45].

B mesons				
Particle	Quark Content	Mass [MeV]	Mean Life [ps]	$c\tau$ [μm]
B^+	$u\bar{b}$	5279.15 ± 0.31	1.638 ± 0.011	491.1
B^0	$d\bar{b}$	5279.53 ± 0.33	1.530 ± 0.009	458.7
B^0_s	$s\bar{b}$	5366.3 ± 0.6	$1.470^{+0.026}_{-0.027}$	441
B^+_c	$c\bar{b}$	6276 ± 4	0.46 ± 0.07	-
b-baryons				
Particle	Quark Content	Mass [MeV]	Mean Life [ps]	$c\tau$ [μm]
Λ^0_b	udb	5620.2 ± 1.6	$1.383^{+0.049}_{-0.048}$	415
Σ^+_b	uub	5807.8 ± 2.7	-	-
Σ^-_b	ddb	5815.2 ± 2.0	-	-
Ξ^0_b / Ξ^-_b	usb / dsb	5292.4 ± 3.0	$1.42^{+0.28}_{-0.24}$	-
Ω^-_b	ssb	6165 ± 23	-	-

6.2 The Standard Model

Astoundingly, in all subsequent experiments, the Standard Model has proven to be an accurate theory, describing all the detected particles and their interactions. The neutral component (Z^0 boson) and the charged components ($W^\pm$ bosons) of the weak interaction, were discovered in 1983 at CERN. Owing to its large mass, the second quark of the third generation—the top quark—was only confirmed in 1995 at Fermilab. The basic concepts of the Standard model can be summarised as follows

- The twelve spin-$\frac{1}{2}$ particles (*fermions*) can be grouped with respect to their interaction properties: the six *quarks*—up (u), down (d), charm (c), strange (s), bottom (b) and top (t)—interact strongly, whereas the six *leptons*—the electron (e) and its neutrino (ν_e), the muon (μ) and its neutrino (ν_μ), and the tau (τ) and its neutrino (ν_τ)—do not partake in the strong interaction.

- The fundamental interactions of these particle are described by the gauge group

$$SU(3)_C \otimes SU(2)_L \otimes U(1)_Y,$$

where C refers to color (QCD), L to left-handed fields and Y denotes the weak hypercharge generators. The interactions are mediated by spin-1 particles (*bosons*): eight massless *gluons* G_α and one massless *photon* γ, for the strong and electromagnetic interaction respectively, and three massive *gauge bosons* W^+, W^- and Z^0 for the weak interaction.

- The $SU(2)_L \otimes U(1)_Y$ group undergoes spontaneous symmetry breaking to the electromagnetic subgroup $U(1)_Q$, giving rise to a physical scalar (spin-0) particle, known as the *Higgs boson* H.

- The matter fields (quarks and leptons) are Dirac fields and obtain their masses from Yukawa couplings to the field of a Higgs particle. The same field also generates masses for the gauge bosons of the weak interaction.

- The charged current couplings for the transition of a down-type quark to an up-type quark are described in terms of a unitary 3×3 matrix, known as *CKM matrix*[1]

$$\mathbf{V}_{CKM} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix}. \tag{6.1}$$

  In the original Standard Model concept with massless neutrinos, the analogous matrix in the lepton sector is a unit matrix.

A minor revision of the Standard Model was necessary to accommodate the evidence for neutrino oscillations in 1998 implying that neutrinos must have a mass. As in the quark sector, the mixing can be described by a 3×3 matrix, called the PMNS matrix[2]. Including the additional neutrino mass terms, the Standard Model has 26 free parameters. To date, one last particle of the Standard Model—the Higgs particle—remains yet to be discovered.

[1] after Cabibbo, Kobayashi and Maskawa
[2] after Pontecorvo, Maki, Nakagawa and Sakata

6.2.1 Flavour Physics and *CP* Violation

Concerning electroweak interactions, the left-handed fermions can be arranged into $SU(2)_L$ doublets, whereas the right-handed fields transform as singlets under $SU(2)_L$. Together they form three families with identical interaction properties. The three generations differ only by their masses and flavour quantum numbers

$$\begin{pmatrix} u \\ d' \end{pmatrix}_L \quad \begin{pmatrix} c \\ s' \end{pmatrix}_L \quad \begin{pmatrix} t \\ b' \end{pmatrix}_L \qquad u_R,\ d_R,\ c_R,\ s_R,\ t_R,\ b_R \tag{6.2}$$

$$\begin{pmatrix} e \\ \nu_e \end{pmatrix}_L \quad \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix}_L \quad \begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}_L \qquad e_R,\ \mu_R,\ \tau_R,\ \nu_{eR},\ \nu_{\mu R},\ \nu_{\tau R}. \tag{6.3}$$

The weak eigenstates d', s', b' are a mixture of the corresponding mass eigenstates d, s, b and are connected through the CKM *quark-mixing matrix*

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \mathbf{V}_{CKM} \begin{pmatrix} d \\ s \\ b \end{pmatrix}. \tag{6.4}$$

The so-called global CKM fit uses all available measurements and imposes SM constraints to determine the magnitudes of the CKM elements, and leads to the following result [46]

$$\mathbf{V}_{CKM} = \begin{pmatrix} 0.97419 \pm 0.00022 & 0.2257 \pm 0.0010 & 0.00359 \pm 0.00016 \\ 0.2256 \pm 0.0010 & 0.97334 \pm 0.00023 & 0.0415^{+0.0010}_{-0.0011} \\ 0.00874^{+0.00026}_{-0.00037} & 0.0407 \pm 0.0010 & 0.999133^{+0.000044}_{-0.000043} \end{pmatrix}. \tag{6.5}$$

The mixing between the second and third family is much smaller than the mixing between the first and second family. The mixing between the first and third family is even more suppressed. The freedom to define the global phase of the quark fields allows to reduce the initial parameters of the unitary 3×3 matrix from nine to four. In the standard parametrisation, recommended by the particle data group, the CKM matrix is represented by the product of three complex rotation matrices. The rotations are characterised by three Euler angles θ_{12}, θ_{13}, θ_{23} and one complex phase δ. Expanding each element in this matrix as a series of $\lambda = \sin\theta_{12} = |V_{us}| \approx 0.22$ leads to the

Wolfenstein parametrisation, an approximate parametrisation that nicely displays the hierarchical pattern of the matrix in powers of λ:

$$\mathbf{V}_{CKM} = \begin{pmatrix} 1-\frac{\lambda^2}{2} & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1-\frac{\lambda^2}{2} & A\lambda^2 \\ A\lambda^3(1-\rho-i\eta) & -A\lambda^2 & 1 \end{pmatrix} + \mathcal{O}(\lambda^4). \tag{6.6}$$

The transitions $b \to c$ and $b \to u$ are suppressed by a factor of λ^2 and λ^3 respectively. The hierarchical structure of the CKM matrix first became apparent, when the lifetimes of B mesons was shown to be much longer than expected ($\sim 10^{-12}\,$s). The complex phase also allows the accommodation of *CP violating phenomena* that have been observed in the neutral kaon system—or more recently in the neutral B meson system—within the flavour mixing matrix.

CP stands for the product of charge conjugation (C) and parity inversion (P).

Weak interactions involving $W^{\pm}$ bosons interact exclusively with left-handed particles or right-handed antiparticles. Therefore, the interaction is not invariant under charge conjugation or parity inversion and hence violates both C-symmetry and P-symmetry in a maximal way. Naturally, it would be assumed that *CP*-symmetry—the combination of the two—would be preserved. However, in 1964 Cronin and Fitch discovered in the decays of neutral kaons, that this is not exactly true and that the weak interaction does violate the *CP*-symmetry. The only possible source of *CP* violation in the Standard Model is provided by the complex phase in the CKM matrix. The standard parametrisation of the CKM matrix—given in equation 6.7—shows that the complex phase is always multiplied by $\sin\theta_{13} = |V_{ub}| \sim 10^{-3}$. Therefore, *CP* violation is clearly suppressed in the Standard Model (independent of the magnitude of the phase δ). Since new sources of *CP* violation are present in most extensions of the Standard Model, measurements of *CP* violation provide an excellent probe of new physics.

$$\mathbf{V}_{CKM} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}e^{-i\delta} & c_{12}c_{23} - s_{12}s_{23}e^{-i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}e^{-i\delta} & -c_{12}s_{23} - s_{12}c_{23}e^{-i\delta} & c_{23}c_{13} \end{pmatrix} \tag{6.7}$$

where $c_{ij} = cos\theta_{ij}$, $s_{ij} = \sin\theta_{ij}$ and $s_{12} = |V_{us}|$, $s_{13} = |V_{ub}|$, $s_{23} = |V_{cb}|$. The angles θ_{ij} are also referred to as quark mixing angles and the complex phase δ accounts for *CP* violation.

Contrary to quark transitions induced by charged currents, the unitarity of the CKM matrix forbids neutral current couplings to the Z^0 changing the flavour but not the charge of a fermion.[1] In the Standard Model *flavour changing neutral current* FCNC processes are thus forbidden at the tree-level. They can, however, proceed through higher order diagrams involving flavour changing $W^{\pm}$ vertices. At the one-loop level, FCNC processes can be described by *penguin* and *box* diagrams, that are composed of a set of basic triple and quartic effective vertices respectively. Nevertheless, these processes are highly suppressed in the Standard Model by the GIM mechanism. The FCNC sector is therefore of particular interest in the search of new physics (NP): whereas the effects of NP in most realistic models can safely be neglected in transitions mediated by Standard Model tree-level processes, NP can have a significant impact on FCNC amplitudes. Beyond the Standard Model (BSM) new particles may enter into the penguin and box diagrams and new tree-level contributions to FCNC processes may be generated.

6.3 The Goals of B-physics

The main focus of B-physics lies on the verification of the Standard Model by exploring the nature of quark mixing and its role in CP violation as well as on probing physics beyond the Standard Model (BSM) in rare FCNC processes in B meson systems.

The CKM picture of quark mixing and CP violation has been confirmed quantitatively in precise measurements of many different B-decay modes, that over-constrain the CKM matrix. The elements of $|V_{ub}|$ and $|V_{cb}|$ have been measured in *exclusive* and *inclusive* analyses of semileptonic B-decays. In exclusive decays, all decay products in the final state are identified and measured. In inclusive decays, all (or a special class) of the accessible final states are summed up. In the heavy quark limit, exclusive and inclusive decays can be treated in the two (distinct) approaches of heavy quark

[1] The unitarity conditions of the CKM matrix and the flavour-changing parameters Z_{ij} of the down-type quarks are connected through the following relations

$$\begin{aligned} V_{ud}^* V_{us} + V_{cd}^* V_{cs} + V_{td}^* V_{ts} &= Z_{ds}, \\ V_{ub}^* V_{ud} + V_{cb}^* V_{cd} + V_{tb}^* V_{ts} &= Z_{bd}, \\ V_{ub}^* V_{us} + V_{cb}^* V_{cs} + V_{tb}^* V_{ts} &= Z_{bs}. \end{aligned}$$

effective theory (HQET) and the heavy quark expansion (HQE), respectively.[1] Since the Standard Model penguin and box diagrams for FCNC processes are dominated by virtual top contributions, the elements $|V_{td}|$, $|V_{ts}|$ and $|V_{tb}|$ of the quark mixing matrix indirectly follow from the measurements of FCNC processes. The system of B mesons also offers a variety of processes to study CP violation that, in the Standard Model, have their only source in the complex phase of the CKM matrix. The parameters ρ and η in the Wolfenstein parametrisation are related to CP violation. The unitarity of the CKM matrix can be represented by the *unitarity triangle* in the (ρ, η)-plane. The angles α, β and γ of the triangle are related to CP violating asymmetries that can be measured in non-leptonic B-decays[2].

CP violation also provides an excellent probe of new physics, since most extensions of the Standard Model exhibit new sources of CP violation. FCNC processes in the B meson system, such as particle-antiparticle mixing and rare decays, are equally important in the quest for new physics. The transition amplitudes can be significantly enhanced by new particles contributing to the box diagrams or even at tree-level.

6.3.1 Facilities for B-physics

CLEO [47] and ARGUS [48] were the first experiments that studied B-decays at e^+e^- colliders by running at the $\Upsilon(4S)$ resonance. This resonance can decay to B_u and B_d but not to B_s. The current e^+e^- storage rings PEP II and KEKB, with their associated experiments BaBar [49] and Belle [50], run at asymmetric energies at the $\Upsilon(4S)$ resonance. The energy asymmetry produces a $\Upsilon(4S)$ boosted along the beam axis and allows the decay vertices of the B meson to be resolved. Recently, the Belle experiment has also been studying B_s decays by taking data at the $\Upsilon(5S)$ resonance [51].

At higher energy e^+e^- collider, such as the LEP, $b\bar{b}$ pairs can also be produced at the Z^0 pole where the full spectrum of b-hadrons is accessible. The production mechanisms at hadron colliders are more complex than at e^+e^- colliders. The momenta and directions of the b-hadrons vary over a large range. Hadron colliders also suffer from a very high background in b events and include a complicated underlying event in addition to the produced b-hadrons. Nevertheless, hadron colliders benefit from the

[1] The former treats the heavy quark in a meson as a static source of the gluon field (similar to considering a hydrogen atom), and in the latter the decay rate can be expanded in inverse powers of m_b, with the leading term describing the decay of a free b quark.

[2] for instance, in the "gold plated" $B^0 \to J/\psi K^0_S$ decay that is used to determine the CKM angle β

production of all species of b-hadrons and they have a much higher $b\bar{b}$ production cross-section that compensate for the clean environment in e^+e^- B factories.[1] At Tevatron, the cross-section is $\sim 100\,\mu$b at $\sqrt{s} = 1.96\,$TeV. At Tevatron both experiments, CDF and D0, pursue a rich B-physics programs that complements the B factories. The area of research include the study of CP violation, mixing and lifetime measurements, rare decays and B production, fragmentation and spectroscopy. In particular CDF has observed $B_s - \bar{B}_s$ mixing and determined the oscillation frequency as $\Delta m_{B_s} = (117.0 \pm 0.8) \times 10^{-10}\,$MeV. With an expected cross-section of $\sigma_{b\bar{b}} \sim 500\,\mu$b and a design luminosity of $\mathcal{L} = 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, the LHC provides an excellent opportunity for B-physics studies. LHCb will also be the first experiment dedicated to B-physics at a hadron collider.

6.4 B Production Mechanisms at the LHC

There are several mechanisms contributing to heavy flavour production at hadron colliders, arising from the following three processes: *flavour creation*, *flavour excitation* and *gluon splitting*.

Flavour creation: The leading order (LO) processes are gluon gluon fusion $gg \to Q\overline{Q}$ or quark annihilation of light quarks $qq \to Q\overline{Q}$ shown in Figures 6.1 and 6.2 respectively. At LHC and Tevatron, gluon-gluon fusion processes are the dominant $b\bar{b}$ production mechanism out of the two hard processes. In the center-of-mass frame, the quark and antiquark are produced back-to-back and are therefore also back-to-back in the plane transverse to the beam direction.

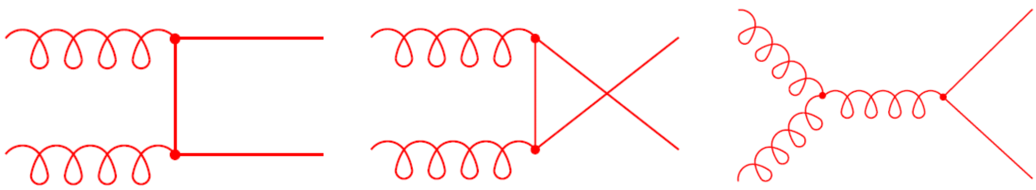

Figure 6.1: Leading order $\mathcal{O}(\alpha_s^2)$ diagrams for $b\bar{b}$ pair production: Gluon-gluon fusion.

[1]The B factories running at the $\Upsilon(4S)$ resonance have a cross-section of $\sigma_{b\bar{b}} \sim 1\,$nb and running at the Z^0 pole at LEP gives a cross-section of $\sigma_{b\bar{b}} \sim 7\,$nb.

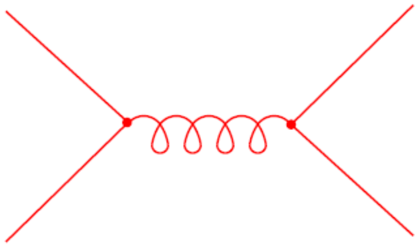

Figure 6.2: Leading order $\mathcal{O}(\alpha_s^2)$ diagrams for $b\bar{b}$ pair production: Quark annihilation.

Flavour excitation: In this next-to-leading order process, a heavy quark is assumed to be already present in the quark sea of the proton. It is then put on mass-shell by scattering against a parton of the other proton, as shown in Figure 6.3 for $Qg \rightarrow Qg$. Since the b is not a valence flavour it must originate from $g \rightarrow Q\overline{Q}$. In flavour excitation, only one of the b quarks undergoes a hard QCD scattering and therefore usually only one of the quarks from the $b\bar{b}$ pair is produced at high p_T.

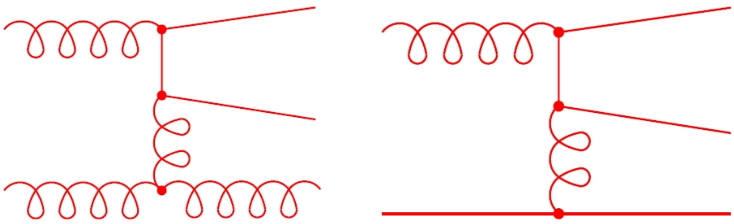

Figure 6.3: Next-to-leading order $\mathcal{O}(\alpha_s^3)$ diagrams for $b\bar{b}$ pair production: Flavour excitation.

Gluon splitting: In this next-to-leading order process, the heavy quarks arise from $g \rightarrow Q\overline{Q}$ in either the initial state or final state shower (see Figure 6.4). Here the dominant source is gluons from the final state showers, and the hard QCD process involves gluons and light quarks and antiquarks. The $b\bar{b}$ pairs from gluon splitting are usually very close in phase space and the p_T spectrum increases logarithmically.

Additional next-to-leading order arise from $\mathcal{O}(\alpha_s^3)$ corrections to the parton fusion process, that include real and virtual gluon emission. The three categories are characterised by 2, 1 or 0 heavy flavour quark(s) participating in the hard interaction. It has been shown, that next-to-leading order processes are actually larger than leading

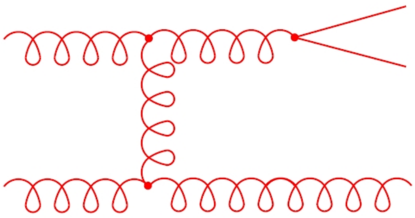

Figure 6.4: Next-to-leading order $\mathcal{O}(\alpha_s^3)$ diagrams for $b\bar{b}$ pair production: Gluon splitting.

order processes at energies larger than m_Q.[1] Figure 6.5 shows the total $b\bar{b}$ cross-section as a function of the center-of-mass energy. The dominant contribution to $\sigma_{b\bar{b}}$ at LHC energies arise from flavour excitation, followed by pair creation and gluon splitting.

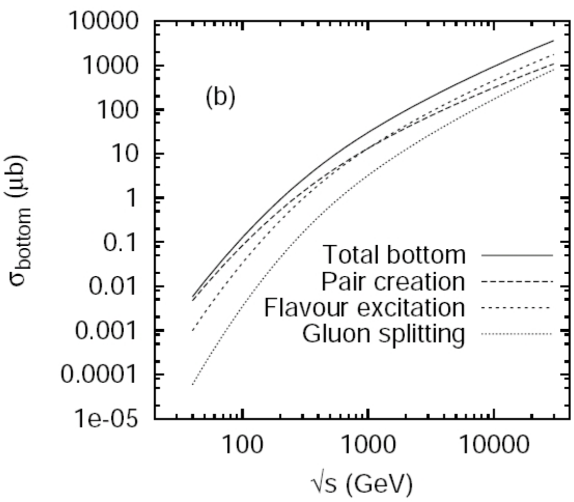

Figure 6.5: The total $b\bar{b}$ cross-section as a function of the center-of-mass energy $E_{CM} = \sqrt{s}$ at pp-collision, and the different contribution from pair creation, flavour excitation and gluon splitting [52].

[1] The cross-section for the production of gluons through $gg \to gg$ is order of magnitude larger than the leading order contribution of $gg \to Q\overline{Q}$

Chapter 7

The Search for $B_s^0 \to \mu^+\mu^-$

The leptonic decays $B_q^0 \to \ell^+\ell^-$ (where $q = s, b$ and $\ell = e, \mu$) have a highly suppressed rate in the Standard Model (SM), since they involve a $b \to s(d)$ transition. In the SM these flavour changing neutral current (FCNC) transitions are forbidden at tree-level and can only proceed through high-order diagrams, that are described by electroweak penguin and box diagrams at the one loop level (see Figure 7.1). The dominant contribution stems from the Z-penguin diagram. There are no contributions from a Standard Model Higgs to the penguin diagram, since a Higgs boson couples to fermions with Yukawa couplings $y_b \propto m_b/M_W$ and $y_\ell \propto m_\ell/M_W$. Photonic penguins also do not contribute to the decay, since the lepton-anti-lepton pair with zero angular momentum has charge conjugation number $C = 1$, whereas the photon has $C = -1$. The box diagram is suppressed by a factor of m_W^2/m_t^2 with respect to the Z-penguin.

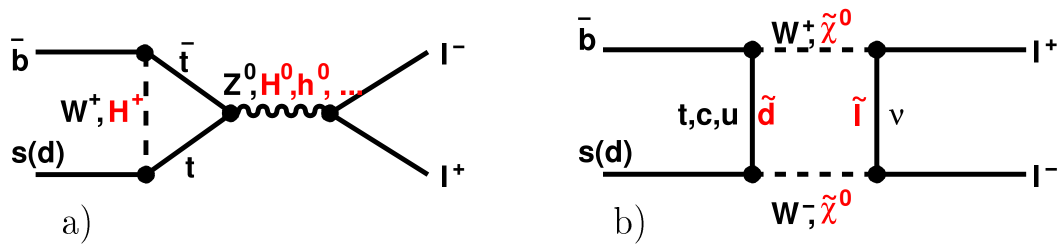

Figure 7.1: Illustration of the rare decays $B_q^0 \to \ell^+\ell^-$. In the SM, these decays proceed through $W^\pm$ and Z^0 bosons in Z-penguin (a) and box (b) interactions. In SM extensions, new particles (*e.g.* neutralinos $\widetilde{\chi}^0$, Higgs bosons and supersymmetric partners of the quarks and leptons) can contribute to the process.

In addition to the electro-weak loop suppression, these decays are helicity suppressed in the SM by a factor of m_ℓ^2/m_B^2. Since these decays are highly suppressed in the Standard Model, they are potentially sensitive probes of physics beyond the SM, where new particles can enter the diagram (see Figure 7.1) and can thereby increase the expected branching fraction by orders of magnitude. To date these decays have not been observed and the current best limits from CLEO [47], Belle [50], *BABAR* [49], D0 [53] and CDF [54] are given in Table 7.1 together with the SM expectation.

Table 7.1: The expected branching ratios for the decays $B_q^0 \to \ell^+\ell^-$ (where $q = s, b$ and $\ell = e, \mu$) in the Standard Model [55] and the current best upper limits (U.L.) at the 90 % C.L. from various experiments.

Mode	$B_s^0 \to \mu^+\mu^-$	$B_d^0 \to \mu^+\mu^-$	$B_d^0 \to e^+e^-$	$B_d^0 \to e^\pm\mu^\mp$
SM Expect. [55]	$(3.86 \pm 0.15) \times 10^{-9}$	$(1.06 \pm 0.04) \times 10^{-10}$	$(2.49 \pm 0.09) \times 10^{-15}$	~ 0
CLEO [47]	-	6.1×10^{-7}	8.3×10^{-7}	15×10^{-7}
BELLE [50]	-	1.6×10^{-7}	1.9×10^{-7}	1.7×10^{-7}
BABAR [49]	-	5.2×10^{-8}	11.3×10^{-8}	9.2×10^{-8}
D0 [53]	7.5×10^{-8}	-	-	-
CDF [54]	4.7×10^{-8}	1.50×10^{-8}	-	-

The searches for the rare B-decays at the $\Upsilon(4S)$ resonance, *i.e.* the CLEO, Belle and *BABAR* experiments, have no sensitivity to B_s decays (see Chapter 6). However, the CDF and D0 experiments at the Tevatron have sensitivity to the decay $B_s^0 \to \mu^+\mu^-$. The D0 experiment cannot discriminate between the decays $B_s^0 \to \mu^+\mu^-$ and $B_d^0 \to \mu^+\mu^-$ because of its limited mass resolution.[1] With $2\,\text{fb}^{-1}$ of integrated luminosity so far, neither D0 nor CDF have found evidence for the decay. The Tevatron will be unlikely to integrate enough luminosity for these experiments to measure this process at the SM expectation. Their current analyses are both tuned for high efficiency and are limited by backgrounds. The lowest experimental upper limit on the SM branching fraction of $B_s^0 \to \mu^+\mu^-$ to date comes from CDF and is about one order of magnitude above the SM prediction.

With an expected cross-section of $\sigma_{b\bar{b}} \sim 500\,\mu\text{b}$ and a design luminosity of $\mathcal{L} = 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$, the LHC provides abundant opportunities to study b hadron decays. Both general purpose experiments, CMS [56; 57] and ATLAS [58], as well as the dedicated B-physics experiment LHCb [59; 60] have studied the sensitivity to the decay $B_s^0 \to \mu^+\mu^-$.

[1] The signal window is $5.047\,\text{GeV} < m_{\mu\mu} < 5.622$.

$\mathbf{B_q^0} \to \ell^+\ell^-$ in the Standard Model

At the quark level $b \to s(d)$ transitions can be described by the corresponding low-energy effective Hamiltonian:

$$\mathcal{H}_{eff} = \frac{G_F}{\sqrt{2}} V_{tb} V_{tq}^* \sum_i C_i(\mu) Q_i(\mu), \qquad \text{for } q = s, d. \tag{7.1}$$

where G_F is the Fermi constant, V_{tb} and V_{tq}^* are the corresponding CKM matrix element, and $\mu = \mathcal{O}(m_b)$ denotes the mass scale, that separates the short and long distance contributions to the decay amplitude. The Wilson coefficients $C_i(\mu)$ contain the short distance physics contributions at scales higher than μ. Due to the asymptotic freedom of QCD, they can be calculated using perturbative methods, as long as μ is not too small. The coefficients C_i include contributions from the top quark and from other heavy particles. Therefore the coefficients C_i generally depend on the mass of the top quark m_t and on the masses of new particles in models beyond the SM. The local operators Q_i contain the long distance contributions to the decay amplitude—which generally cannot be calculated perturbatively anymore. Since the non-perturbative methods have their limitations, the largest theoretical uncertainties in the decay amplitudes of weak decays come from the operators Q_i. However, the purely leptonic decays $B_q^0 \to \ell^+\ell^-$ can be calculated very reliably and are among the theoretically cleanest decays in the field of rare B-decays. They can be described by only three operators

$$Q_A = \left(\bar{b}_L \gamma^\nu q_L\right)\left(\bar{\ell}\gamma_\nu\gamma_5\ell\right), \qquad Q_S = m_b\left(\bar{b}_R q_L\right)\left(\bar{\ell}\ell\right), \qquad Q_P = m_b\left(\bar{b}_R q_L\right)\left(\bar{\ell}\gamma_5\ell\right) \tag{7.2}$$

and their coefficients C_A, C_S and C_P. The corresponding low-energy effective Hamiltonian reads as

$$\mathcal{H}_{eff} = \frac{G_F}{\sqrt{2}} \frac{\alpha}{\pi \sin^2\theta_W} V_{tb} V_{tq}^* \left[C_A Q_A + C_S Q_S + C_P Q_P\right] + h.c. \tag{7.3}$$

where α is the fine structure constant given as $\alpha(M_Z) = 1/128$ and θ_W is the Weinberg angle. The amplitudes of $B_q^0 \to \ell^+\ell^-$ decays were first calculated in [61] and have been updated with the next-to-leading order QCD corrections in [62; 63; 64]. In terms of Wilson coefficients the branching fraction of $B_q^0 \to \ell^+\ell^-$ decays can be expressed as shown in equation 7.4.

$$\mathcal{B}(B_q^0 \to \ell^+\ell^-) = \frac{G_F^2 \alpha^2}{64\,\pi^3 \sin^4\theta_W} |V_{tb}^* V_{tq}|^2 \, \tau_{B_q} \, m_{B_q}^3 \, f_{B_q}^2 \sqrt{1 - \frac{4m_l^2}{m_{B_q}^2}}$$

$$\times \left[\left(m_{B_q} C_P - \frac{2m_l}{m_{B_q}} C_A \right)^2 + \left(1 - \frac{4m_l^2}{m_{B_q}^2} \right) m_{B_q}^2 C_S^2 \right] \tag{7.4}$$

where f_{B_q} is the decays constant and τ_{B_q} is the lifetime of the B_q meson. In the SM, the dominant contribution comes from the coefficient C_A, whereas the coefficients C_S and C_P of the scalar and pseudoscalar couplings respectively, are suppressed by m_b^2/M_W^2 and can safely be neglected [65]. This gives the following Standard model predictions [55]:

$$\mathcal{B}(B_s^0 \to \mu^+\mu^-) = (3.86 \pm 0.15) \cdot 10^{-9} \times \frac{\tau_{B_s}}{1.527\,\text{ps}} \left[\frac{|V_{ts}|}{0.0408} \right]^2 \left[\frac{f_{B_s}}{240\,\text{MeV}} \right]^2 \tag{7.5}$$

$$\mathcal{B}(B_d^0 \to \mu^+\mu^-) = (1.06 \pm 0.04) \cdot 10^{-10} \times \frac{\tau_{B_d}}{1.527\,\text{ps}} \left[\frac{|V_{td}|}{0.0082} \right]^2 \left[\frac{f_{B_d}}{200\,\text{MeV}} \right]^2 \tag{7.6}$$

$$\mathcal{B}(B_s^0 \to e^+e^-) = (9.05 \pm 0.34) \cdot 10^{-14} \times \frac{\tau_{B_s}}{1.527\,\text{ps}} \left[\frac{|V_{ts}|}{0.0408} \right]^2 \left[\frac{f_{B_s}}{240\,\text{MeV}} \right]^2 \tag{7.7}$$

$$\mathcal{B}(B_d^0 \to e^+e^-) = (2.49 \pm 0.09) \cdot 10^{-15} \times \frac{\tau_{B_d}}{1.527\,\text{ps}} \left[\frac{|V_{td}|}{0.0082} \right]^2 \left[\frac{f_{B_d}}{200\,\text{MeV}} \right]^2 \tag{7.8}$$

$\mathbf{B_q^0} \to \ell^+\ell^-$ beyond the SM

Since these decays are highly suppressed in the Standard Model and C_A is additionally helicity suppressed by a factor m_ℓ^2/m_B^2, they are potentially sensitive probes of physics with new scalar or pseudoscalar interactions. Most of the weakly coupled extensions of the Standard model contain extra Higgs multiplets. In the *two-Higgs-doublet model of type II* (2HDM), one Higgs doublet H_u only couples to up-type fermions and the other doublet H_d only couples to down-type fermions, which avoids tree-level contributions to FCNC couplings. Both doublets acquire a vacuum expectation value and the ratio of these values is a free parameter known as $\tan\beta = v_u/v_d$. After electroweak symmetry breaking, five physical Higgs bosons remain: two neutral scalar particles h and H, one pseudoscalar particle A and two charged particles $H^\pm$. If $\tan\beta$ is large, the Yukawa coupling to b quarks is of the order of one and the decay amplitude can be enhanced substantially. In the large $\tan\beta$ limit, C_P and C_S will have sizable contributions from charged and neutral Higgs bosons in the box and penguin diagrams, while C_A remains the same. The diagrams have been calculated individually in [66] and the final result depends only on the charged Higgs mass M_{H^+} and $\tan\beta$:

$$C_S = C_P = \frac{m_\ell}{4M_W^2}\tan^2\beta\frac{\ln r}{r-1}, \qquad \text{with } r = \frac{M_{H^+}^2}{\overline{m}_t^2}. \tag{7.9}$$

With the current upper limit on the branching ratio from CDF and the lower bound on M_{H^+} from the branching fraction of the inclusive radiative decay $\overline{B} \to X_s\gamma$ [67; 68], the branching fraction in 2HDM can only be enhanced if $\tan\beta > 60$ (see Figure 7.2).

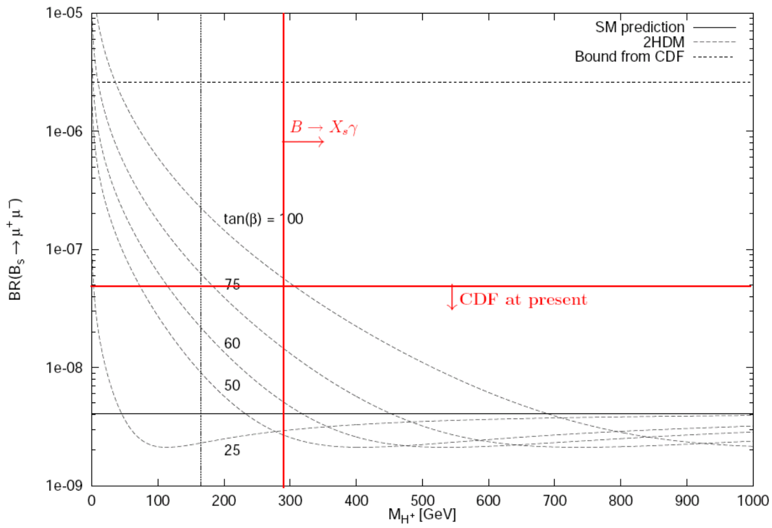

Figure 7.2: $\mathcal{B}(B_s^0 \to \mu^+\mu^-)$ in 2HDM as a function of the charged Higgs mass M_{H^+} for different values of $\tan\beta$ [66] and the updated experimental bounds on the branching fraction from CDF and on M_{H^+} from $\mathcal{B}(\bar{B} \to X_s\gamma)$.

The tree-level Yukawa couplings in the Higgs sector of the *Minimal Supersymmetric Standard Model* (MSSM) are the same as in the 2HDM. At the loop level though both doublets couple to all fermions. In the MSSM, the coefficients depend on the mass $M_A \sim M_H$ of the neutral heavy Higgs bosons and on $\tan^3\beta$, giving the following dependence of the branching fraction on M_A and $\tan\beta$ in the MSSM

$$\mathcal{B}(B_q^0 \to \ell^+\ell^-)_{\text{MSSM}} \propto \frac{m_b^2 m_\ell^2 \tan^6\beta}{M_A^4}. \tag{7.10}$$

The branching fraction of $B_q^0 \to \ell^+\ell^-$ decays can therefore be enhanced by orders of magnitude in the MSSM , especially at large $\tan\beta$. In principle, the MSSM branching fraction could exceed the SM expectation by three orders of magnitude [69; 70], and in supersymmetric models with modified minimal flavour violation at large $\tan\beta$ [71], the branching fraction can be increased by up to four orders of magnitude. Hence, the experimental upper limit from CDF already cuts severely into the MSSM parameter space. In specific models containing leptoquarks [72] and supersymmetric models without R-parity [73] $B_s^0 \to \ell^+\ell^-$ and $B_s^0 \to \ell^+\ell^-$ decays can be enhanced separately even at a low $\tan\beta$.

The strong dependence of $\mathcal{B}(B_s^0 \to \mu^+\mu^-)$ on $\tan^6\beta$ in the MSSM also provides sensitivity to $\tan\beta$. Recently, there has been significant interest [74; 75; 76] in using the decay mode $B_s^0 \to \mu^+\mu^-$ to "measure" the key parameter $\tan\beta$ of the MSSM and to constrain other extensions of the SM. The determination of $\tan\beta$ is difficult—there is no general technique to measure it at hadron colliders—yet all supersymmetric observables, in particular in the MSSM, depend on it. It has been shown, that with very general assumptions, that do not depend on specific models, it is possible to put significant lower bounds on $\tan\beta$. Based on very general principles $\tan\beta$ is also constrained from above [77], so a lower bound on $\tan\beta$ is tantamount to a measurement.

7.1 Event Simulation

The Monte Carlo (MC) event samples were generated, simulated and reconstructed as part of the 'Computing Software and Analysis Challenge 2007' (CSA07) and the predecessor production, named Spring07. The CSA07 event samples were generated using PYTHIA 6.409 [78] and were reconstructed in the CMS software [79] release CMSSW_1_6_X, assuming $100\,\mathrm{pb}^{-1}$ alignment conditions. The Spring07 event samples were generated using PYTHIA 6.227 [80] and were reconstructed with CMSSW_1_3_X. Since the level-1 and high-level trigger information was not available in the samples from the Spring07 production, each sample from Spring07 was reprocessed, executing the various trigger paths with CMSSW_1_3_1_HLT6. Pile-up events were not included in either of the productions. Table 7.2 provides a summary of all the MC event samples used in this analysis. In all event samples, a generator filter required two muons (or hadrons for rare decays), each with transverse momentum $p_\perp > 2.5\,\mathrm{GeV}$ and to be in the central part of the detector $-2.5 < \eta < 2.5$. Details about the different software releases and parametrisations involved in the Spring07 and CSA07 productions are summarised in Table 7.3, along with the production details of the event samples used in the previous $B^0_s \to \mu^+\mu^-$ study in CMS [57].

In PYTHIA there are two ways to generate $b\bar{b}$ events. Using a steering card `MSEL=5`, the $b\bar{b}$ pairs are mainly generated through gluon-gluon fusion and each event contains at least one $b\bar{b}$ pair ($\sigma_{\mathrm{MSEL5}} \approx 500\,\mu\mathrm{b}$). Using the card `MSEL=1` produces the generic QCD $2 \to 2$ subprocesses, which are also referred to as minimum bias events ($\sigma_{\mathrm{MSEL1}} \approx 55\,\mathrm{mb}$). In this study, all signal and background events are selected from `MSEL=1` card and present a mixture of gluon-gluon fusion, flavour excitation, and gluon splitting. The different contributions of these processes to the signal, background and normalisation samples are discussed in more detail in the following subsections.

The event generation through minimum-bias processes is very time-consuming, but necessary for this analysis, as isolation variables have been found crucial for background reduction [53; 54]. It is essential to also include gluon splitting and flavour excitation for $b\bar{b}$ production, when studying the impact of these variables: the two b quarks in gluon-fusion events tend to be back-to-back, while those from gluon-splitting are closer together in phasespace; this has strong influences on the hadronic activity around the dimuon direction.

Table 7.2: Spring07 and CSA07 production event samples used in the analysis. The generated number of events after the generator-level selection described in the text, the equivalent integrated luminosity, the visible cross section, the expected number of events in $1\,\mathrm{fb}^{-1}$, and the branching fraction is given. The visible cross-sections include fragmentation, branching fractions, $p_\perp$ and $|\eta|$ selection criteria. The numbers N_{exp} do not yet include any selection criteria, however, the muon misidentification probability for pions, kaons and protons is already included in N_{exp} for rare decays.

	Sample	N_{gen}	$\mathcal{L}_{\mathrm{gen}}\,[\mathrm{fb}^{-1}]$	$\sigma_{\mathrm{vis}}\,[\mathrm{fb}]$	N_{exp} in $1\,\mathrm{fb}^{-1}$	$\mathcal{B}$	Ref.
Spring07	$B_s^0 \to \mu^+\mu^-$	87041	1.32×10^3	64.0	64.0	3.9×10^{-9}	[55]
Spring07	$b\bar{b} \to \mu^+\mu^- + X$	674727	0.005	1.32×10^8	1.32×10^8		
Spring07	$c\bar{c} \to \mu^+\mu^- + X$	23579	0.002	1.21×10^7	1.21×10^7		
Spring07	$B^\pm \to J/\psi(\to \mu^+\mu^-)K^\pm$	413770	1.97×10^{-1}	2.10×10^6	2.10×10^6	1.0×10^{-3}	[45]
Spring07	$b \to J/\psi(\to \mu^+\mu^-)X$	409574	0.003	1.34×10^8	1.34×10^8		
CSA07	$B_s^0 \to \mu^+\mu^-$	18000	1.75×10^2	102.8	102.8	3.9×10^{-9}	[55]
CSA07	$b\bar{b} \to \mu^+\mu^- + X$	2623900	0.008	3.24×10^8	3.24×10^8		
CSA07	$c\bar{c} \to \mu^+\mu^- + X$	958424	0.010	9.30×10^7	9.30×10^7		
CSA07	*Stew*	12420568					
CSA07	$B_s^0 \to K^+K^-$	7417	0.017	4.29×10^5	52.0	2.4×10^{-5}	[68]
CSA07	$B_s^0 \to \pi^+\pi^-$	8469	1.1	7.70×10^3	0.277	5.0×10^{-7}	[68]
CSA07	$B_s^0 \to K^-\pi^+$	7891	0.115	6.88×10^4	4.5	5.0×10^{-6}	[81]
CSA07	$B_s^0 \to K^-\mu^+\nu_\mu$	5976	0.004	1.62×10^6	1.78×10^4	1.4×10^{-4}	[45]
CSA07	$B_s^0 \to \mu^+\mu^-\gamma$	71500	557	1.28×10^2	128	2.0×10^{-8}	[82]
CSA07	$B^0 \to \pi^+\pi^-$	9538	0.045	2.12×10^5	7.6	5.2×10^{-6}	[68]
CSA07	$B^0 \to \pi^-K^+$	9433	0.013	7.47×10^5	49.3	1.9×10^{-5}	[68]
CSA07	$B^0 \to \pi^-\mu^+\nu_\mu$	8412	0.002	5.08×10^6	3.05×10^4	1.4×10^{-4}	[45]
CSA07	$B^0 \to \pi^0\mu^+\mu^-$	7856	23.8	3.30×10^2	330	2.0×10^{-8}	[82]
CSA07	$\Lambda_b^0 \to \pi^-p^+$	10081	0.308	3.27×10^4	0.393	3.5×10^{-6}	[81]
CSA07	$\Lambda_b^0 \to K^-p^+$	10948	0.197	5.54×10^4	1.2	5.6×10^{-6}	[81]
CSA07	$B^+ \to \mu^+\mu^-\mu^+\nu_\mu$	5508	0.053	1.04×10^5	1.04×10^5	5.0×10^{-6}	[83]
CSA07	$B_c^+ \to \mu^+\mu^-\mu^+\nu_\mu$	9087	36.7	2.48×10^2	248	5.0×10^{-6}	[83]
CSA07	$B_c^+ \to J/\psi(\to \mu^+\mu^-)\mu^+\nu_\mu$	3113	0.080	3.91×10^4	3.91×10^4	2.0×10^{-2}	[84]

Table 7.3: Monte Carlo event sample productions.

	CSA07	Spring07	SM06/private
Generation	PYTHIA 6.409	PYTHIA 6.227	PYTHIA 6.227
Interface	-	-	CMKIN_6_0_0
Simulation	CMSSW_1_4_X	CMSSW_1_2_3(4)	OSCAR_3_6_5
Reconstruction	CMSSW_1_6_X	CMSSW_1_3_X	ORCA_8_7_3
Trigger	CMSSW_1_6_X	CMSSW_1_3_1_HLT6	private code
Alignment conditions	100 pb^{-1}	ideal	ideal
Average pile-up events	-	-	5

7.1.1 Signal

Figure 7.3 illustrates the production mechanisms contributing to the signal sample. In addition to the generator-level requirements described above, the events are required to have a reconstructed dimuon candidate, where the two muons have different electric charge and are reconstructed as global muons (for more details see section 7.3). No other selection requirements are applied. The dominant subprocess is flavour excitation (52 %), followed by gluon splitting (29 %) and gluon-gluon fusion (17 %).

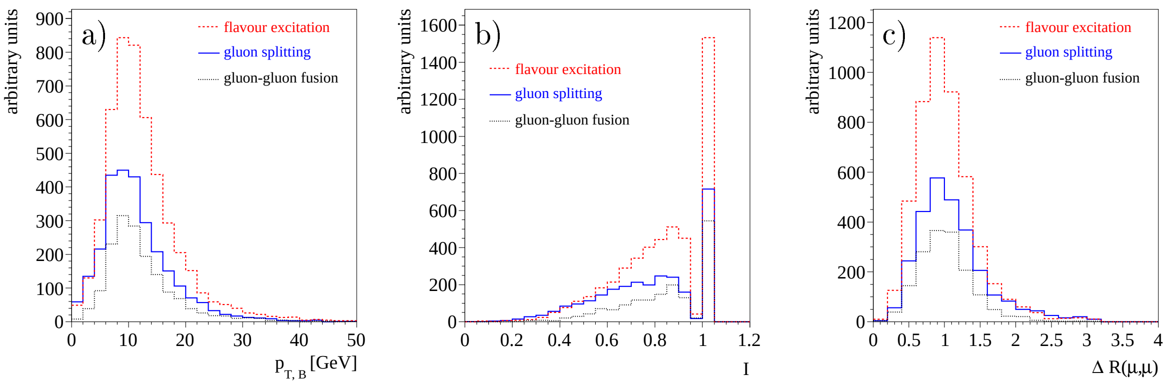

Figure 7.3: Contributions of different partonic processes to the signal sample: flavour excitation (52 %), gluon splitting (29 %) and gluon-gluon fusion (17 %). For the reconstructed muon pairs the graphs show a) transverse momentum, b) isolation variable, c) $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ separation.

The $p_\perp$ spectrum, and probably the isolation distribution, of the B_s-mesons will be reweighed to account for differences between the MC simulation and data. These weights will be obtained from the comparison of the corresponding spectra of $B^\pm \to J/\psi K^\pm$ in MC simulation and data.

In the PYTHIA signal sample, both B_s^0 and $\bar{B}_s^0$ are forced to decay into a muon pair. Therefore, events containing two B_s^0 mesons will contain two signal decays. These events have been artificially removed from the analysis, since the leptonic decay of the second B_s^0 meson biases the hadronic activity in the event. To correct for this removal, the number of signal events in Table 7.2 has been increased by 7.6 % (CSA07) and 11 % (Spring07). These numbers were determined on the signal MC samples and constitute the fraction of events with two B_s^0 mesons compared to events with exactly one B_s^0 meson. While this fraction depends on f_s, the generator-level filters also strongly affect this number.

7.1.2 Background

The main challenge in the measurement of the $B_s^0 \to \mu^+\mu^-$ decay rate is background suppression. Many background sources can mimic the signal topology. First, $q\bar{q}$ events (where $q = b, c$) with $q \to \mu X$ (prompt or cascade) decays of both q-hadrons or a single $q \to \mu X$ decay combined with misidentified muon (punch-through or in-flight decay of a hadron). Second, events where a true muon is combined with a misidentified hadron. Since the available MC event samples from Spring07 and CSA07 do not contain an adequate simulation of a substantial background source, this type of background has not been studied in the scope of this work. An estimation of the contribution based on a generator-level simulation can be found in [85]. Finally, rare B_d, B_u, B_s and Λ_b decays, comprising of hadronic, semileptonic, and radiative decays. Some of these decays constitute a resonant background, such as $B_s \to K^+K^-$, $\Lambda_b \to pK^-$, others have a continuum dimuon invariant mass distribution.

7.1.2.1 Muon-enriched QCD Background ('Stew')

A possibility for studying the generic QCD background in the CSA07 production is potentially provided by the 'Stew'—the 'soup' containing muon-enriched samples of onia, non-prompt J/ψ, and minimum bias events (ppMuX) [86]. However, the equivalent luminosity of this particular background in this 'soup' is only of the order $0.4\,\mathrm{pb}^{-1}$, thus rendering the 'Stew' unusable.

7.1.2.2 Background from Semileptonic Heavy Quark Decays

In the non-peaking dimuon background samples $b\bar{b} \to \mu^+\mu^- + X$ and $c\bar{c} \to \mu^+\mu^- + X$, both heavy quarks are forced to decay in a multitude of semimuonic decay channels (more details can be found in the configuration files in the CMSSW CVS repository [87]). For $b\bar{b} \to \mu^+\mu^- + X$, no constraints on the decay of the charm meson are applied, and therefore events where one b hadron decays into two muons (one from the direct $b \to c\mu^-\bar{\nu}$ and from $b \to c \to s\mu^+\nu$) are included as well. Semimuonic charm decays after hadronic B decays are not contained in this event sample, as all B mesons are forced to decay semimuonically. Of the remaining background events after the full analysis, the background is composed entirely of muons from direct B decays (see section 7.4).

The production mechanisms of the background $b\bar{b} \to \mu^+\mu^- + X$ are illustrated in Figure 7.4. In addition to the generator-level requirements described above, the events are required to have a reconstructed dimuon candidate, where the two muons have different electric charge and are reconstructed as global muons (for more details see section 7.3). The dominant subprocesses are gluon splitting (45 %) and flavour excitation (38 %), followed gluon-gluon fusion (15 %).

7.1.2.3 Rare b-Hadron Decays Background

Rare b-hadron decays could potentially lead to sizable background contributions. The following two cases can be distinguished:

- Peaking background from rare decays, where a heavy particle decays into a pair of hadrons. Examples for these decays include $B_s \to K^+K^-$, $\Lambda_b \to pK^-$.
- Non-peaking background from rare B_d, B_u, and B_s decays, comprising hadronic, semileptonic, and radiative decays. The invariant mass distribution for these decays is a continuum with an upper edge at the mass of the decaying particle; the finite momentum resolution could lead to events reconstructed in the $B_s^0 \to \mu^+\mu^-$ signal mass window. Because semileptonic decays have branching fractions several orders of magnitude above $\mathcal{B}(B_s^0 \to \mu^+\mu^-)$, this background could be problematic.

For each decay channel, events were generated and analyzed without requiring explicit muon identification. The misidentification probability was applied as weighting factors at the end.

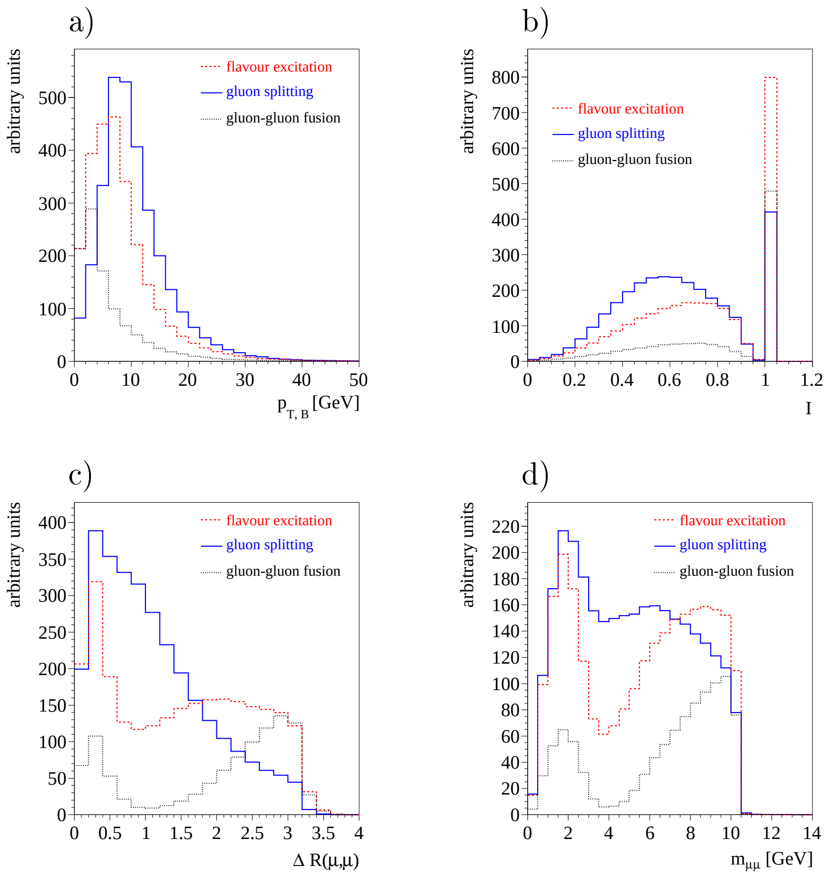

Figure 7.4: Contribution of different partonic processes to the background sample $b\bar{b} \to \mu^+\mu^- + X$: gluon splitting (45 %) flavour excitation (38 %), and gluon-gluon fusion (15 %). For the reconstructed muon pairs the graphs show a) transverse momentum, b) isolation variable, c)$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ separation, d) invariant mass.

7.1.3 Normalisation

To minimise the dependence on the unknown $b\bar{b}$ production cross section and luminosity measurements, a relative normalisation to the well-measured decays $B^\pm \to J/\psi K^\pm$ is used in this analysis. Choosing a decay channel with a signature similar to the signal decay $B_s^0 \to \mu^+\mu^-$, such as $B^\pm \to J/\psi K^\pm$, has the advantage, that many systematic errors cancel to first order, when deriving the upper limit normalising to a similar decay channel measured in data. The upper limit on the branching fraction is (schematically) determined by

$$\mathcal{B}(B_s^0 \to \mu^+\mu^-;\ 90\,\%\ \text{C.L.}) = \frac{N(B_s^0 \to \mu^+\mu^-;\ 90\,\%\ \text{C.L.})/\varepsilon_{B_s}}{N(B^\pm \to J/\psi K^\pm)/[\varepsilon_{B^+} \cdot \mathcal{B}(B^\pm \to J/\psi K^\pm) \cdot \mathcal{B}(J/\psi \to \mu^+\mu^-)]} \times \frac{f_u}{f_s}, \quad (7.11)$$

where ε_{B_s} and ε_{B^+} are the combined acceptance, trigger, and selection efficiencies for the signal and normalisation samples respectively. $N(B_s^0 \to \mu^+\mu^-;\ 90\,\% C.L.)$ is the expected 90 % C.L. upper limit on the number of signal decays and $N(B^\pm \to J/\psi K^\pm)$ is the number of reconstructed $B^\pm \to J/\psi K^\pm$ candidates. f_u and f_s describe the probability that a b quark hadronises into a B^+ or B_s^0 meson.

Figure 7.5 illustrates the production mechanisms contributing to the signal sample of the normalisation channel. In addition to the generator-level requirements described above, the events are required to have a B^+ candidate reconstructed from two muon tracks, where the two muons have different electric charge, and a third track selected from a cone around the dimuon direction. No other selection requirements are applied. The dominant subprocess is flavour excitation (53 %), followed by gluon splitting (28 %) and gluon-gluon fusion (18 %).

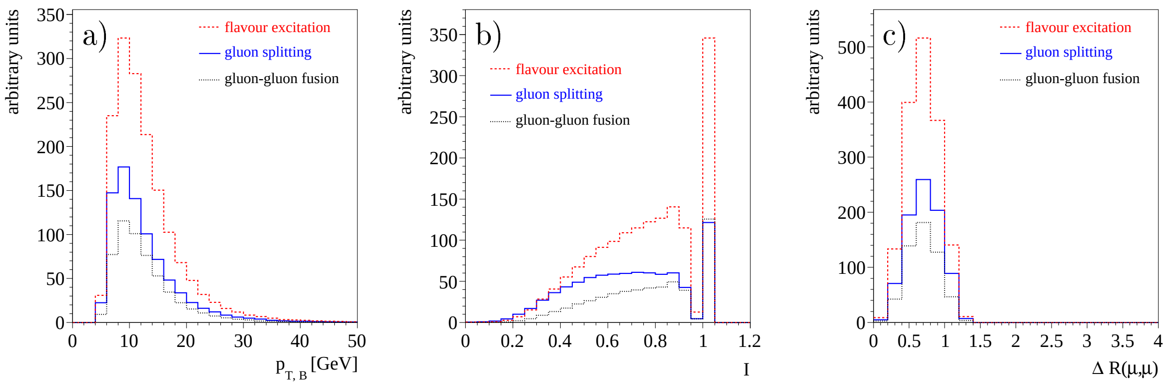

Figure 7.5: Contribution of different partonic processes to the signal sample of the normalisation channel: flavour excitation (53 %), gluon splitting (28 %) and gluon-gluon fusion (18 %). For the reconstructed muon pairs the graphs show a) transverse momentum, b) isolation variable, c) $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ separation.

7.2 Trigger

This analysis is not primarily targeted at the initial very low-luminosity start-up period of the LHC but requires about $1\,\mathrm{fb}^{-1}$. Therefore the trigger strategy is based on an instantaneous luminosity of at least $10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ as provided in [88].

7.2.1 Level-1 Trigger

The level-1 (L1) muon trigger provides fast identification of muon candidates and an estimate of their transverse momentum $p_\perp$ based on signals from the drifttubes (DT, $|\eta| < 1.2$), cathode strip chambers (CSC, $0.9 < |\eta| < 2.4$) and resistive plate chambers (RPC, $|\eta| < 2.1$, in the start-up phase $|\eta| < 1.6$). The DT and CSC subtriggers compare segment slopes in successive layers for their $p_\perp$ estimate, while the RPC subtrigger is based on predefined hit patterns to classify the muon trajectory. The Global Muon Trigger matches the DT and CSC candidates with RPC candidates and rejects unconfirmed candidates. The four highest-quality muon candidates with the largest $p_\perp$ are passed on to the global trigger, which sorts them by rank. The rank is determined by $p_\perp$ and quality. In the global trigger, separate threshold requirements could be applied to each of the four muon candidates; other requirements for the azimuthal angle or pseudorapidity of single muon candidates are also possible.

In this analysis, the L1 condition is based on `A_DoubleMu3` requiring two L1 muons anywhere in the muon detector, each with $p_\perp > 3\,\mathrm{GeV}$. No isolation or charge requirements are applied.

7.2.2 High Level Trigger

The high-level trigger (HLT) condition is based on the displaced dimuon trigger `BJpsiMuMu` described in detail in section 7.5.1 of [88]. The HLT starts with the level-2 (L2) muon reconstruction. L1 muon candidates serve as seeds for the reconstruction of (standalone) tracks in the muon chambers with higher $p_\perp$ resolution compared to L1. A transverse momentum requirement of $p_\perp > 3\,\mathrm{GeV}$ is applied to these L2 muons. In the next step, L2 muons are used to determine regions of interest where tracks in the central tracker are reconstructed and combined with the L2 muons. This constitutes slightly different muons than the standard level-3 (L3) muons. The combined muon track has to satisfy $p_\perp > 3\,\mathrm{GeV}$. The two muons are fit to a common decay vertex where a good vertex quality is required with $\chi^2 < 10$. The significance of the transverse decay length is required to be above 3 and the angle α between the reconstructed dimuon momentum vector and the vector from the primary to the decay vertex has to fulfill $\cos\alpha > 0.9$. The primary vertex at the HLT is determined with pixel tracks using the divisive method [89].

7.2.3 Determination of Trigger Efficiency

The determination of the trigger efficiencies in data comprises several components: (i) the single muon efficiency at level-1, (ii) the single-muon efficiency at HLT (either L3 or an independent version as implemented in the $b \to J/\psi \to \mu^+\mu^-$ HLT trigger path), (iii) and finally the selection efficiency of additional criteria applied at the HLT. The following describes the 'tag and probe' method, where one well-identified ('tag') muon is used to seed the reconstruction of a J/ψ candidate, which serves as a source of unbiased ('probe') muons.

To determine the L1 single muon efficiency in data, an *unbiased* muon sample must be available. The decay $J/\psi \to \mu^+\mu^-$ provides this possibility. The event sample is triggered by single relaxed muons, passed through HLT with prescales ranging from 1–4000 with overall event rates of < 1 Hz as described in [88]. A single well-identified muon of specific charge, matched to the L1-trigger primitive, is combined with other tracks of $p_\perp > 2\,\mathrm{GeV}$ to form J/ψ candidates. A fit to the invariant mass distribution with a Gaussian and polynomial provides an estimate for the total number N_{tot} of J/ψ candidates. This yield can be compared to the number N_{L1L1} of J/ψ candidates where the second track is matched to a L1-trigger primitive. The muon L1 trigger efficiency follows as $\varepsilon_\mu = N_{L1L1}/N_{tot}$.

The HLT muon efficiency is determined in a similar way. It remains to be seen which prescaled sample provides the best statistical sensitivity: higher-$p_\perp$ single muons with lower prescale factors, or lower-$p_\perp$ single muons with higher prescale factors. This will not be a problem given the very open triggers during the start-up phase.

The HLT efficiency for additional selection criteria can be determined for the normalisation sample $B^\pm \to J/\psi K^\pm$ in data and MC simulation. The comparison of these efficiencies will provide an estimate of the systematic error to be applied for the HLT efficiency for the signal $B^0_s \to \mu^+\mu^-$. The best sample for this study is a prescaled L1-dimuon sample.

7.3 Muon Reconstruction

7.3.1 Muon Reconstruction

The track parameters of the muons are measured in two CMS sub-detectors: the inner tracker and the muon system. Independent of the subsystem, the trajectories are reconstructed using the same track parametrisation and the same tracking algorithm as in section 2.7. Depending on the sub-system involved in the reconstruction of the high-level muon physics object, there are three different types of muons [34]: stand-alone, global and tracker muons.

7.3.1.1 Stand-alone Muons

The stand-alone reconstruction uses only the hits in the muon spectrometer. Seeds are generated based on DT and CSC. The seed is propagated to the innermost compatible layer in the muon system. A pre-filter is applied in the inside-out direction, using the track segments provided by the DT and CSC for the fit and imposing only a loose χ^2 cut. In the final filter the trajectory is built in the outside-in direction, using the hits composing the track segment with a tighter χ^2 cut. At each filter step the trajectory parameters are propagated from one layer in the muon system to the next, including multiple scattering and energy losses due to ionisation and bremsstrahlung in the return yoke and the muon chamber. A trajectory is only accepted as a muon track if there are at least two measurements present in the fit, where one of them has to be DT or CSC type. The inclusion of the RPC measurements can improve the reconstruction efficiency of low momentum muons. After the trajectory cleaning, the remaining tracks are extrapolated to the point of closest approach to the beam line and a beam spot constraint is applied to improve the $p_\perp$ resolution.

7.3.1.2 Global Muons

Global muons are reconstructed by combining tracks reconstructed in the tracker system (‘tracker tracks’) and tracks reconstructed in the muon system (‘muon tracks’). Since the momentum resolution of muon tracks with $p_T < 200\,\mathrm{GeV}$ is dominated by multiple scattering, the resolution at low momentum is significantly improved by including the information from the tracker. The track reconstruction in the tracker starts with the seed generation. A track seed can be defined by a hit pair or a hit triplet in the pixel. Since a hit pair does not constrain the momentum, an additional vertex constraint is applied. Seeds from hit pairs can have a high ghost rate whereas seeds from hit triplets have a high purity but a significantly lower efficiency. Therefore, in the standard track reconstruction only seeds from hit pairs are used. The track candidates from triplets on the other hand, allow a simple and efficient primary vertex reconstruction and can be used in the online selection. The pattern recognition is based on the combinatorial Kalman filter method and proceeds as described in section 2.7. To account for the possibility that a track did not leave a hit in a specific layer, an additional trajectory without an associated hit ('invalid hit') is created in each layer. To limit the otherwise exponentially growing number of candidates, the number of candidates is truncated at each layer by limiting the maximum number of candidates, the minimum number of hits per track, the number of invalid hits, the maximum χ^2 and the minimum transverse momentum. After track building, ambiguities in trajectories sharing more than 50 % of their hits[1] are resolved by discarding the track with less hits or, in case of equal numbers of hits, the track with the higher χ^2.

The track matching between tracker tracks and muon tracks proceeds in two steps. In the first step, a region of interest in the $\eta - \phi$ space is defined: the origin of this region is defined by the primary vertex from the pixel algorithm. The direction around which the region of interest will be opened is taken from the stand-alone muon. The sizes $\Delta\eta$ and $\Delta\phi$ of the region of interest are determined from the error estimates of the stand-alone muon direction, where the values of $\Delta\eta$ and $\Delta\phi$ are limited to keep the region of interest of reasonable size [2]. Only tracks that are within the region of interest and have a $p_\perp$ above 60 % of the $p_\perp$ of the stand-alone muon track are selected. In the second step, the subset of selected tracker tracks are matched to the muon tracks by comparing the five parameters describing the trajectories. The trajectories of either tracker track or the muon track are propagated onto a common surface. For low p_T muons this is the detector surface of outermost tracker track hit and for high p_T muons it is the detector surface of the innermost muon track hit. The best match is chosen by applying more stringent momentum and spatial matching criteria on a combination of discriminating variables, that are determined by the position and momentum of the two tracks.

Finally a global refit is performed for each combination of a tracker muon and a stand-alone muon by combining the corresponding collections of tracker and muon

[1] relative to the number of hits in the trajectories with the least number of hits

[2] the size of the region of interest has a strong impact on the reconstruction efficiency and the fake rate

hits.[1] Since only the global muon track with the best χ^2 is kept, in case there is more than one possible global muon track, there is a maximum of one global muon reconstructed for each stand-alone muon.

As shown in Figure 2.14, the tracker system is essential to ensure a good momentum resolution at low transverse momentum, where the resolution in the muon chambers is dominated by multiple scattering. At high transverse momentum the best momentum resolution is given by the resolution obtained with the muon system.

7.3.1.3 Tracker Muons

Stand-alone muon reconstruction only becomes highly efficient for muons with a $p_\perp$ of more than $6 - 7\,\mathrm{GeV}$. Muons with a lower $p_\perp$ do not leave enough hits in the muon spectrometer to be reconstructed as stand-alone muons, or do not reach the muon system at all (see Table 7.4).

Table 7.4: The required minimum $p_\perp$ for a muon to reach the first muon station in different η regions assuming a homogeneous magnetic field of 4 T, where R_T^{min} is the minimal radial distance to the first muon chamber in the corresponding η region.

	R_T^{min}	$p_T^{min} = 0.3 B R_T^{min}$
$0 < \|\eta\| < 1.2$	4 m	4.8 GeV
$1.2 < \|\eta\| < 1.5$	3 m	3.6 GeV
$1.5 < \|\eta\| < 2.4$	1 m	1.2 GeV

The complementary approach of tracker muons is therefore particularly useful in the reconstruction of low $p_\perp$ muons. The reconstruction of tracker muons considers all tracker tracks and searches for compatible segments in the muon system. In the first step, each track is propagated in the calorimeter and the energy deposited in ECAL crystals and HCAL towers are calculated. In the second step, the track is extrapolated into the muon detectors. In both steps the magnetic field inhomogeneities, multiple scattering and energy losses are taken into account. While the trajectory is propagated through the muon system, the algorithm collects and stores all relevant information.

[1] The resolution of high energy muons can be improved by omitting selected hits in the muon system where the measurements can be degraded by electromagnetic showers.

Based on this information the tracker tracks can be matched to hits in the muon segments. By design, the association between tracker tracks and the muon segments is kept very loose. Unlike for global muons, no combined refit is performed for the tracker muons. In case several tracks that are close to each other have been associated to the same segment, the best track-segment combination is determined using arbitration algorithm [34]. In the default configuration the minimum $p_{\perp}$ threshold is 1.5 GeV for tracker muons and the minimum number of matched segment is one.

The reconstruction efficiencies of tracker, stand-alone and global muons are illustrated in Figure 7.6 for different $p_{\perp}$ samples as a function of pseudorapidity [34]. The drops in efficiency correspond to discontinuities in the geometrical structure of the CMS detector:

- $|\eta| \simeq 0$: gaps between the barrel pixel sensors on the ladders at $z = 0$
- $|\eta| \simeq 0.3$: discontinuity between the DT central wheel and its neighbours
- $0.8 < |\eta| < 1.2$: overlap between DT and CSC (leading to failures in the seed finding algorithm)
- $|\eta| \simeq 1.8$: transition from the TID to the TID/TEC subsystem

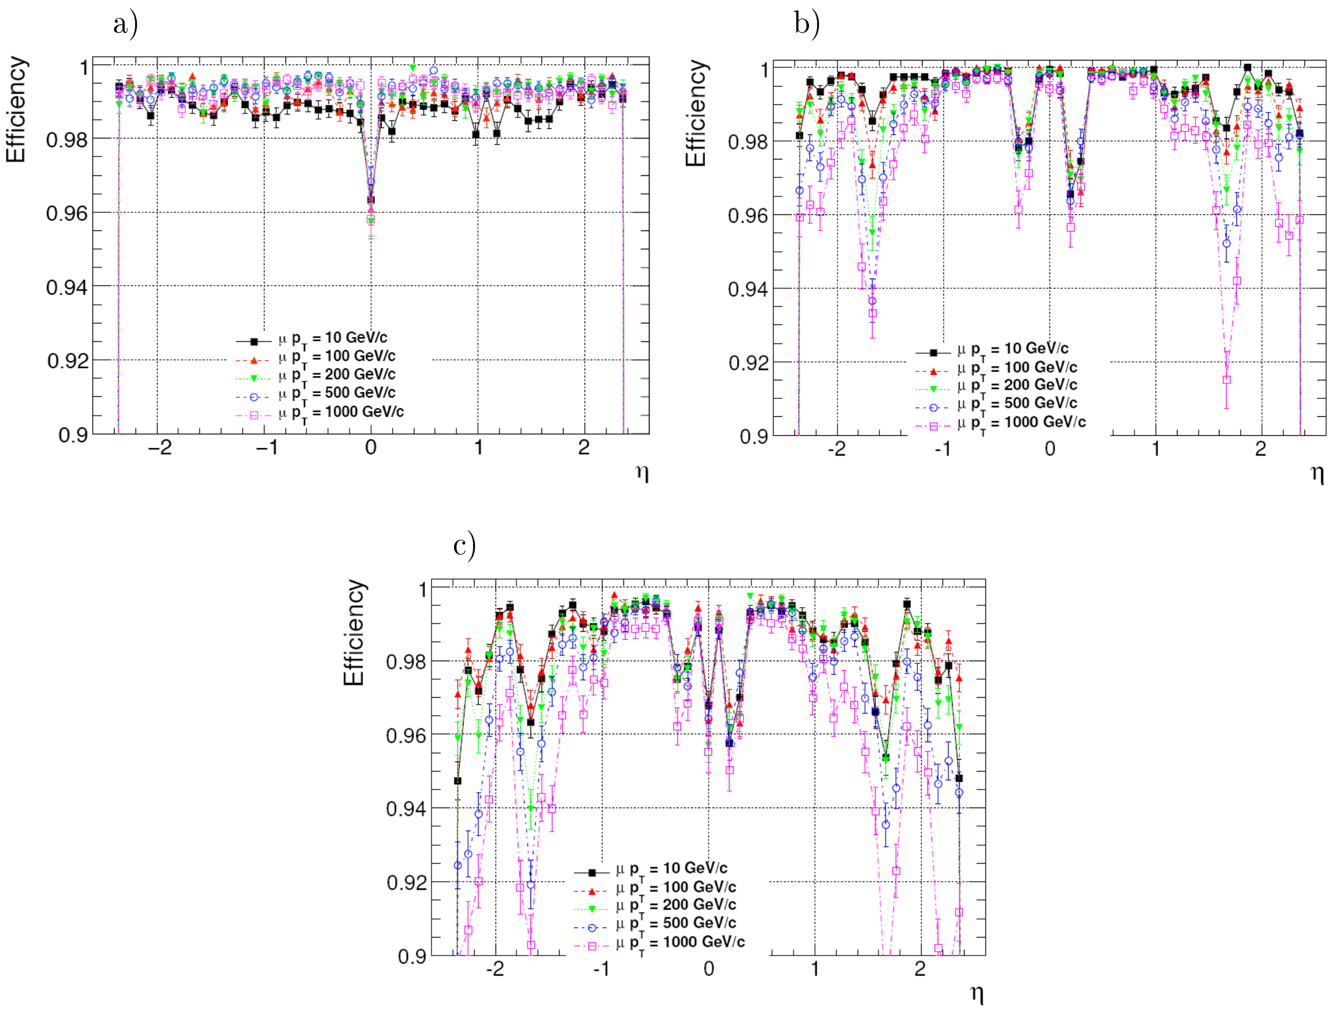

Figure 7.6: Reconstruction efficiencies of a) tracker, b) stand-alone and c) global muons for different $p_\perp$ samples as a function of pseudorapidity [34].

7.3.2 Muon Identification

In this analysis, muon candidates are selected from global muons. If less than two global muons are found in an event, additional muon candidates are added from tracker muons, if available. Figures 7.7 and 7.8 illustrate the muon identification efficiencies for tracker and global muons determined on the CSA07 MC event samples used in this analysis (see Table 7.2).

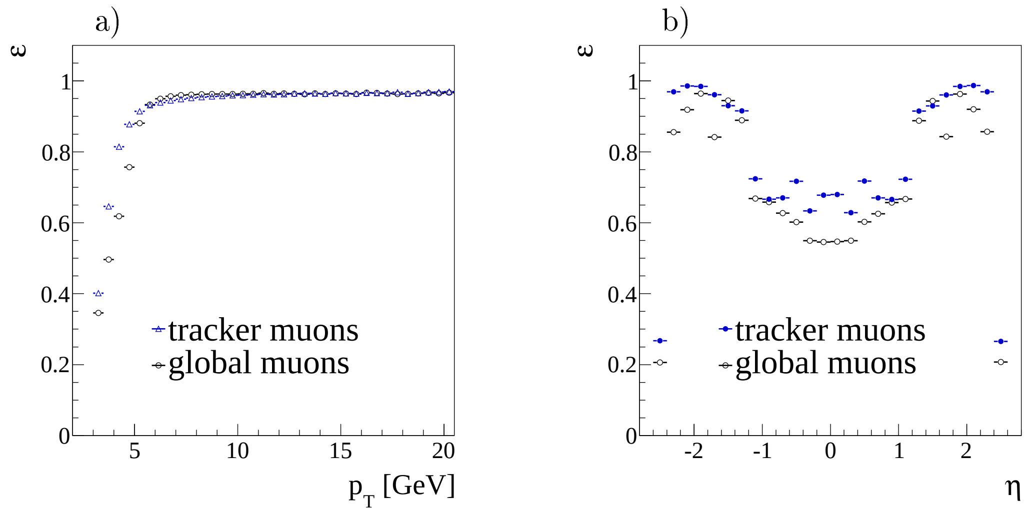

Figure 7.7: Muon identification efficiency for global muons and tracker muons from CSA07 samples as a function of a) transverse momentum $p_\perp$, b) pseudorapidity η.

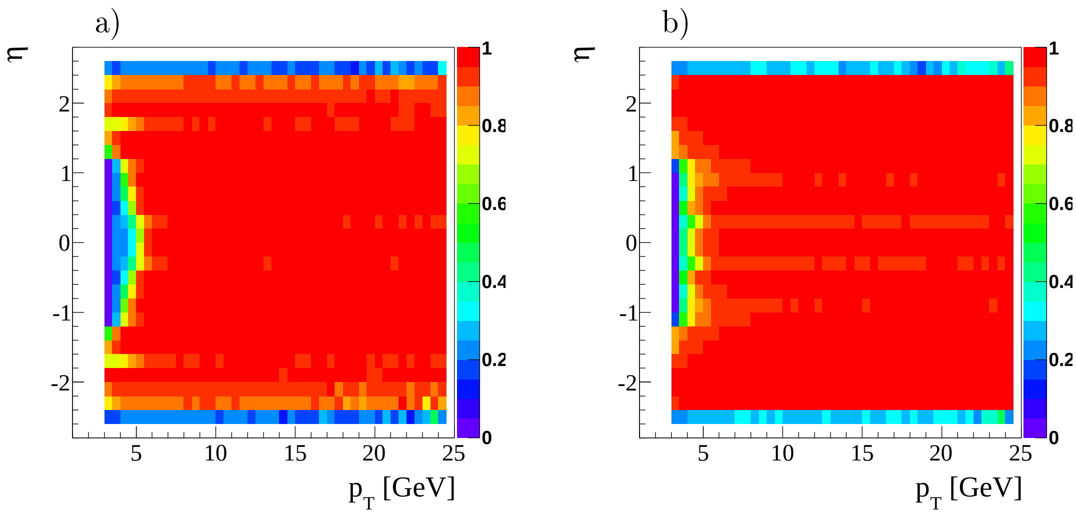

Figure 7.8: Muon identification efficiency from CSA07 samples as a function of pseudo-rapidity η as a function of the transverse momentum $p_\perp$ for a) global muons, b) tracker muons.

7.3.3 Muon Misidentification

Hadrons can be misidentified as muons primarily for two reasons:

- **Punch-through hadrons:** High-momentum hadrons can traverse the calorimeters without hadronic interaction (with a probability $p = \exp(-x/\lambda)$, where x is the distance travelled and λ is the hadronic interaction length) and then interact in the muon system, thus faking a muon signature.
- **In-flight decays of hadrons:** Hadrons, in particular charged kaons, decay dominantly into muons, which will be measured in the muon system.

In the following, the contribution from both effects are included in the misidentification rates, and will not be treated separately. The probability of hadron misidentification is momentum dependent and illustrated in Figure 7.9. All CSA07 event samples have been used to determine whether particles produced as hadrons close to the interaction region have been identified as muons, using the full simulation and reconstruction chain as described in section 7.1. From these figures (conservative) average misidentification probabilities have been extracted for the three charged hadron species $\varepsilon_\pi = 0.6\%$, $\varepsilon_K = 1.1\%$, $\varepsilon_p = 0.2\%$. The misidentification probabilities are used as scaling weights for the rare background contributions, which are dominated by hadrons that have been misidentified (see section 7.1.2.3).

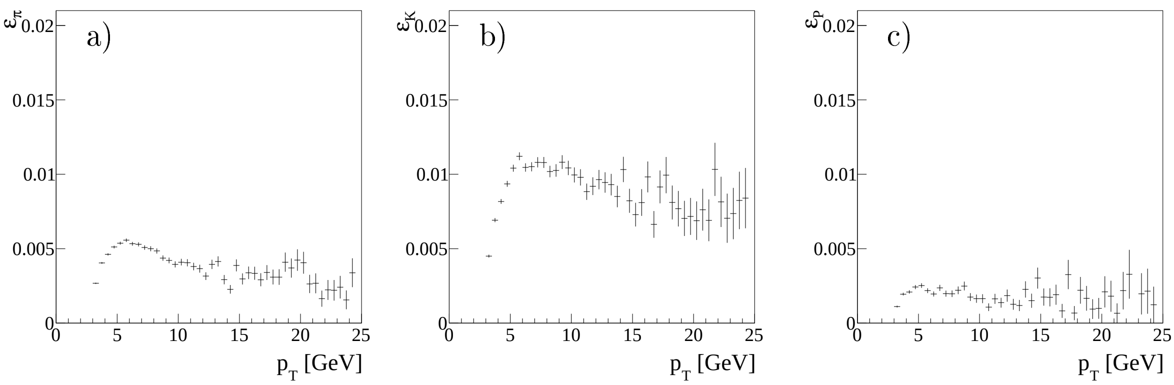

Figure 7.9: Muon misidentification rate for different hadrons as a function of transverse momentum: a) pions, b) kaons, c) protons. Contributions from hadron punch-through and from in-flight decays are included in these illustrations. From these plots, average misidentification rates are determined as follows: $\varepsilon_\pi = 0.6\%$, $\varepsilon_K = 1.1\%$, $\varepsilon_p = 0.2\%$

The muon misidentification probabilities are needed to estimate the contributions in the signal region from peaking rare backgrounds only (for example $B^0 \to \pi^+\pi^-$). Non-peaking backgrounds will be estimated from the sidebands. The kaon and pion misidentification probabilities can also be determined in data, e.g. with $D^0 \to K\pi$ samples obtained in partially reconstructed semileptonic B-decays [90].

7.3.4 Muon Identification Efficiency Determination

The muon identification efficiency is determined with the 'tag and probe' (TNP) method [91], also used for the determination of the trigger efficiency (described in section 7.2.3).

Well-identified global muons μ of a specific charge, matched to the relaxed single muon trigger primitives at both L1 and HLT, are used to seed the reconstruction of J/ψ candidates. Tracks t of the opposite charge and $p_\perp > 2\,\text{GeV}$ within $\Delta R < 1.5$ are combined with the muon and retained if the invariant mass is between $2.5 < m_{\mu t} < 3.5\,\text{GeV}$, illustrated in Figure 7.10a). The J/ψ candidate mass distribution formed by two identified muons is shown in Figure 7.10b). The efficiency can be determined in two ways, which give consistent results and provide a systematic cross-check:

$$\varepsilon \;=\; \frac{N_{\mu\mu}}{N_{\mu t}} = \frac{N_{\mu\mu}}{N_{\mu\mu} + N_{\mu\bar{\mu}}},$$

where $N_{\mu\mu}$ ($N_{\mu t}$) is the number of J/ψ mesons extracted from a fit to data with both the tag and probe leptons (only the tag lepton) identified as muon. In the second approach, $N_{\mu\bar{\mu}}$ quantifies the number of J/ψ mesons, again extracted from a fit to the data, where the probe explicitly failed muon identification. In the above equations, all yields are evaluated as integrals of single Gaussians above a linear background, in intervals of transverse momentum and pseudorapidity. In principle $N_{\mu t} = N_{\mu\mu} + N_{\mu\bar{\mu}}$, but the practical determination of the numbers differs as they are extracted from fits to different histograms.

Figure 7.11 illustrates the muon identification efficiency as a function of transverse momentum in three pseudorapidity bins. The TNP method is compared against two MC-truth based methods. For truth-matching, the standard CMSSW algorithm 'TrackAssociatorByChi2' has been used to match the inner-tracker track of the global muon to a generator-level charged particle. The histogram labelled 'MC' is the muon identification efficiency determined on all muons using MC-truth to identify any muon. This histogram provides a cross-check that the TNP-selected muons do not induce a

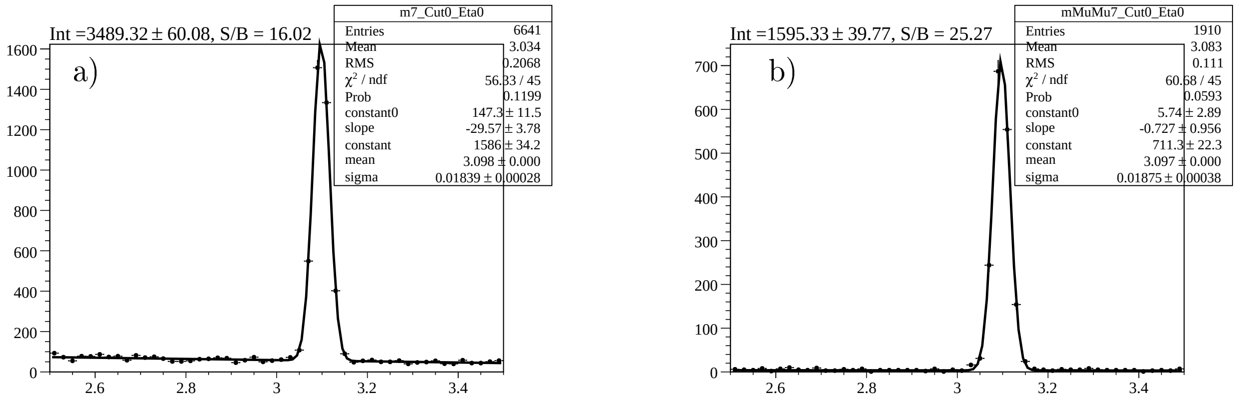

Figure 7.10: Reconstructed J/ψ candidates in an event sample containing a luminosity-weighted combination of non-prompt J/ψ and muons from B-decays [91]. a) J/ψ candidates formed from one muon candidate and one track, b) J/ψ candidates formed from two muon candidates.

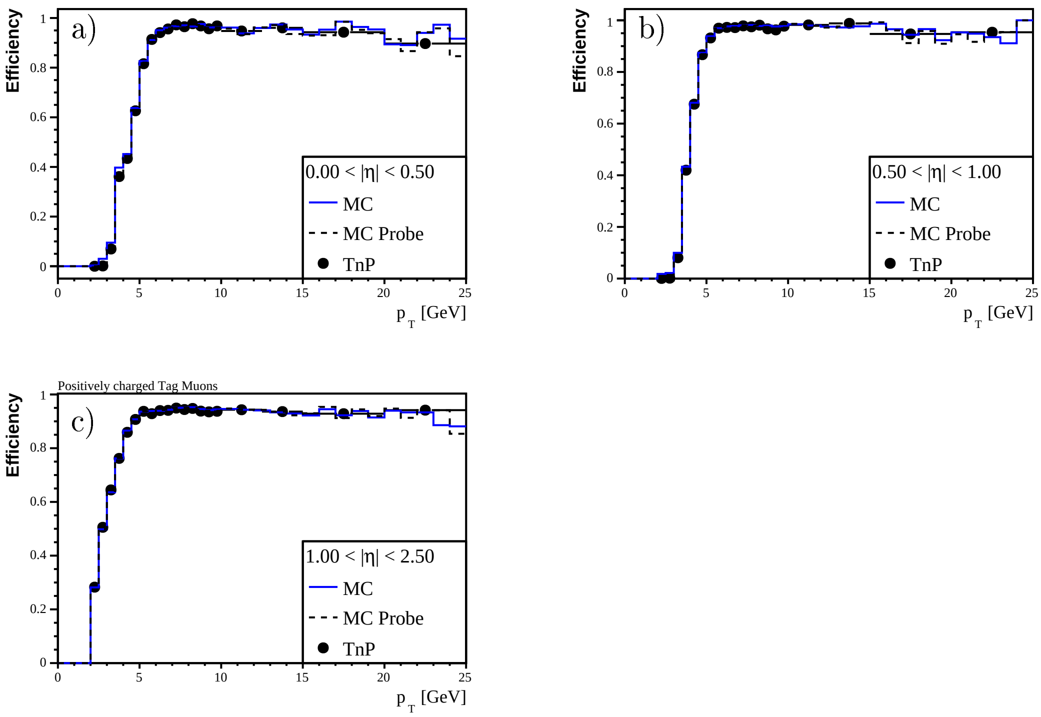

Figure 7.11: Muon identification in three $|\eta|$ bins: a) $|\eta| < 0.5$, b) $0.5 < |\eta| < 1$ and c) $1 < |\eta| < 2.5$, measured with the 'tag and probe' method [91] (labelled 'TnP'). The histogram labelled 'MC' is the muon identification efficiency determined on all muons using MC-truth to identify any muon. The histogram labelled 'MC Probe' uses MC-truth on the restricted set of muons which accompany a tag muon.

bias in the muon selection. The histogram labelled 'MC Probe' uses MC-truth on the restricted set of muons which accompany a tag muon and establishes that the yield determination from the fit is unbiased.

7.4 Event Selection for $B_s^0 \to \mu^+\mu^-$

7.4.1 Selection Variables

For the offline event selection, variables related to the primary vertex, the muon candidates, and the B_s^0 candidate with its associated secondary vertex are calculated. In the following a description of the calculation of all relevant variables is provided. Tables 7.5 and 7.6 summarise the numerical values for all selection criteria applied on these variables for signal and various background samples. For the figures illustrating the distributions used in the analysis, all previous selection requirements have been applied. Appendix C provides more illustrations where the distributions are shown after the HLT. In all figures of this section and the Appendix, the background is composed of $b\bar{b} \to \mu^+\mu^- + X$. The background contributions from rare decays (peaking and non-peaking) will be discussed later in section 7.6 and are not included in this section. The most important selection criteria have been optimised in a grid search for best upper limit. This is described in subsection 7.4.4. The primary vertex is determined with the standard algorithm [92] used in CMS.

7.4.1.1 Muon Selection

Muon candidates are selected from the global muon collection. If more than two muon candidates are found, the pair with the smallest hf separation is chosen. Alternative selection schemes, e.g. the two leading muons, or the leading muon plus the closest muon, lead to comparable signal selection efficiencies, albeit with (insignificantly) lower signal/background ratios. Both muons are required to have transverse momentum $p_\perp > 4.0\,\mathrm{GeV}$ and to be in the central part of the detector $-2.4 < \eta < 2.4$. For the signal reconstruction, both muons are required to have opposite charges. The $\eta\phi$ separation of the two muons:

$$\Delta R(\mu\mu) = \sqrt{(\eta_{\mu_1} - \eta_{\mu_2})^2 + (\phi_{\mu_1} - \phi_{\mu_2})^2} \tag{7.12}$$

is a powerful discriminator against gluon-gluon fusion background with both b-hadrons decaying semileptonically: the muons of those b-hadrons tend to be back-to-back, while the signal shows a peaked distribution with a maximum at $\Delta R(\mu\mu) \sim 1$. Figure 7.12 illustrates signal and background distributions of muon variables.

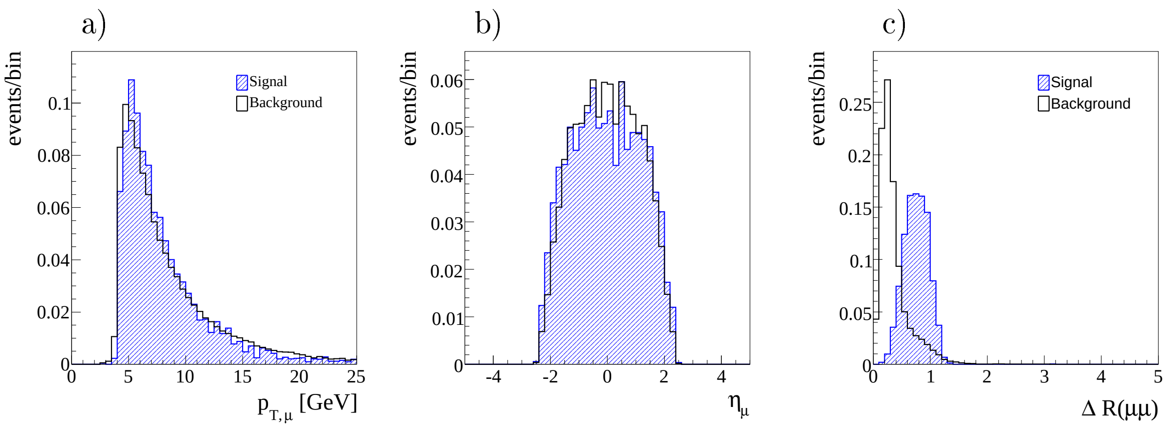

Figure 7.12: Muon variable distributions in the mass region $4.8 < m < 6.0\,\text{GeV}$ (after HLT): a) transverse momentum, b) pseudo-rapidity, c) $\eta\phi$ separation of the two muons. The histograms are normalised to unity.

7.4.1.2 B_s^0 Candidate Selection

B_s candidates are formed by vertexing the two muon candidates. The B_s candidate is required to fulfill $p_\perp > 5\,\text{GeV}$. Figures 7.15a) and 7.15b) show the transverse momentum and pseudorapidity distribution of the reconstructed B_s^0 candidates respectively. The reconstructed mass of the B_s candidate is a powerful handle to reduce backgrounds. Figure 7.13 illustrates the mass resolution obtained on the signal MC event sample at various stages of the analysis. The distribution is fit with two Gaussians, the quoted width $\sigma = 41.7\,\text{MeV}$ is determined according to

$$\sigma^2 = \frac{N_n^2\sigma_n^2 + N_w^2\sigma_w^2}{N_n^2 + N_w^2}, \tag{7.13}$$

where $\sigma_n = 35.5\,\text{MeV}$ ($\sigma_w = 70.2\,\text{MeV}$) and $N_n = 0.17$ ($N_w = 0.06$) are the width and normalisation of the narrow (wide) Gaussian respectively. The mass resolution, in particular its strong $|\eta|$-dependence (see Figure 7.14), is limited by an inconsistent treatment in simulation and reconstruction of inhomogeneities in the magnetic field.[1]

[1] This problem has been resolved in the releases CMSSW_1_8_0.

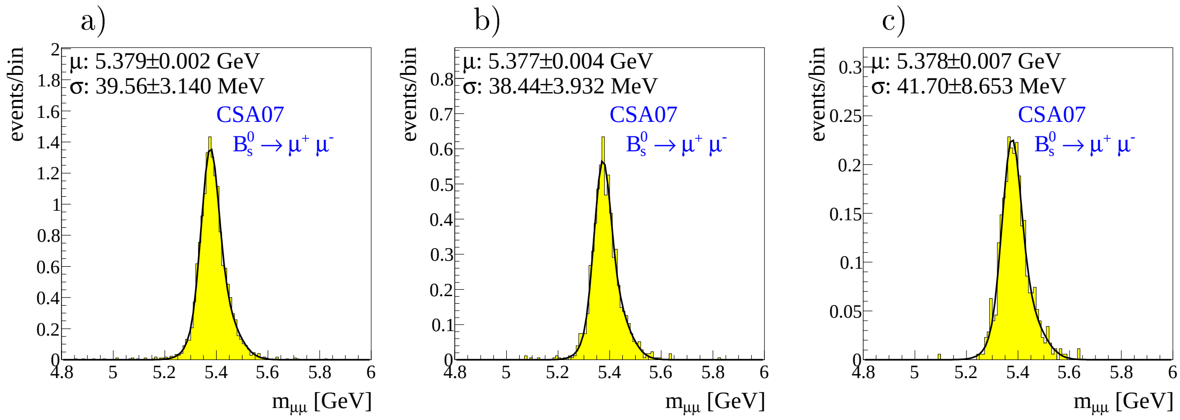

Figure 7.13: Reconstructed B_s candidates $m_{\mu\mu}$ distribution in signal MC, normalised to $1\,\mathrm{fb}^{-1}$ a) after HLT, b) before vertex and isolation cuts, c) after all analysis cuts.

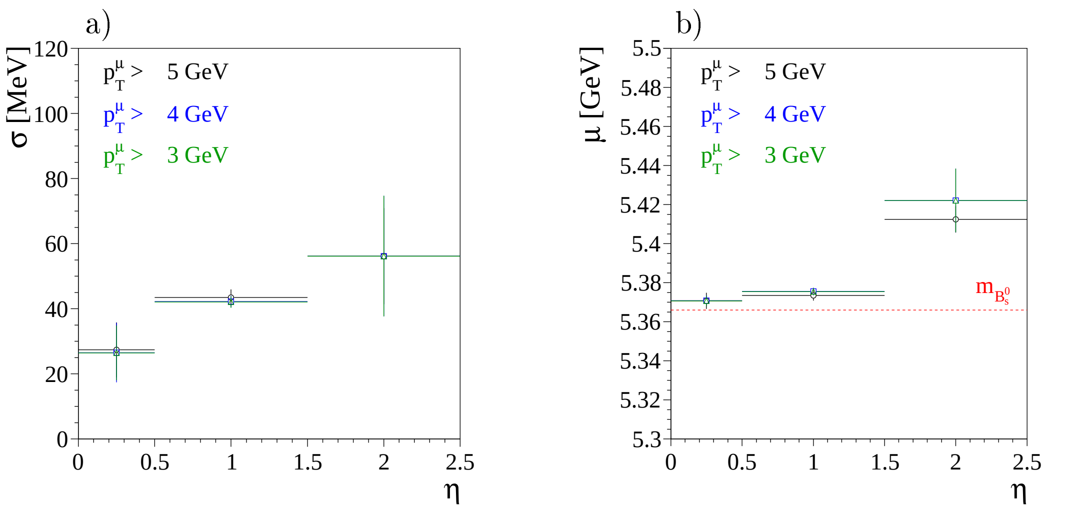

Figure 7.14: $|\eta|$-dependence of the invariant mass distribution of the reconstructed B_s candidates after HLT for different muon $p_\perp$-threshold. a) width σ and b) mean of $m_{\mu\mu}$ distribution.

7.4.1.3 B_s^0 Candidate Vertexing

Signal events are distinguished by two muons originating from the same secondary vertex while the muons in the $b\bar{b} \to \mu^+\mu^- + X$ background sample stem from separate vertices. Vertexing the two muons therefore provides a powerful handle in this background reduction. The transverse momentum vector of the B_s candidate must be close to the displacement of the secondary vertex from the primary vertex: the cosine of the opening angle between the two vectors must fulfill $\cos(\alpha) > 0.9985$, corresponding to an angular separation of about 3.1°. The flight length significance of the B_s^0 candidate is an excellent handle against (prompt) combinatorial background. The significance of the (unsigned) flight length l_{3D} is defined as l_{3D}/s_{3D}, where s_{3D} is the error on the flight length. Both the flight length and its error are determined by the standard CMSSW tool `VertexDistance3D`. The vertex quality is quantified by the fit-χ^2; for a vertex with two tracks the number of degrees of freedom is always 1. Figure 7.16 illustrates the distributions relevant for vertexing. It should be noted that the $b\bar{b} \to \mu^+\mu^- + X$ background distribution displays two peaks in this distribution: a second peak is off-scale at $\cos(\alpha) \sim 1$, this peak is absent for the signal sample.

7.4.1.4 B_s^0 Candidate Isolation

In high-$p_\perp$ gluon-splitting events the $b\bar{b}$ quark pair moves closely together due to their boost, and the two decay vertices of the resulting b-hadrons cannot be well separated in all cases. However, because of the other hadrons in semileptonic decays of both b-hadrons, the hadronic activity around the dimuon direction is enhanced compared to the signal decay. This is exploited in isolation requirements. The isolation I, as applied in the searches at the Tevatron, is determined from the B_s candidate transverse momentum and charged tracks with $p_\perp > 0.9\,\text{GeV}$ in a cone with half-radius $r = 1.0$ around the dimuon direction as follows:

$$I = \frac{p_\perp(B_s)}{p_\perp(B_s) + \sum_{trk} |p_\perp|}, \tag{7.14}$$

where all track parameters are evaluated at the origin. Figure 7.15c) illustrates the distribution of isolation variable I. For B_s candidates without any charged tracks above the transverse momentum cutoff, $I = 1$. The pronounced dip in the distribution just below I = 1 arises from the minimum transverse momentum requirement that implies a maximum value of I, depending on the transverse momentum of the B_s^0 candidate.

7.4.2 Factorising Selection Requirements

The efficiency for event selection on the signal and $b\bar{b} \to \mu^+\mu^- + X$ background is provided in Table 7.5. The application of all selection requirements leaves no remaining background event. Given the limited luminosity of the generated background sample, this does not allow to determine a reliable background estimate. However, the relatively mild correlation to the other selection criteria [85] allows a factorisation of the isolation I and χ^2 requirements from the other cuts: their efficiencies are determined on an

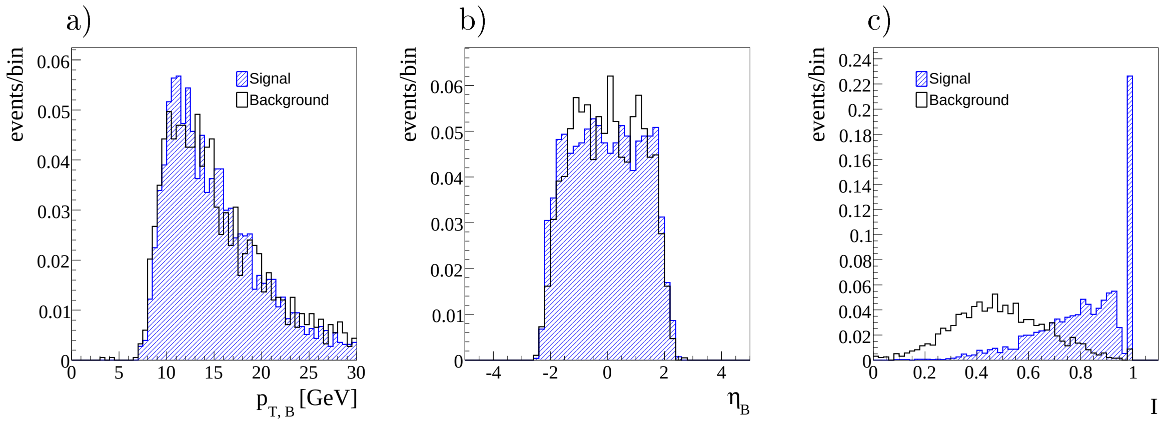

Figure 7.15: Reconstructed B_s candidates in the mass region $4.8 < m < 6.0\,\mathrm{GeV}$ (after HLT): a) transverse momentum, b) pseudo-rapidity, c) isolation. The histograms are normalised to unity.

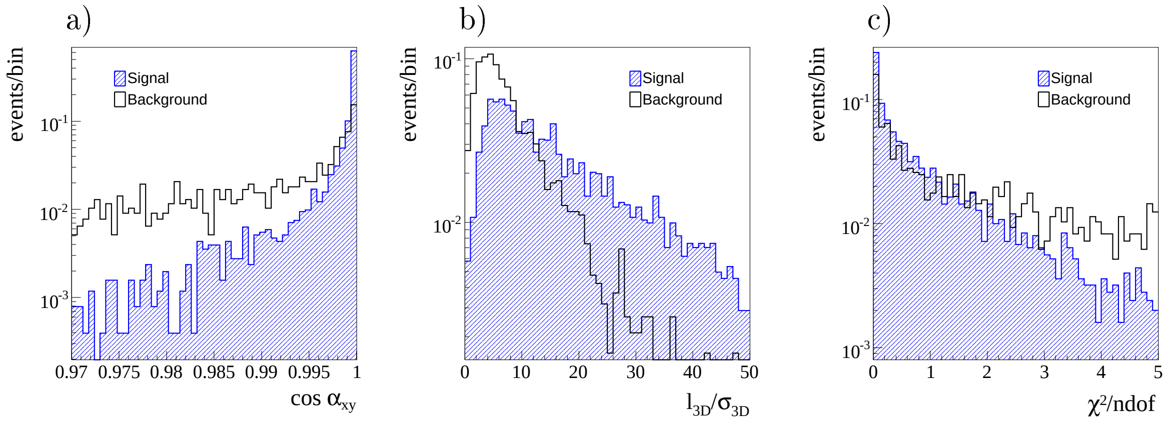

Figure 7.16: Secondary vertex distributions in the mass region $4.8 < m < 6.0\,\mathrm{GeV}$ (after HLT): a) cosine of the angle between the B_s candidate flight direction and secondary vertex in the transverse plane, b) flight length significance, c) χ^2/ndof of the secondary vertex fit. The histograms are normalised to unity.

event sample where the dimuon mass is $4.8 < m < 6.0\,\mathrm{GeV}$ and the significance of the secondary vertex separation is $l_{3D}/\sigma_{3D} > 7$. The expected $b\bar{b} \to \mu^+\mu^- + X$ background event yield is then obtained by multiplying the isolation and χ^2 efficiencies with the event yield after all the other cuts. This preselection is quite loose to provide enough statistics to allow tight χ^2 or I cuts, but still retains mostly those background events that mimic the signal event signature.

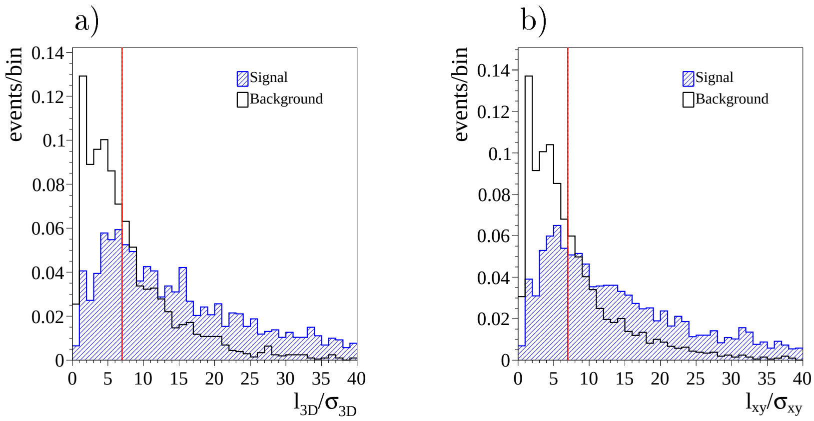

Figure 7.17: Signal and background distributions of the flight length significance after HLT in the mass region $4.8 < m < 6.0\,\text{GeV}$: a) in three dimensions and b) in the transverse plane. The preselection criteria was defined as $l_{3d}/\sigma_{3d} > 7$ and $l_{xy}/\sigma_{xy} > 7$ respectively.

7.4.3 Event Selection Summary

The total signal efficiency amounts to $\varepsilon = (2.64 \pm 0.120) \times 10^{-2}$, assuming factorisation of the I and χ^2 selection criteria it is $\varepsilon = (2.66 \pm 0.121) \times 10^{-2}$, arguably consistent with the former. Both errors are statistical only. For the $b\bar{b} \to \mu^+\mu^- + X$ dimuon background sample, the efficiency is determined to be $\varepsilon = (4.24 \pm 0.192) \times 10^{-8}$, assuming factorisation of these two criteria (statistical error only). Due to the limited MC statistics in the background sample, the simultaneous application of all cuts results in no remaining events and does not allow a determination of the background rejection without the factorisation assumption.

At this stage the $b\bar{b} \to \mu^+\mu^- + X$ background event yields have been obtained in the full mass window $4.8 < m_{\mu\mu} < 6.0\,\text{GeV}$. For the determination of the final sensitivity only the background yield in the signal window $m_{B_s} \pm 100\,\text{MeV}$ is relevant. This reduction factor $f = 0.17$ is determined by loosening the selection cuts to those at the HLT, and then determining the ratio of background events in that window to the total. With a linear background parametrisation, f varies only very weakly with the fit parameters. Figure 7.18 illustrates the non-peaking $b\bar{b} \to \mu^+\mu^- + X$ background $m_{\mu\mu}$ distribution after kinematic cuts and after HLT requirements.

The origin of the remaining background events has also been studied. In the mass window $0 < m_{\mu\mu} < 10\,\text{GeV}$, and before trigger requirements, all production processes (gluon fusion, flavour excitations, and gluon splitting) are present. Gluon fusion pro-

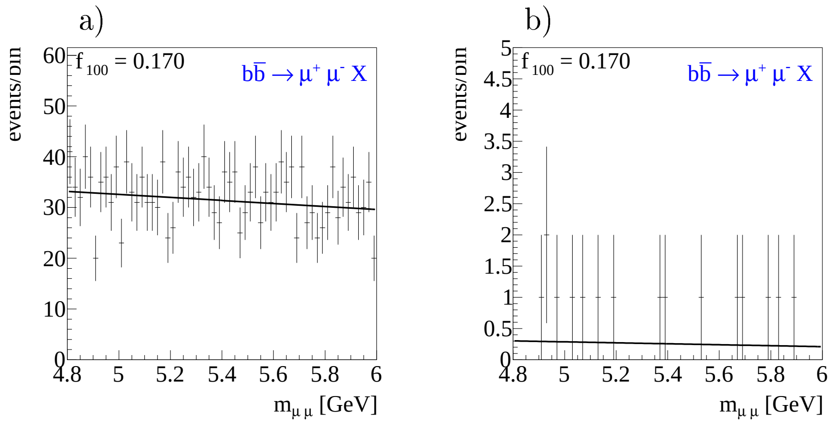

Figure 7.18: Background $m_{\mu\mu}$ distribution. a) after HLT, b) before vertex and isolation requirements.

cesses contribute at high mass with direct muons from both B-decays, while at low masses $m_{\mu\mu} \approx 2\,\text{GeV}$ (see Figure 7.4) muon pairs from direct and cascade decays contribute. After the HLT, gluon fusion processes no longer contribute significantly. In the mass window $4.8 < m_{\mu\mu} < 6.0\,\text{GeV}$ and after the full analysis chain, the remaining 16 events are from 14 gluon-splitting and 2 flavour-excitation events.

The non-peaking background from $c\bar{c}$ production has also been analyzed both in the Spring07 and CSA07 samples. However in both cases no events are left: in Spring07 there is no event left after the application of the flight length significance criterion, and in CSA07 the isolation requirement eliminates all remaining events (even in the factorising version). Table 7.6 summarises this together with the event reduction obtained in the Stew background sample.

Table 7.5: Event reduction and efficiency for the offline selection. The events are counted in the mass interval $4.8 < m_{\mu\mu} < 6.0\,\mathrm{GeV}$ and are normalised to a luminosity of $1\,\mathrm{fb}^{-1}$. The efficiencies for χ^2 and I, quoted in the middle part of the table, are determined relative to the event sample after the requirements of $4.8 < m_{\mu\mu} < 6.0\,\mathrm{GeV}$ and $l_{3D}/\sigma_{3D} > 17.0$ (different normalisation). The other efficiencies are cumulative. The total event selection efficiency and event yield are provided with and without the assumption of factorisation of the χ^2 and I cuts.

		Signal		$b\bar{b} \to \mu^+\mu^- + X$	
Description	Selection Criteria	Events	Efficiency	Events	Efficiency
gen. kinematics	see text	103	—	3.24×10^8	—
L1	see text	51.7	0.503	1.52×10^8	0.469
HLT (w/o mass cut)	see text	17.6	0.171	5.07×10^6	0.016
Good events	rec. candidate, PV	15.2	0.148	4.84×10^6	0.015
Mass cut	$4.8 < m_{\mu\mu} < 6.0\,\mathrm{GeV}$	15.1	0.147	2.30×10^5	7.09×10^{-4}
Pointing angle	$\cos(\alpha) > 0.9985$	11.0	0.107	2.46×10^4	7.58×10^{-5}
Flight distance	$l_{3d}/\sigma_{3d} > 17.0$	6.2	0.060	1979	6.10×10^{-6}
Vertex fit (diff. norm.)	$\chi^2 < 5.0$	—	0.940	—	0.406
Isolation (diff. norm.)	$I > 0.850$	—	0.469	—	0.017
Total	w/o factorisation	2.7	0.026	0.0	0.0
Total	w/ factorisation	2.7	0.027	13.8	4.24×10^{-8}
Signal window	$m_{B_s} \pm 100\,\mathrm{MeV}$	2.6 ± 0.079	0.025	$2.3^{+0.662}_{-0.516}$	7.20×10^{-9}

Table 7.6: Event reduction and efficiency for additional background samples. For other details see caption of Table 7.5.

		$c\bar{c} \to \mu^+\mu^- + X$		Stew	
Description	Selection Criteria	Events	Efficiency	Events	Efficiency
gen. kinematics	see text	9.29×10^7	—	5.95×10^8	—
L1	see text	4.22×10^7	0.454	3.76×10^8	0.149
HLT (w/o mass cut)	see text	1.26×10^6	0.014	1.75×10^7	0.007
Good events	rec. candidate, PV	9.86×10^5	0.011	1.68×10^7	0.007
Mass cut	$4.8 < m_{\mu\mu} < 6.0\,\mathrm{GeV}$	1.19×10^5	0.001	1.14×10^5	4.52×10^{-5}
Pointing angle	$\cos(\alpha) > 0.9985$	9213	9.91×10^{-5}	1.94×10^4	7.65×10^{-6}
Flight distance	$l_{3d}/\sigma_{3d} > 17.0$	97.0	1.04×10^{-6}	1.10×10^4	4.35×10^{-6}
Vertex fit (diff. norm.)	$\chi^2 < 5.0$	—	0.419	—	0.149
Isolation (diff. norm.)	$I > 0.850$	—	0.0	—	0.0
Total	w/o factorisation	0.0	0.0	0.0	0.0
Total	w/ factorisation	0.0	0.0	0.0	0.0
Signal window	$m_{B_s} \pm 100\,\mathrm{MeV}$	0.0	0.0	0.0	0.0

7.4.4 Selection Optimisation

The selection requirements presented in the previous sections are the result of a multi-dimensional grid search for the lowest achievable upper limit in regions around the values of the selection criteria used in the previous study [57].

As shown in Figure 7.19 the $\eta\phi$ separation and the invariant mass of the two muons in $b\bar{b} \to \mu^+\mu^- + X$ background sample are correlated. If only the events in the mass region $4.8 < m_{\mu\mu} < 6.0\,\text{GeV}$ are considered, the $\Delta R(\mu\mu)$ distribution of the background events becomes very similar to the one of the signal events. In fact, it was found that by omitting the $\Delta R(\mu\mu)$ selection criteria neither the overall background rejection power nor the overall signal selection efficiencies change. Therefore the $\Delta R(\mu\mu)$ selection criteria was removed from this analysis, but should eventually be reconsidered in samples if larger statistics are available.

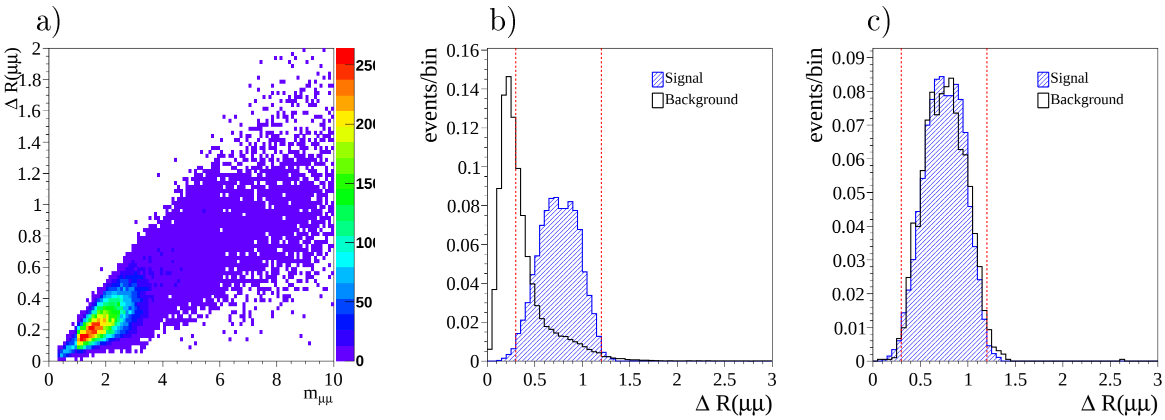

Figure 7.19: Correlation between $\eta\phi$ separation and invariant mass of the selected muons: a) two dimensional distribution of $\Delta R(\mu\mu)$ and $m_{\mu\mu}$ in background sample, $\Delta R(\mu\mu)$ distribution of signal and background b) after HLT, c) after HLT in mass window $4.8 < m_{\mu\mu} < 6.0\,\text{GeV}$. The histograms are normalised to unity.

After removing the $\Delta R(\mu\mu)$ cut, the remaining selection criteria were optimised by determining the lowest achievable upper limit on the $B_s^0 \to \mu^+\mu^-$ branching fraction in $1\,\text{fb}^{-1}$ in a multi-dimensional grid search. The value of each selection criteria was varied in a certain interval, as listed in Table 7.7 along with the number of division per interval. For each permutation of cut variables, the upper limit was calculated evaluating the final number of signal and background events using factorising vertex and isolation selection requirements. The grid search was performed twice, once using a decay length significance criterion in three dimension and once in the transverse plane.

Table 7.7: Optimisation of selection requirements.

Variable	Range	Number of steps	Step size
$p_\perp(\ell)$	3 ... 4 GeV	2	1
$p_\perp(B_s)$	5 ... 8 GeV	4	1
l_{xy}/σ_{xy}	5 ... 24	20	1
l_{3d}/σ_{3d}	5 ... 24	20	1
$\cos(\alpha)$	0.9980 ... 0.9995	4	0.0005
I	0.85 ... 0.95	4	0.05
χ^2	1 ... 8	8	1

The best combination of selection criteria is compared in Table 7.8 to the previous analysis [57]. The most important changes include the relaxation of the vertex χ^2 requirement and the tightening of the pointing angle requirement. The resulting best upper limits, when applying a decay length significance cut in either the transverse plane or in three dimensions, do not differ significantly. A choice was made in favour of a three dimensional decay length significance cut, since it gives a slightly better upper limits.

Table 7.8: Optimised selection requirements.

Previous analysis [57]	Present analysis
$p_\perp(\ell) > 3.0$	$p_\perp(\ell) > 4.0$
$0.3 < R_{\mu\mu} < 1.2$	removed
$4.8 < m_{\mu\mu} < 6.0\,\text{GeV}$	$4.8 < m_{\mu\mu} < 6.0\,\text{GeV}$
$p_\perp(B_s) > 5.0\,\text{GeV}$	$p_\perp(B_s) > 5.0\,\text{GeV}$
$\cos(\alpha) > 0.9950$	$\cos(\alpha) > 0.9985$
$l_{xy}/\sigma_{xy} > 18.0$	$l_{3d}/\sigma_{3d} > 17.0$
$\chi^2 < 1.0$	$\chi^2 < 5.0$
$I > 0.850$	$I > 0.850$

Figures 7.20-7.22 illustrate the previous and present selection criteria. The distributions show the the corresponding variable on a loosely preselected event selection before the cut is applied.

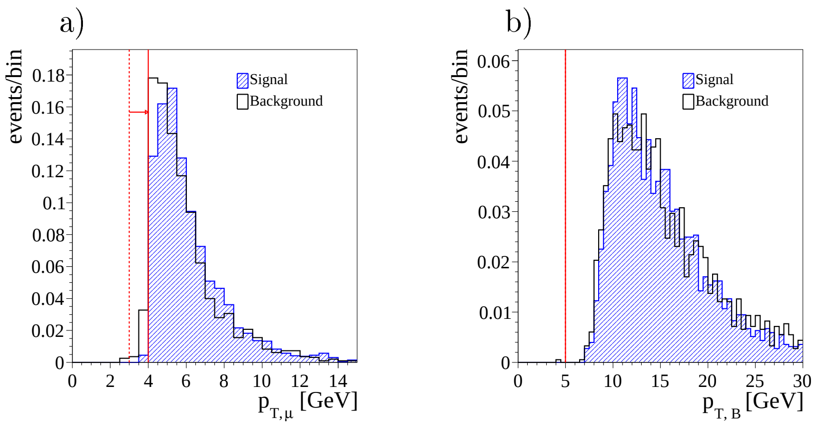

Figure 7.20: Transverse momentum of a) the muons and b)the reconstructed B_s candidates before the application of the $p_\perp(\ell)$ and $p_\perp(B_s)$ cut respectively. The previous and optimised selection criteria are indicated by the dashed and solid lines respectively.

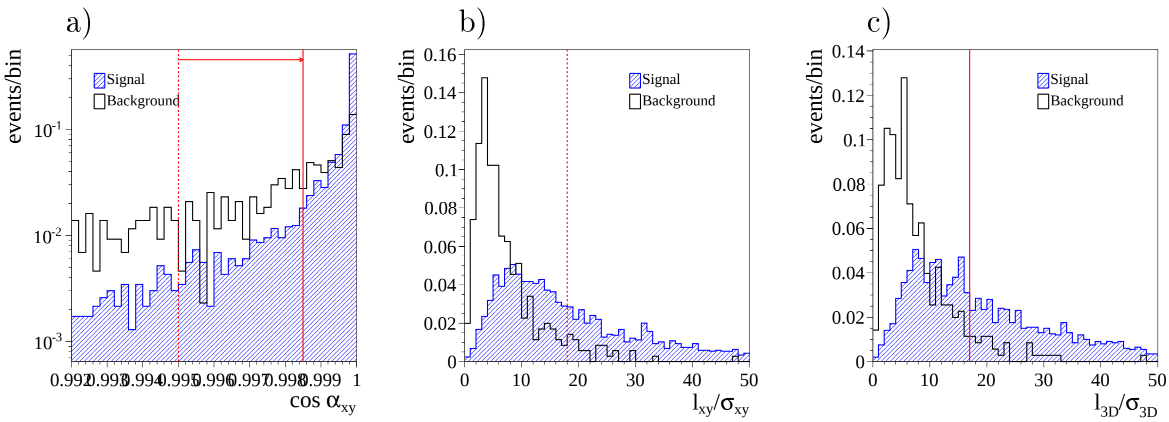

Figure 7.21: Distributions of a) pointing angle, b) decay length significance in the transverse plane and c) in three dimension, before the application the corresponding cut. The previous and optimised selection criteria are indicated by the dashed and solid lines respectively.

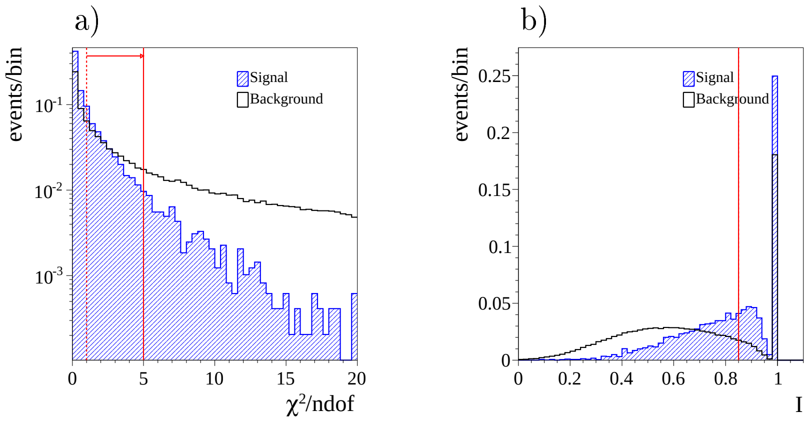

Figure 7.22: Distributions to determine the efficiencies of the two factorising cuts after applying the preselection criterion $l_{3d}/\sigma_{3d} > 7$, a) vertex fit χ^2/ndof and b) isolation. The previous and optimised selection criteria are indicated by the dashed and solid lines respectively.

7.5 The Normalisation Sample $B^\pm \to J/\psi\, K^\pm$

To minimise the dependence on the unknown $b\bar{b}$ production cross section and luminosity measurements, the analysis uses a normalisation sample $B^\pm \to J/\psi\, K^\pm$ with a signature similar to the signal decay $B_s^0 \to \mu^+\mu^-$. Many systematic errors cancel to first order when the upper limit is derived using a normalisation to a similar decay channel measured in data. The $B^\pm \to J/\psi\, K^\pm$ sample will furthermore allow a detailed comparison of the detector performance and analysis selection efficiencies in data and MC simulation. It will also allow the reweighing of the B^+ transverse momentum spectra so that the MC simulation reproduces the data.

The decay $B^\pm \to J/\psi\, K^\pm$ has a large and well-measured branching fraction with only one additional track in the final state compared to the signal decay. However, the hadronisation of the B^+ mesons can be different from the B_s^0 meson, affecting for instance the isolation variable. The dominant uncertainty here will be in the ratio f_s/f_u, which is of the order 15 %. The decay is reconstructed using requirements as similar to the signal mode as possible: the B^+ decay vertices are reconstructed using only the two muons and no mass-constraint on the J/ψ mass is applied. Table 7.9 summarises the selection criteria and their efficiencies of the normalisation analysis.

Figure 7.23 illustrates the mass resolution obtained on the $B^\pm \to J/\psi\, K^\pm$ event sample at various stages of the analysis. Figure 7.24 illustrates the combinatorial background to be expected from b-hadron decays into J/ψ mesons after subsequent

Table 7.9: Event reduction and efficiency for the offline selection applied to the normalisation $B^\pm \to J/\psi\, K^\pm$. For other details see caption of Table 7.5.

		Normalisation Signal		Background	
Description	Selection Criteria	Events	Efficiency	Events	Efficiency
gen. kinematics	see text	2.10×10^6	—	8.31×10^7	—
L1	see text	9.78×10^5	0.465	3.60×10^7	0.269
HLT (w/o mass cut)	see text	4.86×10^5	0.231	1.51×10^7	0.113
Good Event	rec. candidate, PV	3.66×10^5	0.174	1.07×10^7	0.080
Mass cut	$4.8 < m_{\mu\mu K} < 6.0\,\mathrm{GeV}$	1.82×10^5	0.086	8.64×10^5	0.006
Pointing angle	$\cos(\alpha) > 0.9985$	1.38×10^5	0.066	4.34×10^5	0.003
Flight distance	$l_{3d}/\sigma_{3d} > 17.0$	9.42×10^4	0.045	2.36×10^5	0.002
Vertex fit (diff. norm.)	$\chi^2 < 5.0$	—	0.900	—	0.501
Isolation (diff. norm.)	$I > 0.850$	—	0.412	—	0.387
Total	w/o factorisation	3.14×10^4	0.015	4.02×10^4	3.00×10^{-4}
Total	w/ factorisation	3.49×10^4	0.017	4.56×10^4	3.41×10^{-4}
Signal window	$m_{B^\pm} \pm 100\,\mathrm{MeV}$	3.29×10^4	0.016	6045	4.52×10^{-5}

requirements. While the background is not negligible, it is not expected to pose a significant problem for the extraction of the normalisation yield. The background shape is well described by an exponential function; the experience at CDF and D0 does not indicate any evidence that Cabibbo-suppressed $B^+ \to J/\psi\,\pi^+$ decays appear at a significant level. The signal yields determined in the $B^\pm \to J/\psi\, K^\pm$ signal MC sample, given in the third-last row in table 7.9, agree within the statistical uncertainties well with the signal yields obtained in the non-prompt J/ψ sample in Figure 7.24.

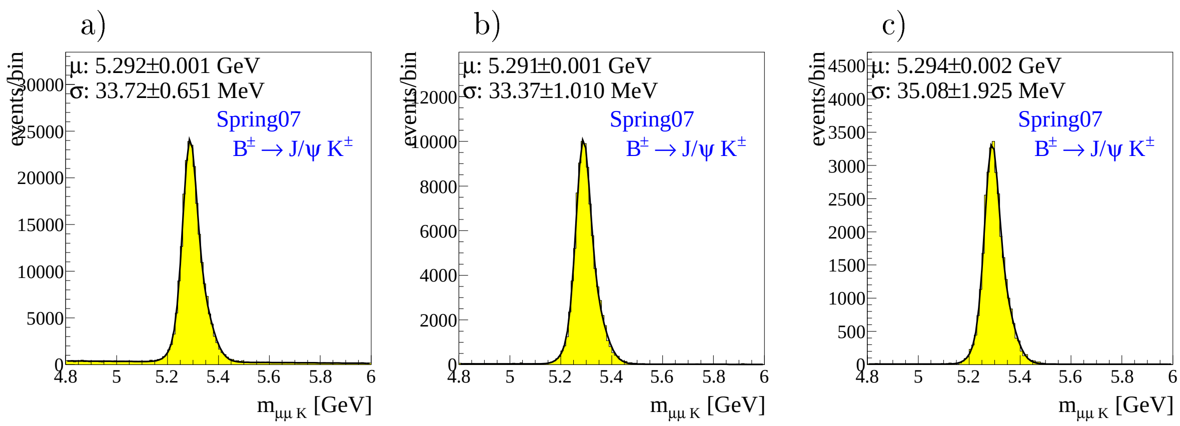

Figure 7.23: Reconstructed B^+ candidates $m_{\mu\mu K}$ distribution in signal MC, normalised to $1\,\mathrm{fb}^{-1}$ a) after HLT, b) before vertex and isolation cuts, c) after all analysis cuts.

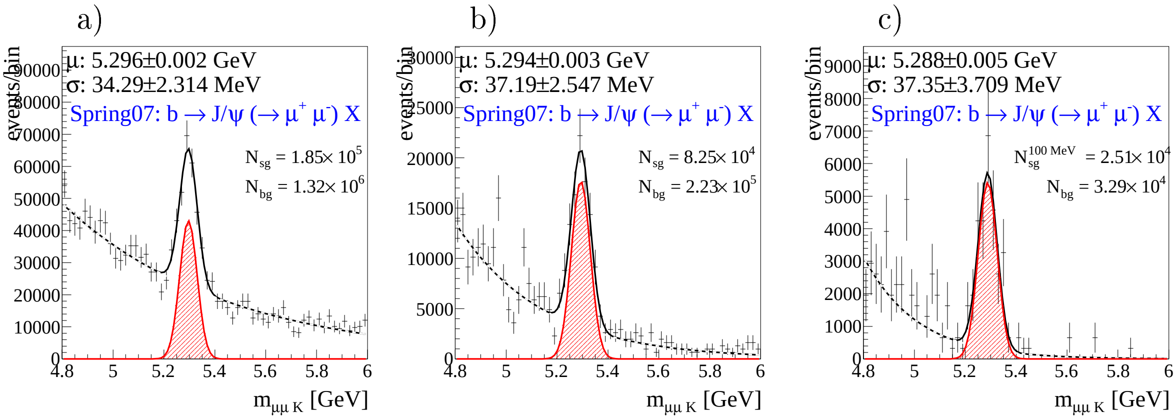

Figure 7.24: Reconstruction of $B^{\pm} \to J/\psi\, K^{\pm}$: Signal and background (combinatorial background in $b \to J/\psi\,(\to \mu^+\mu^-)X$ sample), normalised to $1\,\text{fb}^{-1}$. a) after HLT, b) before vertex and isolation cuts, c) after all cuts (w/o factorisation).

The various distribution for $B_s^0 \to \mu^+\mu^-$ and $B^{\pm} \to J/\psi\, K^{\pm}$ after the HLT requirement and in the mass window $4.8 < m_{\mu\mu(K)} < 6.0\,\text{GeV}$ are illustrated in Appendix D. The agreement between the respective distributions is quite good, indicating that the reconstruction of the normalisation sample and the signal sample is very similar.

7.6 Background Study

This section quantifies the background from rare decays of one b-hadron, with or without muons in the final state. The background contributions from the combination of one muon with a misidentified hadron have been investigated with a generator-level study in [85]. A variety of rare decay channels—as listed in Table 7.2—has been studied. In Table 7.10 the efficiency for the event selection in the rare b-hadron background samples is presented. The mass distributions of the rare b-hadron decay backgrounds before the application of selection criteria are illustrated in Appendix E. A few remarks on these mass distributions can be made:

- Often the invariant mass distributions shows two different components: one at lower invariant masses due to the decay channel under study, and one at higher in variant masses due to the combination of one final state particle with another muon from the semileptonic decay of the other b-hadron in the event.

- Semileptonic decays are not a problem as the good mass resolution provides for sufficient separation between the upper edge of the continuum mass distribution and the $B_s^0 \to \mu^+\mu^-$ signal region.

- Decays of Λ_b hadrons constitute a peaking background in the signal region. Their rate, however, is very strongly suppressed and their expected background contribution is very small even before any selection criteria (see Table 7.2).

- The good mass resolution also significantly reduces background from rare hadronic B-decays, so that only a minor fraction of the tail (the central value of their mass distribution is shifted because of the wrong mass hypothesis) is leaking into the $B_s^0 \to \mu^+\mu^-$ signal region.

Table 7.10: Rare background contributions expected in $1\,\text{fb}^{-1}$ calculated using a misidentification probability of $\varepsilon_\pi = 0.6\%$ for pions, $\varepsilon_K = 1.1\%$, for kaons and $\varepsilon_p = 0.2\%$ for protons. The initial number of events is the number of expected background events reduced by the misidentification probability and the HLT efficiency.

	Final State	$2h$ from b-hadron		$1h + 1\ \mu + X$		$2\ \mu + X$	
Description	Selection Criteria	Events	Efficiency	Events	Efficiency	Events	Efficiency
generated events	see text	19.7	—	8258	—	2.47×10^4	—
Pointing angle	$\cos(\alpha) > 0.9985$	2.6	0.131	89.4	0.011	38.5	0.002
Flight distance	$l_{3D}/\sigma_{3D} > 17.0$	1.5	0.074	54.4	0.007	25.3	0.001
Vertex fit (diff. norm.)	$\chi^2 < 5.0$	—	0.933	—	0.916	—	0.958
Isolation (diff. norm.)	$I > 0.850$	—	0.600	—	0.410	—	0.149
Total	w/o factorisation	0.829	0.042	19.3	0.002	2.3	9.25×10^{-5}
Total	w/ factorisation	0.819	0.041	20.3	0.002	3.7	1.48×10^{-4}
Signal window	$m_{B_s} \pm 100\,\text{MeV}$	0.399	0.020	1.1	1.30×10^{-4}	0.019	7.50×10^{-7}

Figure 7.25 illustrates the background contributions of the different rare b-hadron decays expected in $1\,\text{fb}^{-1}$, using factorising selection criteria. For each decay, the invariant mass distributions before the application of the factorising selection criteria is shown, where the numbers have been scaled using the efficiency of the factorising vertex and isolation requirements, and applying the muon misidentification probabilities from section 7.3.3 as scaling weights ($\varepsilon_\pi = 0.6\%$ for pions, $\varepsilon_K = 1.1\%$ for kaons and $\varepsilon_p = 0.2\%$ for protons). The number quoted in brackets next to each decay corresponds to the number of events in the signal window $m_{B_s} \pm 100\,\text{MeV}$. In total, rare b-hadron decay backgrounds contribute to the overall background with an additional $n_B^{\text{rare}} = 1.5$ events.

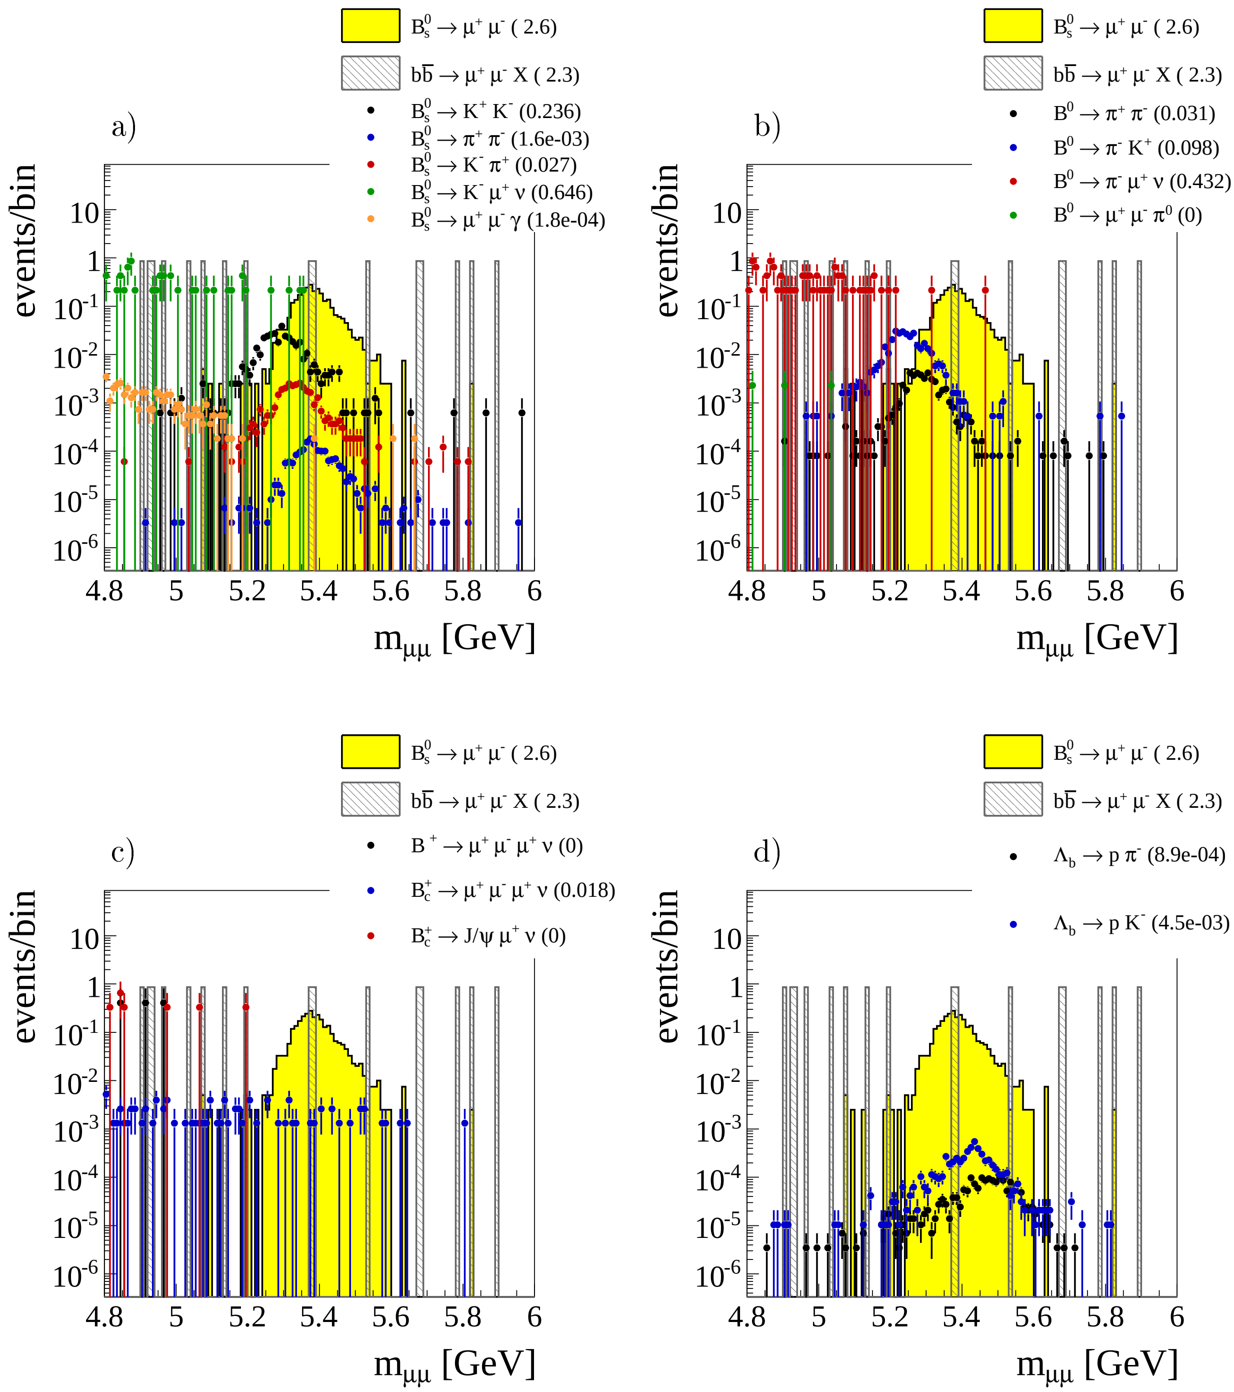

Figure 7.25: Signal and background $m_{\mu\mu}$ distributions normalised to a luminosity of $1\,\text{fb}^{-1}$, after the application of all selection criteria (using factorising selection criteria) and including muon misidentification probabilities. Background contributions of a) B_s decays, b) B_d decays, c) B_c and B_u decays, d) Λ_b decays. The number quoted in brackets next to each decay corresponds to the number of events expected in the signal window $m_{B_s} \pm 100\,\text{MeV}$ in $1\,\text{fb}^{-1}$.

7.7 Systematics

The upper limit is affected by the statistical and systematic errors through the resulting uncertainty in the signal efficiency and the background yield. The calculation of the upper limit, described in section 7.8, requires as input the errors on signal efficiency and background yield.

7.7.1 Muon Identification

The uncertainty on the muon identification efficiency has no influence on the signal efficiency, as it cancels to first order in the ratio with the normalisation sample. It affects the background uncertainty, however. It is assumed that the muon identification efficiency will be determined with an error of 5 %. The hadron misidentification probabilities for the determination of the hadron background have been varied by ±20 %; the background uncertainty amounts to 6 %. Kaon misidentification is the dominant source for this uncertainty.

7.7.2 Tracking

The tracking efficiency uncertainty is assumed to be 5 %. It will be determined by a dedicated study group of the CMS tracking POG. The effects of this uncertainty are on the one hand in the signal track reconstruction, and on the other hand the isolation criteria is affected. Since the normalisation sample has one additional kaon track in the final state, the tracking uncertainty will affect this directly by 5 %. The uncertainty due to the tracker misalignment is estimated based on the efficiency difference of the vertex χ^2 requirement between the perfectly aligned Spring07 signal sample (summarised in Table 7.11) and the $100\,\mathrm{pb}^{-1}$ alignment conditions in the CSA07 signal sample. This gives a signal efficiency uncertainty of 3 % and a background uncertainty of 5 %.

7.7.3 Factorising Selection Requirement

Because of the limited statistics in the background samples, the selection requirements for the vertex fit χ^2 and isolation are studied independently on an enlarged dataset. The signal efficiency differs by 1 % between the factorising and simultaneous analysis efficiency. In the normalisation sample a difference of 10 %, and 12 % for the normalisation background was found. On the background sample, the two efficiencies cannot

Table 7.11: Event reduction and efficiency for the offline selection for signal and background in the Spring07 event samples (perfect alignment). For other details, see the caption of Table 7.5.

		Signal		$b\bar{b} \to \mu^+\mu^- + X$	
Description	Selection Criteria	Events	Efficiency	Events	Efficiency
gen. kinematics	see text	103	—	1.32×10^8	—
L1	see text	52.6	0.512	3.89×10^7	0.295
HLT (w/o mass cut)	see text	23.4	0.228	3.97×10^6	0.030
Good events	rec. candidate, PV	17.8	0.173	2.89×10^6	0.022
Mass cut	$4.8 < m_{\mu\mu} < 6.0\,\mathrm{GeV}$	14.4	0.140	1.62×10^5	0.001
Pointing angle	$\cos(\alpha) > 0.9985$	10.0	0.098	1.74×10^4	1.32×10^{-4}
Flight distance	$l_{3d}/\sigma_{3d} > 17.0$	6.76	0.066	2739	2.07×10^{-5}
Vertex fit (diff. norm.)	$\chi^2 < 5.0$	—	0.910	—	0.391
Isolation (diff. norm.)	$I > 0.850$	—	0.491	—	0.017
Total	w/o factorisation	2.82	0.027	0.0	0.0
Total	w/ factorisation	3.02	0.029	17.9	1.36×10^{-7}
Signal window	$m_{B_s} \pm 100\,\mathrm{MeV}$	2.80 ± 0.37	0.027	$3.1^{+0.934}_{-0.715}$	2.32×10^{-8}

be compared as no event survives the complete analysis chain. The systematic error for the background yield is assumed to be 20 %.

7.7.4 Trigger Efficiency

An uncertainty of 5 % (relative) for each the L1 and HLT efficiency is assumed. This propagates directly into a 5 % uncertainty on signal efficiency and background yield.

7.7.5 Hadronisation Uncertainties in the Normalisation Sample

The normalisation for this analysis will rely on the measurement of a control sample in data (as in the analyses of CDF and D0). The largest external uncertainties here is from the ratio of fragmentation probabilities f_s and f_d. The uncertainty amounts to 15 %.

7.7.6 Summary

Combining the systematic error, summarised in Table 7.12, quadratically with the statistical error, the signal efficiency is known to about 18 %, while the background yield uncertainty amounts to about 37 %.

Table 7.12: Summary of systematic uncertainties.

Source	$\Delta\varepsilon_{\text{Signal}}$	$\Delta\varepsilon_{\text{Background}}$
Muon ID efficiency	-	5 %
Muon misID probability	-	6 %
L1 Efficiency	5 %	5 %
HLT Efficiency	5 %	5 %
Misalignment	3 %	5 %
Kaon tracking efficiency	5 %	-
Factorising selection	1 %	20 %
f_s/f_u	15 %	-
Total	18 %	23 %

7.8 Results

Using the event and candidate selection described in section 7.4 the total cumulative selection efficiency for signal events is $\varepsilon_S = 0.025$ and the background reduction factor is $\varepsilon_B = 7.20 \times 10^{-9}$. With this selection, the first $1\,\text{fb}^{-1}$ of integrated luminosity will yield $n_S = 2.7$ signal events and $n_B = 2.3$ background events in the signal window $m_{B_s} \pm 100\,\text{MeV}$. Additional background events in the mass window arise from rare decays of b-hadrons as described in section 7.6. The total contribution of these events is $n_B^{\text{rare}} = 1.5$, giving a total background contribution of $n_B^{\text{tot}} = 3.8$. As described in section 7.7, the combined statistical and systematic uncertainties of the background estimate is 37 % and for the signal efficiency it is 18 %.

Using the tools in [93] the signal can be extracted with a significance ScP = 0.6, which is too low to claim a significant observation. Therefore, the main result of this analysis is the expected upper limit that can be achieved in this data sample.

The upper limit on the number of observed signal events is determined following the Bayesian procedure described as in [94], using the function

```
blimit(double c, int n, double a, double aS, double b, double bS, double g)
```

where `b` is the confidence level (0.9 in our case for a 90 % C.L.). The (expected) number of observed events `n` $= n_S + n_B$ is computed from the expected signal yield and the non-peaking plus peaking backgrounds. By setting the signal acceptance `a` = 1, the function `blimit` will return the number of observed signal events and not the number of produced signal events. The acceptance error `aS` is set to the relative efficiency error (quadratic sum of statistical and systematic error). The background yield `b` contains both peaking and non-peaking contributions, its error `bS` is the quadratic sum of statistical and systematic error. With $\mathtt{g} = 1$ a flat prior is obtained.

The upper limit on the branching fraction has been determined $B_s^0 \to \mu^+\mu^-$ in the two ways, described in sections 7.8.1 and 7.8.2. The first method relies on an absolute normalisation and will not be used in data. The second method has been described in subsection 7.1.3 and relies on a normalisation sample. In both approaches, the number $N(n_{obs}, n_B, n_S)$ is the number of signal candidate $B_s^0 \to \mu^+\mu^-$ decays at the 90 % C.L., estimated using the Bayesian approach of [94], where n_{obs} is the expected number of observed events given n_B and n_S expected background and signal events.

7.8.1 Result with Absolute Normalisation

In this first approach, the upper limit is determined as

$$\mathcal{B}(B_s^0 \to \mu^+\mu^-) \leq \frac{N(n_{obs}, n_B, n_S)}{\varepsilon_{\text{gen}}\, \varepsilon_{\text{total}}\, N_{B_s}}.$$

The a priori expected limit is given by the average of all possible observations, randomly sampled from a Poisson distribution with mean $n_{obs} = n_B + n_S$. The number of produced B_s mesons, $N_{B_s} = 1.05 \times 10^{11}$, is computed from the 'known' cross section and luminosity of the MC event sample (in the real analysis with data, this will be normalised to a control sample with well-measured branching fraction). The efficiency is divided into two parts: $\varepsilon_{\text{gen}} = 2.54 \times 10^{-1}$ is the kinematic acceptance that a produced B_s meson decays into two muons satisfying the generator level cuts described in section 7.1. The efficiency $\varepsilon_{\text{total}} = 0.025$ is the cumulative efficiency of the complete analysis chain. For the expected signal and background event numbers $n_S = 2.6$ and $n_B^{\text{tot}} = 3.8$, the expected number of observed events is $n_{obs} = 6.4$ and the corresponding number of signal candidate decays is $N(n_{obs}, n_B^{tot}, n_S) = 8.5$ at the 90 % C.L. The gives the following upper limit on the branching fraction

$$\mathcal{B}(B_s^0 \to \mu^+\mu^-) \leq 1.3 \times 10^{-8}.$$

7.8.2 Result with Normalisation Sample

The second method of extracting the upper limit is based on the normalisation sample as described in subsection 7.1.3. The upper limit at the 90 % C.L. is calculated from

$$\mathcal{B}(B_s^0 \to \mu^+\mu^-; 90\,\%\mathrm{C}.L.)$$
$$= \frac{N(n_{obs}, n_B, n_S)}{N(B^\pm \to J/\psi\, K^\pm)} \cdot \frac{f_u}{f_s} \cdot \frac{\alpha_{B^+}}{\alpha_{B_s^0}} \cdot \frac{\varepsilon^{\mathrm{trg}}_{B^+}}{\varepsilon^{\mathrm{trg}}_{B_s^0}} \cdot \frac{\varepsilon^{\mathrm{ana}}_{B^+}}{\varepsilon^{\mathrm{ana}}_{B_s^0}} \cdot \mathcal{B}(B^\pm \to J/\psi\, K^\pm) \cdot \mathcal{B}(J/\psi \to \mu^+\mu^-),$$

where $\alpha_{B_s^0}$ (α_{B^+}) is the generator-level acceptance for signal (normalisation) events, $\varepsilon^{\mathrm{trg}}_{B_s^0}$ ($\varepsilon^{\mathrm{trg}}_{B^+}$) is the trigger efficiency for signal (normalisation) events, $\varepsilon^{\mathrm{ana}}_{B_s^0}$ ($\varepsilon^{\mathrm{ana}}_{B^+}$) is the analysis efficiency for signal (normalisation) events, and $\mathcal{B}(B^\pm \to J/\psi\, K^\pm) = (1.007 \pm 0.035) \times 10^3$ and $\mathcal{B}(J/\psi \to \mu^+\mu^-) = (5.93 \pm 0.06) \times 10^2$, and finally $f_s = (10.50.9)\,\%$ and $f_u = (40.20.9)\,\%$.

Using the event and candidate selection described in section 7.5 the total cumulative selection efficiency for signal events in the normalisation channel is $\varepsilon_{\mathrm{tot,N}} = 0.016$. By normalising to the the number of $B^\pm \to J/\psi\, K^\pm$ events $n_N = 3.29 \times 10^4$, the resulting on the branching fraction is given by

$$\mathcal{B}(B_s^0 \to \mu^+\mu^-) \leq 1.3 \times 10^{-8}.$$

This determination of the upper limit is consistent with the upper limit resulting from the absolute normalisation.

Conclusions

The Large Hadron Collider (LHC), scheduled to start collisions in October 2009, will open up the door to a new energy regime and put the Standard Model of Particle Physics and its extensions to the test. One of the most appealing extensions of the Standard Model is the Minimal Supersymmetric Standard Model (MSSM) which requires two Higgs doublets (H_u, H_d) and gives rise to three neutral Higgs bosons h, H, A and two charged Higgs bosons $H^{\pm}$ after symmetry breaking. The size of the coupling to down-type fermions is controlled by the free parameter $\tan\beta$, given by the ratio v_u/v_d of the two vacuum expectation values acquired by H_u and H_d. $\tan\beta$ is very difficult to measure yet is crucial for predicting and interpreting the observables of the MSSM. Flavour changing neutral currents (FCNC) mediated by Higgs bosons may provide a powerful tool to constrain the parameters of the MSSM. In particular, the observation of the flavour changing process $B_s^0 \to \mu^+\mu^-$ allows the determination of a lower bound on $\tan\beta$ and at the same time constrains the mass scale of M_A.

In this work, a search strategy for the rare decay $B_s^0 \to \mu^+\mu^-$ with the CMS experiment has been presented. Based on Monte Carlo samples from the official CSA07 and Spring07 productions, the study predicts an expected upper limit on the Standard Model branching fraction of 1.3×10^{-8} at the 90% confidence level with $1\,\text{fb}^{-1}$ of integrated luminosity. The upper limit is calculated using the decay $B^{\pm} \to J/\psi\, K^{\pm}$ as normalisation channel and includes the background contributions arising from semimuonic decays of both b-hadrons as well as from rare b-hadron decays. Systematic and statical uncertainties have been included in the upper limit estimation using a Bayesian approach.

The online event selection is based on a Level-1 dimuon trigger requiring both muons to have a $p_{\perp} > 3\,\text{GeV}$ and on a dedicated high-level trigger based on a displaced dimuon trigger. The offline analysis is based on selection variables related to the muon candidates, and the reconstructed B_s candidates and their associated secondary vertices.

The selection criteria were optimized by determining the lowest achievable upper limit on the branching fraction in a multi-dimensional grid search. In addition, the determination of the muon identification efficiency from data has been demonstrated using the 'tag and probe' method.

The results of this study promise an interesting start-up analysis with the possibility of setting tight constraints on the MSSM. With sufficient integrated luminosity, the precision measurement of the $B_s^0 \to \mu^+\mu^-$ branching fraction will set constraints on models of new physics. The results obtained are also comparable to those of ATLAS [95].

In the hardware related part of this work, the module qualification procedure has been presented. This procedure comprised of two test suites which included a multitude of functionality, calibration, performance and optimisation tests. In addition, a set of grading criteria defining three categories of module quality has been derived. These criteria were based on the performance and lifetime requirements of the pixel detector in the expected experimental conditions of CMS. Both test suite procedures, post-processing of data generated by the test algorithms and grade evaluation based on ROC and sensor performance were fully automatised.

In total, 971 fully assembled detector modules entered the process of module qualification. 813 modules were considered acceptable (grades A and B). 158 modules were deemed unacceptable for use in the detector system (grade C). The prevalent reasons for this, were high sensor leakage currents and/or large amounts of pixel functionality defects on one or several ROCs. The final yields are as follows:

- Of the 848 tested full modules, 63 % were of grade A, 21 % of grade B and 16 % of grade C.
- Of the 123 tested half-modules, 65 % were of grade A, 19 % of grade B and 16 % of grade C.

768 modules have been implemented in the final barrel pixel detector (672 full and 96 half modules). Of these, 75 % are of grade A quality and 25 % of grade B quality. The differentiation between grade A and B modules can primarily be attributed to sensor IV-characteristics, with bump bonding defects having a second order effect. In July 2008, the pixel detector system was installed in CMS and eagerly awaits the exploration of a new world of physics at the TeV energy scale.

Appendix A

DACs and Registers

Table A.1: DACs and registers sorted by category.

	Category	Bit	Name	Action
	Voltage Regulators	8	Vana	analogue voltage
		4	Vdig	digital voltage
		4	VComp	supply voltage of comparator
		8	Vsf	linear behaviour of the pulse height in the low Vcal range
Pixel Unit Cell	Analogue	8	VwllPr	preamplifier feedback
		4	VrgPr	preamplifier feedback
		8	VwllSh	shaper feedback
		4	VrgSh	shaper feedback
		8	Vtrim	trim bits scale factor
		8	VthrComp	comparator threshold
		8	VhldDel	hold delay
		8	Vleak_comp	sensor leakage current compensation
		8	VIColOr	current sent to periphery
	Trigger	8	Vnpix	min. number of pixel hits per d. c.
		8	VSumCol	min. number of double columns
	Calibrate	8	Vcal	pulse height of calibration signal
		8	CalDel	delay of calibration signal

Table A.2: DACs and registers sorted by category (continued).

	Category	Bit	Name	
D. C. Periphery	Pixel readout	4	Vbias_sf	shifts the pulse height range
		8	VIbias_bus	threshold for the voltage conversion of pixel address currents
	Double column readout	8	VIbiasOp	
		8	VoffsetOp	shifts the pulse height range
		8	VIon	stretches the pulse height range
		8	VOffsetRO	shifts the pulse height range
	Control and Interface Block	8	Ibias_DAC	analogue level of ROCs
		8	VIbias_PH	stretches the pulse height range
		8	VIbias_roc	stretches the pulse height range and the address levels
	Registers	8	CtrlReg	low/high Vcal range, full/half speed and chip enabled/disabled
		8	WBC	trigger latency
		8	RangeTemp	temperature measurement range
	TBM	8	Inputbias	scales the signal
		8	Outputbias	scales the signal
		8	Dacgain	analogue level of TBM

Appendix B

DAC Default Settings

Table B.1: Default settings and dynamic optimisation (denoted with *) of DACs.

DAC	Default	Optimisation criteria
Vana	150*	analog current is 24 mA
Vdig	6	address levels: linear behaviour of amplifier and below external voltage (2.5 V)
VComp	10	reliable operation, fallback solution if trimming doesn't work anymore (after irradiation)
Vsf	150*	optimise linearity in low range while keeping digital current below 5 μA
Vleak_Comp	0	compensation of leakage current after irradiation
VwllPr VwllSh VrgPr VrgSh	35 35 0 0	compromise between maximum pulse height and minimum time walk (the four preamplifier/shaper system DACs are set simultaneously, the two DAC pairs are set to the same value by design)
Vtrim	7*	lower highest pixel threshold to lowest pixel threshold on ROC while all trim bits are on
VthrComp	90*	different settings during trimming (determine Vtrim/trim bit while VthrComp is set to minimum of pixel threshold distributions at Vcal = 60) and calibration (stable point in VthrComp-CalDel readout distribution of one pixel at Vcal = 200)

Table B.2: Default settings and dynamic optimisation (denoted with *) of DACs.

DAC	Default	Optimisation criteria
VhldDel	160	stable sampling point for different pixel and different Vcal values: flat distribution around maximum pulse height and distinguishable for Vcal
VIColOr	99	arbitrary, no influence on pulse height above 20
Vnpix VSumCol	0 0	only self triggering mode with Marlon Trigger Chip (MTC)—not used
Vcal	200	-
CalDel	70*	center of readout range at $\text{VthrComp}\|_{\text{threshold}} + 50$
VIbias_bus	30	reliable address level conversion
Vbias_sf	10	reliable operation (for pulse height shift see VoffsetR0/VoffsetOp)
VIbiasOp	50	no influence on linearity with respect to VoffsetOp over the whole range (but no signal below ≈ 20)
VoffsetOp	40*	shift pulse height range to target ADC range after VIbias_PH optimisation (linearity high range)
VIon	130	no influence on linearity with respect to VoffsetRO and VIbiasOp for VIon > 110 (for pulse height stretch see VIbias_PH instead)
VOffsetR0	120	pulse height range can be shifted to any ADC range with VoffsetOp (linearity high range)
Ibias_DAC	90*	set ultrablack levels of ROCs to TBM ultrablacks (fixes the position of all other levels, maximum level at +1000)
VIbias_PH	220*	stretch/squeeze pulse height height to 2000 (from −1000 to +1000)
VIbias_roc	220	maximum address level stretch and pulse height ADC range (see VIbias_PH instead)
Inputbias	128	no influence on pulse height above 110
Outputbias	128	no influence on pulse height above 110
Dacgain	128*	TBM ultrablack below -1000 for both channels with least difference to -1000 (fixes the position of address levels)

Appendix C

Signal and Background Distributions

Figures C.1-C.5 illustrate the distributions of different selection variables for the signal $B_s^0 \rightarrow \mu^+\mu^-$and background $b\bar{b} \rightarrow \mu^+\mu^- + X$, after the HLT requirement and in the mass window $4.8 < m_{\mu\mu} < 6.0\,\text{GeV}$. The histograms are normalised to unity.

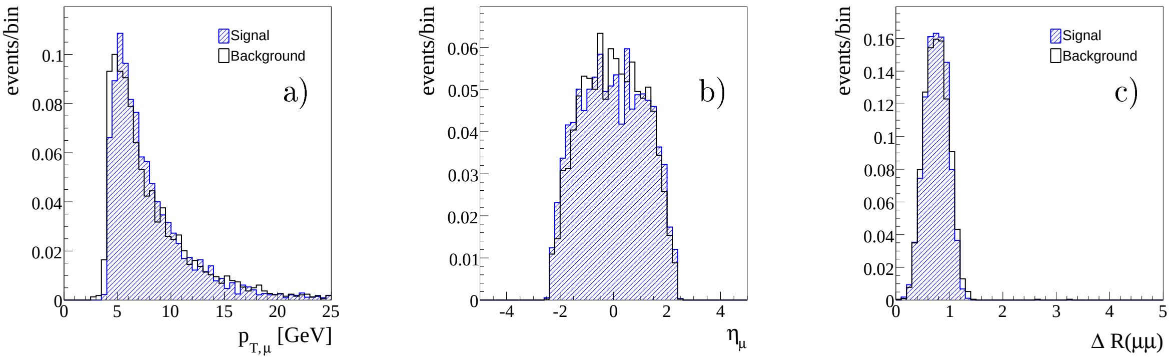

Figure C.1: Muon variable distributions: a) transverse momentum, b) pseudo-rapidity, c) $\eta\phi$ separation of the two muons.

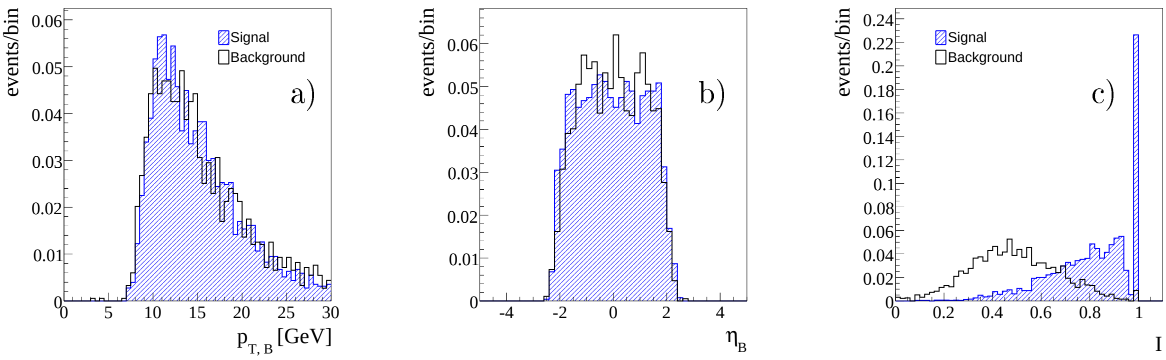

Figure C.2: Reconstructed B_s^0 candidates: a) transverse momentum, b) pseudo-rapidity, c) Isolation of the B_s^0 candidate.

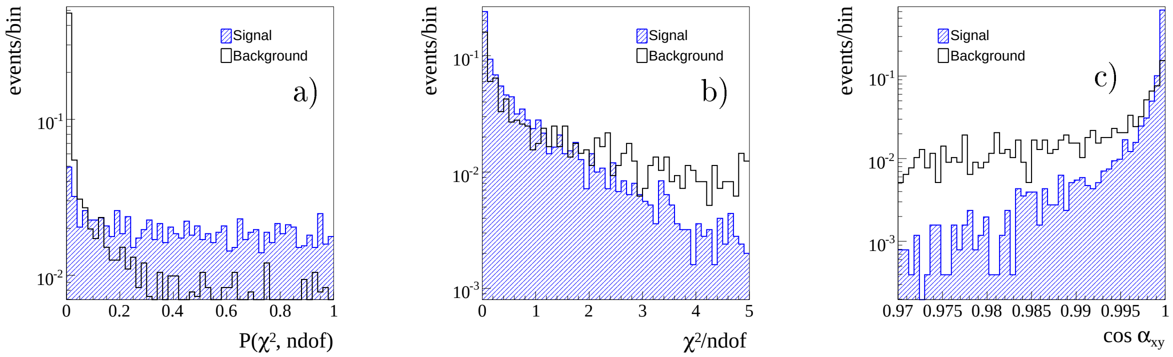

Figure C.3: Secondary vertex distributions: a) χ^2-probability of fit, b) χ^2/ndof of fit and c) cosine of the angle between the B_s^0 candidates flight direction and secondary vertex in the transverse plane.

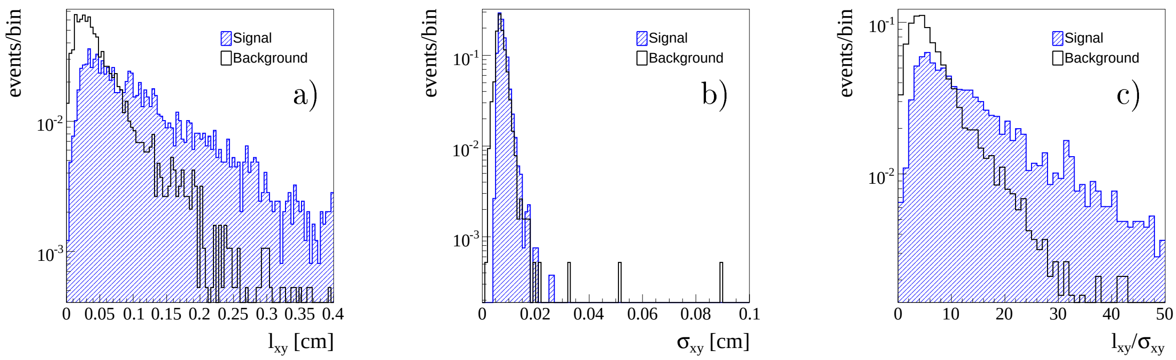

Figure C.4: Flight length distributions in the transverse plane: a) flight length, b) error on the flight length, c) flight length significance.

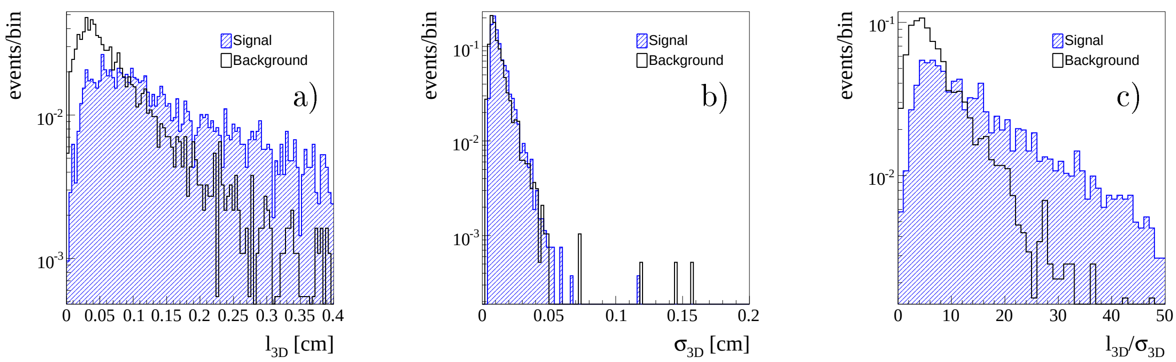

Figure C.5: Flight length distributions in three dimensions: a) flight length, b) error on the flight length, c) flight length significance.

Appendix D

Normalisation Distributions

Figures D.1-D.5 illustrate the distributions of different selection variables for the signal $B_s^0 \to \mu^+\mu^-$and normalisation $B^\pm \to J/\psi K^\pm$, after the HLT requirement and in the mass window $4.8 < m_{\mu\mu(K)} < 6.0\,\text{GeV}$. The histograms are normalised to unity.

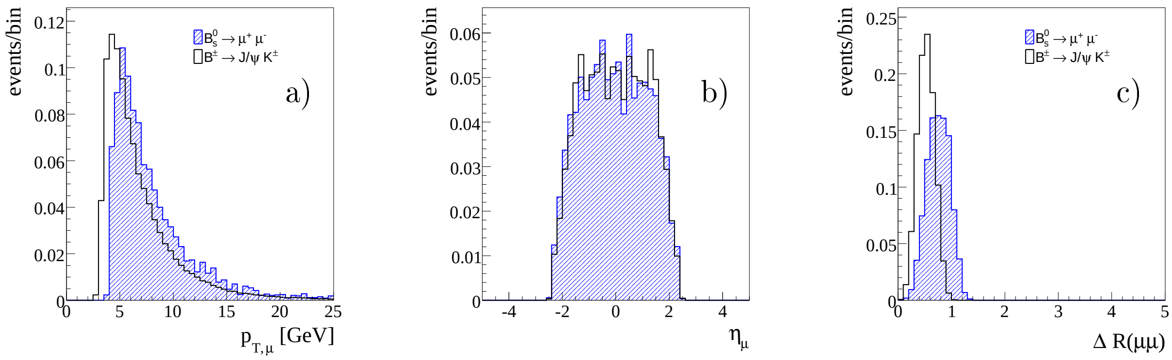

Figure D.1: Muon variable distributions: a) transverse momentum, b) pseudo-rapidity, c) $\eta\phi$ separation of the two muons.

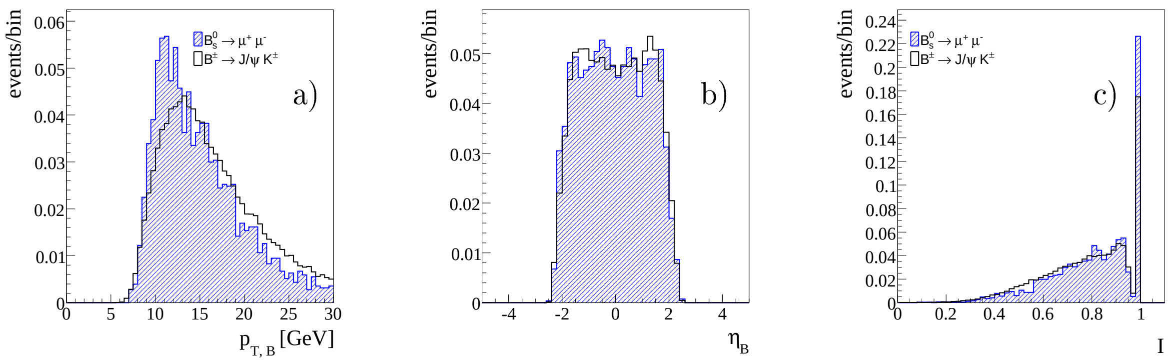

Figure D.2: Reconstructed B_s^0/B^+ candidates: a) transverse momentum, b) pseudo-rapidity, c) Isolation of the B_s^0/B^+ candidate.

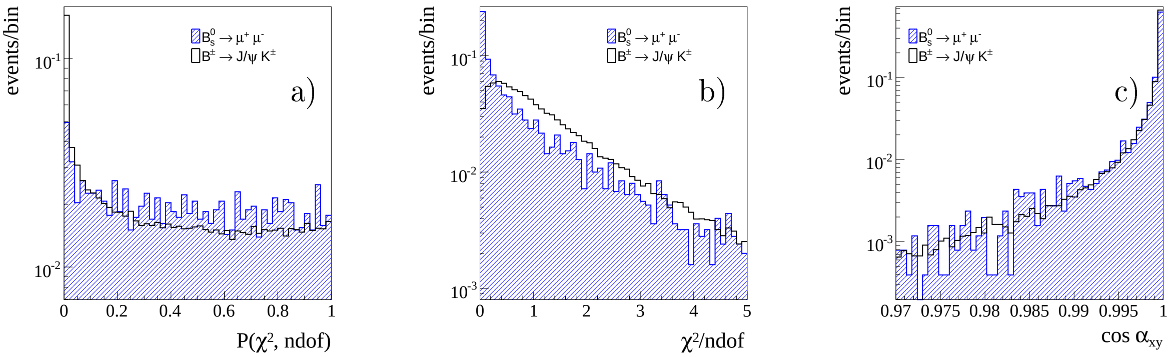

Figure D.3: Secondary vertex distributions: a) χ^2-probability of fit, b) χ^2/ndof of fit and c) cosine of the angle between the B_s^0/B^+ candidates flight direction and secondary vertex in the transverse plane.

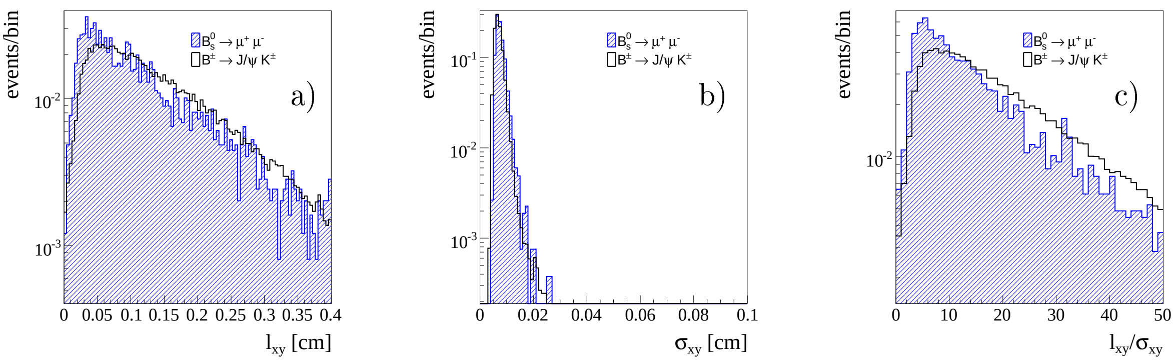

Figure D.4: Flight length distributions in the transverse plane: a) flight length, b) error on the flight length, c) flight length significance.

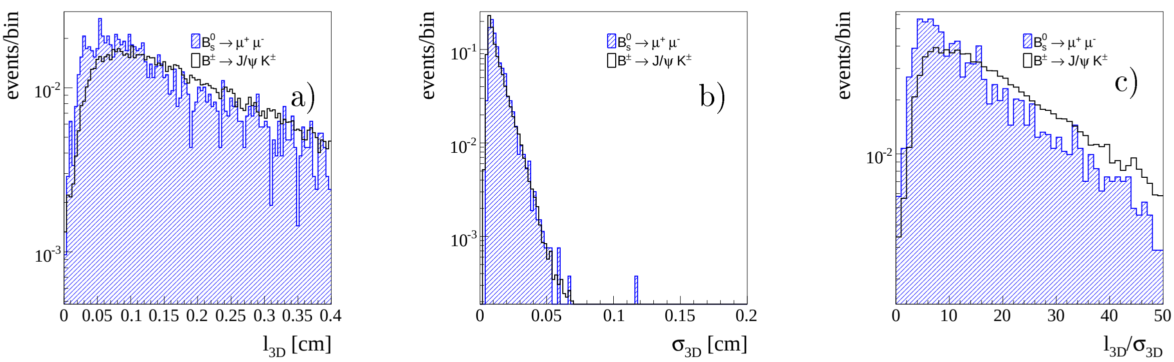

Figure D.5: Flight length distributions in three dimensions: a) flight length, b) error on the flight length, c) flight length significance.

Appendix E

Rare Background Distributions

Figures E.1- E.5 are absolutely normalised and illustrate the background distributions before the application of selection criteria (muon identification, in particular). The signal distribution is normalised to the same area as the background distribution.

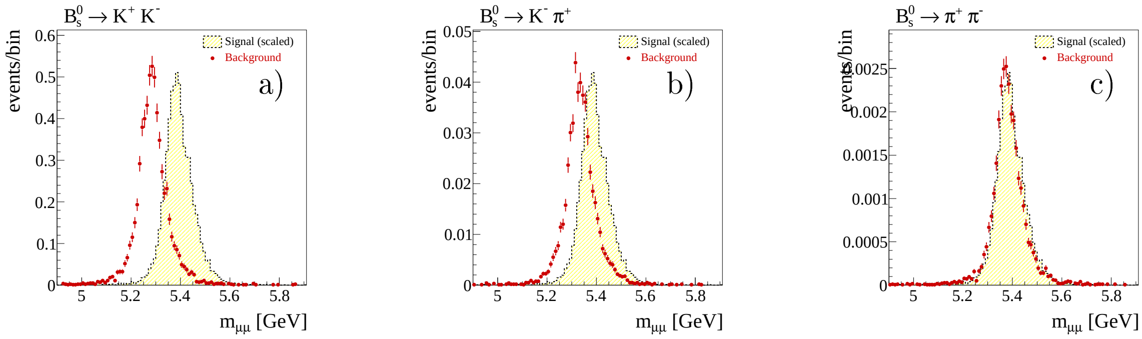

Figure E.1: Background $m_{\mu\mu}$ distributions before the application of selection criteria for different channels: a) $B_s \rightarrow K^+K^-$, b) $B_s \rightarrow K^+\pi^-$, c) $B_s \rightarrow \pi^+\pi$.

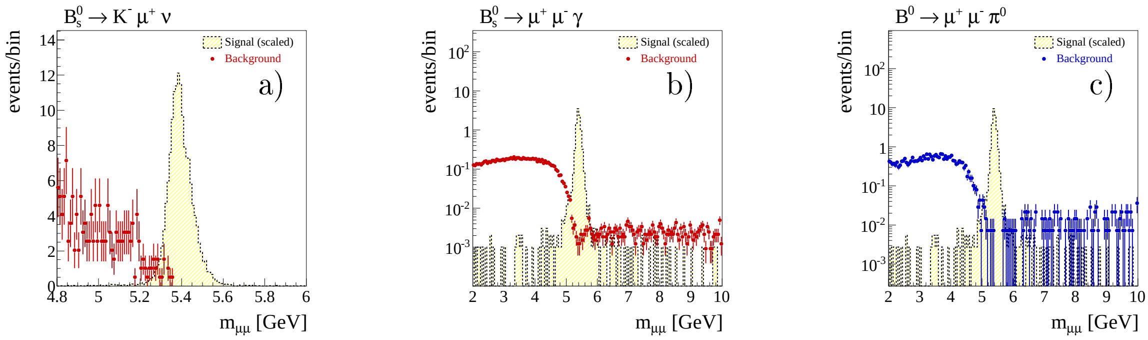

Figure E.2: Background $m_{\mu\mu}$ distributions before the application of selection criteria for different channels: a) $B_s \rightarrow K^-\mu^+\nu$, b) $B_s \rightarrow \mu^+\mu^-\gamma$, c) $B_d \rightarrow \pi^0\mu^+\mu^-$.

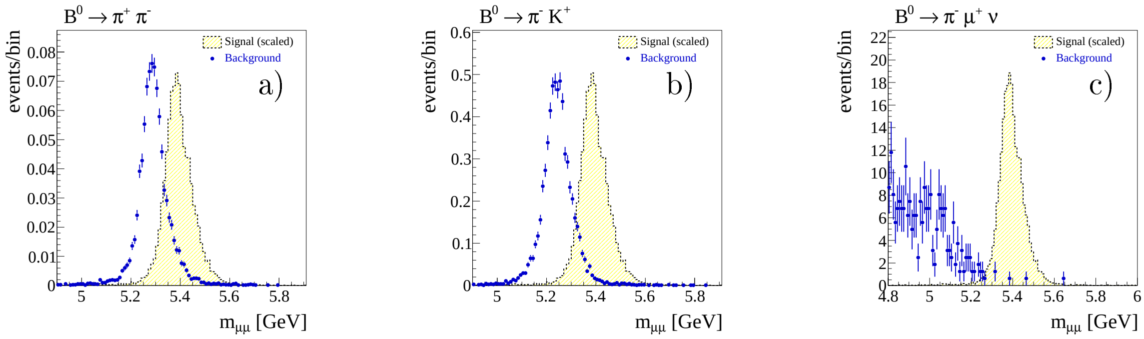

Figure E.3: Background $m_{\mu\mu}$ distributions before the application of selection criteria for different channels: a) $B_d \to \pi^+\pi^-$, b) $B_d \to K^+\pi^-$, c) $B_d \to \pi^-\mu^+\nu$.

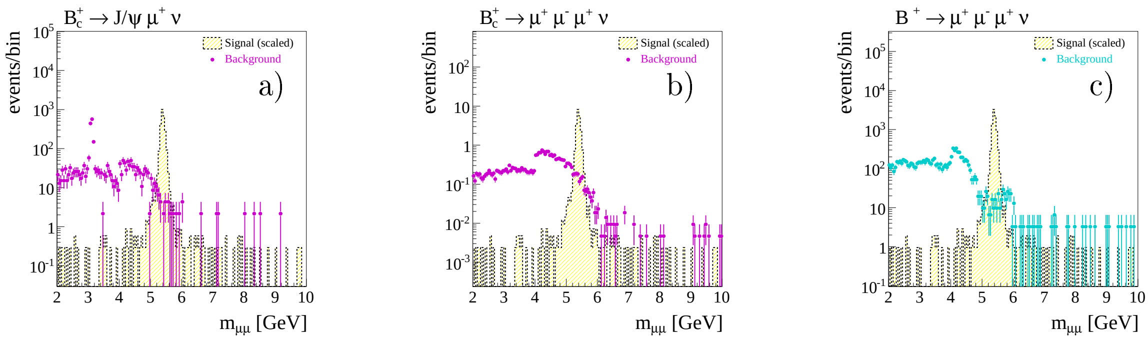

Figure E.4: Background $m_{\mu\mu}$ distribution before the application of selection criteria for different channels: a) $B_c^+ \to J/\psi(\to \mu^+\mu^-)\mu^+\nu_\mu$, b) $B_c^+ \to \mu^+\mu^-\mu^+\nu_\mu$, c) $B^+ \to \mu^+\mu^-\mu^+\nu_\mu$.

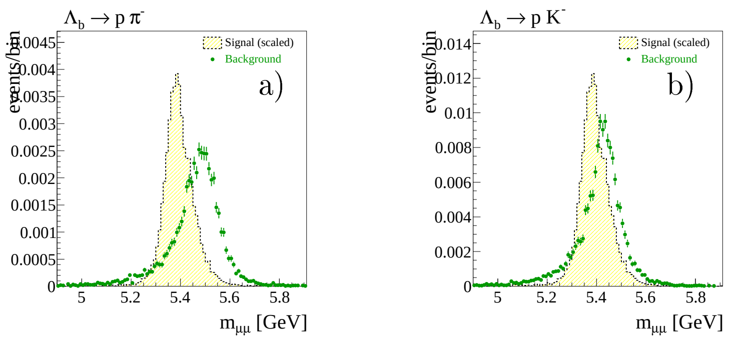

Figure E.5: Background $m_{\mu\mu}$ distribution before the application of selection criteria for different channels: a) $\Lambda_b \to p\pi^-$, b) $\Lambda_b \to pK^-$.

Bibliography

[1] O. Bruning *et al.*, "LHC design report. Vol. I: The LHC main ring", CERN-2004-003-V-1. 3

[2] L. Evans and P. Bryant eds, "LHC Machine", *JINST* **3** (2008) S08001. 3, 5

[3] J. Brau *et al.*, "International Linear Collider reference design report. 1: Executive summary. 2: Physics at the ILC. 3: Accelerator. 4: Detectors", ILC-REPORT-2007-001. 4, 10

[4] H. Braun *et al.*, "CLIC 2008 parameters", CERN-OPEN-2008-021. 4, 10

[5] R. Adolphi *et al.*, "The CMS experiment at the CERN LHC", *JINST* **3** (2008) S08004. 4, 15, 16, 19, 20, 26, 34

[6] "CMS, the Compact Muon Solenoid: Technical proposal", CERN-LHCC-94-38. 4

[7] "ATLAS: Detector and physics performance technical design report. Volume 1", CERN-LHCC-99-14. 4

[8] G. Aad *et al.*, "The ATLAS Experiment at the CERN Large Hadron Collider", *JINST* **3** (2008) S08003. 4

[9] S. Amato *et al.*, "LHCb technical proposal", CERN-LHCC-98-04. 4

[10] A. A. Alves *et al.*, "The LHCb Detector at the LHC", *JINST* **3** (2008) S08005. 4

[11] "ALICE: Technical proposal for a large ion collider experiment at the CERN LHC", CERN-LHCC-95-71. 4

[12] K. Aamodt *et al.*, "The ALICE experiment at the CERN LHC", *JINST* **0803** (2008) S08002. 4

[13] G. Anelli *et al.*, "The TOTEM experiment at the CERN Large Hadron Collider", *JINST* **3** (2008) S08007. 4, 8

[14] O. Adriani *et al.*, "Technnical proposal for the CERN LHCf experiment: Measurement of photons and neutral pions in the very forward region of LHC", CERN-LHCC-2005-032. 4

[15] O. Adriani *et al.*, "The LHCf detector at the CERN Large Hadron Collider", *JINST* **3** (2008) S08006. 4

[16] J. M. Campbell, J. W. Huston and W. J. Stirling, "Hard interactions of quarks and gluons: A primer for LHC physics", *Rept. Prog. Phys.* **70** (2007) 89. 4, 9

[17] O. S. Bruning *et al.*, "LHC luminosity and energy upgrade: A feasibility study", CERN-LHC-PROJECT-REPORT-626. 7

[18] M. Benedikt *et al.*, "Preliminary accelerator plans for maximizing the integrated LHC luminosity", CERN-AB-2006-018. 7

[19] R. Garoby, "Upgrade Issues for the CERN Accelerator Complex", *Conf. Proc.* **C0806233** (2008) 7

[20] D. Bortoletto, "The ATLAS and CMS Plans for the LHC Luminosity Upgrade", (2008) 8

[21] M. Deile *et al.*, "Diffraction and total cross-section at the Tevatron and the LHC", (2006) 8

[22] J. R. Cudell *et al.*, "Benchmarks for the forward observables at RHIC, the Tevatron-run II and the LHC", *Phys. Rev. Lett.* **89** (2002) 201801. 8

[23] J. D. Jackson, *Classical Electrodynamics*. Wiley (1999). 10

[24] C. Ankenbrandt *et al.*, "Muon Collider Task Force Report", (2007) FERMILAB-TM-2399-APC. 11

[25] G. Ambrosio *et al.*, "Design study for a staged very large hadron collider", SLAC-R-591. 11

[26] G. Acquistapace *et al.*, "CMS, the magnet project: Technical design report", CERN-LHCC-97-10. 15

[27] "CMS, tracker technical design report", CERN-LHCC-98-06. 15, 17

[28] "CMS: The electromagnetic calorimeter. Technical design report", CERN-LHCC-97-33. 22

[29] P. Adzic *et al.*, "Energy resolution of the barrel of the CMS electromagnetic calorimeter", *JINST* **2** (2007) P04004. 23

[30] "CMS: The hadron calorimeter technical design report", CERN-LHCC-97-31. 24

[31] E. Yazgan, "The CMS barrel calorimeter response to particle beams from 2-GeV/c to 350-GeV/c", CERN-CMS-CR-2008-040. 25

[32] G. L. Bayatian *et al.*, "CMS physics: Technical design report", CERN-LHCC-2006-001. 25, 27, 30

[33] "CMS, the Compact Muon Solenoid. Muon technical design report", CERN-LHCC-97-32. 26

[34] G. Abbiendi *et al.*, "Muon Reconstruction in the CMS Detector", CMS AN 2008/97. 28, 136, 139, 140

[35] R. Fruhwirth, "Application of Kalman filtering to track and vertex fitting", *Nucl. Instrum. Meth.* **A262** (1987) 444-450. 29

[36] C. Broennimann *et al.*, "Development of an indium bump bond process for silicon pixel detectors at PSI", *Nucl. Instrum. Meth.* **A565** (2006) 303-308. 35, 75

[37] P. Trueb, "CMS Pixel Module Qualification and Monte Carlo Study ", *Diss. ETH No. 17985* (2008) 41, 53, 56, 64, 78

[38] S. Dambach, "Optimization of CMS Pixel Barrel Detector Modules", *Diss. ETH No. 18203* (2008) 42, 57, 58, 59, 80, 87

[39] M. Stueckelberger, "Threshold Calibration of the CMS Pixel Detector", *ETHZ-IPP Internal Report 2009-03* (2008) 43

[40] A. Starodumov *et al.*, "Qualification procedures of the CMS pixel barrel modules", *Nucl. Instrum. Meth.* **A565** (2006) 67-72. 49, 64

[41] W. Erdmann *et al.*, "LHC-rate beam test of CMS pixel barrel modules", *Nucl. Instrum. Meth.* **A572** (2007) 57-58. 56

[42] H. Bichsel, "Straggling in thin Silicon Detectors", *Rev. Mod. Phys.* **60** (1988) 663-699. 56

[43] S. Konig, C. Hormann, R. Horisberger, S. Streuli and R. Weber, "Assembly of the CMS pixel barrel modules", *Nucl. Instrum. Meth.* **A565** (2006) 62-66. 75

[44] S. Cucciarelli and D. Kotlinski, "Pixel Hit Reconstruction", *CMS IN 2004/014* (2004) 87

[45] C. Amsler *et al.*, "Review of particle physics", *Phys. Lett.* **B667** (2008) 1. 111, 128

[46] A. Ceccucci, Z. Ligeti and Y. Sakai, "The CKM quark-mixing matrix", 113

[47] T. Bergfeld *et al.*, "Search for decays of B0 mesons into pairs of leptons: $B^0 \to e^+e^-$, $B^0 \to \mu^+\mu^-$, $B^0 \to e^\pm\mu^\mp$", *Phys. Rev.* **D62** (2000) 091102. 116, 122

[48] H. Albrecht *et al.*, "ARGUS: A Universal Detector at DORIS-II", *Nucl. Instrum. Meth.* **A275** (1989) 1-48. 116

[49] B. Aubert *et al.*, "Study of $B^0 \to \ell^+\ell^-$ decays ($l = e, \mu$)", *Phys. Rev.* **D77** (2008) 032007. 116, 122

[50] M. C. Chang *et al.*, "Search for $B^0 \to \ell^+\ell^-$ at Belle", *Phys. Rev.* **D68** (2003) 111101. 116, 122

[51] A. Drutskoy *et al.*, "Measurements of exclusive B/s0 decays at the Upsilon(5S)", *Phys. Rev.* **D76** (2007) 012002. 116

[52] E. Norrbin and T. Sjostrand, "Production and hadronization of heavy quarks", *Eur. Phys. J.* **C17** (2000) 137-161. 119, 194

[53] V. M. Abazov *et al.*, "Search for $B_s^0 \to \mu^+\mu^-$ at D0", *Phys. Rev.* **D76** (2007) 092001. 122, 127

[54] T. Aaltonen *et al.*, "Search for $B_s^0 \to \mu^+\mu^-$ and $B_d^0 \to \mu^+\mu^-$ Decays with $2\,\mathrm{fb}^{-1}$ of $p\bar{p}$ Collisions", *Phys. Rev. Lett.* **100** (2008) 101802. 122, 127

[55] M. Artuso *et al.*, "B, D and K decays", *Eur. Phys. J.* **C57** (2008) 309-492. 122, 124, 128

[56] A. Nikitenko, A. Starodumov and N. Stepanov, "Observability of $B^0_{s(d)} \to \mu^+\mu^-$ decay with the CMS detector", (1999) 122

[57] C. Eggel, U. Langenegger and A. Starodumov, "Discovery Potential for $B_s^0 \to \mu^+\mu^-$", CMS analysis note 2006/097. 122, 127, 153, 154

[58] G. Aad *et al.*, "Expected Performance of the ATLAS Experiment - Detector, Trigger and Physics", (2009) 122

[59] D. Martinez, J. A. Hernando and F. Teubert, "LHCb potential to measure / exclude the branching ratio of the decay $B_s^0 \to \mu^+\mu^-$", CERN-LHCB-2007-033. 122

[60] D. Martinez Santos, "$B_s^0 \to \mu^+\mu^-$ in LHCb", CERN-LHCB-2008-018. 122

[61] B. A. Campbell and P. J. O'Donnell, "Mass of the Top Quark and Induced Decay and Neutral Mixing of B Mesons", *Phys. Rev.* **D25** (1982) 1989. 123

[62] G. Buchalla and A. J. Buras, "QCD corrections to rare K and B decays for arbitrary top quark mass", *Nucl. Phys.* **B400** (1993) 225-239. 123

[63] G. Buchalla and A. J. Buras, "The rare decays K –> pi nu anti-nu, B –> X nu anti-nu and B –> l+ l-: An update", *Nucl. Phys.* **B548** (1999) 309-327. 123

[64] M. Misiak and J. Urban, "QCD corrections to FCNC decays mediated by Z-penguins and W-boxes", *Phys. Lett.* **B451** (1999) 161-169. 123

[65] W. Skiba and J. Kalinowski, "$B_s \to \tau^+\tau^-$ decay in a two Higgs doublet model", *Nucl. Phys.* **B404** (1993) 3-19. 124

[66] H. E. Logan and U. Nierste, "B/s,d –> l+ l- in a two-Higgs-doublet model", *Nucl. Phys.* **B586** (2000) 39-55. 125

[67] F. Borzumati and C. Greub, "Lessons from anti-B –> X/s gamma in two Higgs doublet models", (1998) 125

[68] E. Barberio *et al.*, "Averages of b-hadron and c-hadron Properties at the End of 2007", (2008) 125, 128

[69] K. S. Babu and C. F. Kolda, "Higgs-mediated $B_s^0 \to \mu^+\mu^-$ in minimal supersymmetry", *Phys. Rev. Lett.* **84** (2000) 228-231. 126

[70] S. R. Choudhury and N. Gaur, "Dileptonic decay of B/s meson in SUSY models with large tan(beta)", *Phys. Lett.* **B451** (1999) 86-92. 126

[71] C. Bobeth, T. Ewerth, F. Kruger and J. Urban, "Enhancement of $\mathcal{B}(B_d^0 \to \mu^+\mu^-)/\mathcal{B}(B_s^0 \to \mu^+\mu^-)$ in the MSSM with minimal flavour violation and large $\tan\beta$", *Phys. Rev.* **D66** (2002) 074021. 126

[72] S. Davidson, D. C. Bailey and B. A. Campbell, "Model independent constraints on leptoquarks from rare processes", *Z. Phys.* **C61** (1994) 613-644. 126

[73] D. P. Roy, "Squark and gluino mass limits in R-parity violating SUSY model from the Tevatron dilepton data", *Phys. Lett.* **B283** (1992) 270-277. 126

[74] G. L. Kane, C. Kolda and J. E. Lennon, "$B_s^0 \to \mu^+\mu^-$ as a probe of $\tan\beta$ at the Tevatron", (2003) 126

[75] S. Baek, "A two-loop contribution to $B_s^0 \to \mu^+\mu^-$ at large $\tan\beta$ in the MSSM", *Phys. Lett.* **B595** (2004) 461-468. 126

[76] A. Dedes and B. T. Huffman, "Bounding the MSSM Higgs sector from above with the Tevatron's $B_s^0 \to \mu^+\mu^-$", *Phys. Lett.* **B600** (2004) 261-269. 126

[77] F. Zwirner, "The quest for low-energy supersymmetry and the role of high-energy e+ e- colliders", (1991) 126

[78] T. Sjostrand, S. Mrenna and P. Skands, "PYTHIA 6.4 physics and manual", *JHEP* **05** (2006) 026. 127

[79] The CMS Offline WorkBook, see https://twiki.cern.ch/twiki/bin/view/CMS/WorkBook. 127

[80] T. Sjostrand, L. Lonnblad and S. Mrenna, "PYTHIA 6.2: Physics and manual", (2001) 127

[81] T. Aaltonen *et al.*, "Observation of New Charmless Decays of Bottom Hadrons", (2008) 128

[82] D. Melikhov and N. Nikitin, "Rare radiative leptonic decays $B_{s(d)}^0 \to \ell^+\ell^-\gamma$ ", *Phys. Rev.* **D70** (2004) 114028. 128

[83] N. V. Nikitin, S. Y. Sivoklokov, L. N. Smirnova, D. A. Tlisov and K. S. Toms, "Possibility of recording rare muonic decays of B mesons at the ATLAS detector in the case of LHC operation in the initial-luminosity mode", *Phys. Atom. Nucl.* **70** (2007) 2086-2102. 128

[84] V. V. Kiselev, "Exclusive decays and lifetime of B_c meson in QCD sum rules", (2002) 128

[85] C. Eggel, U. Langenegger and A. Starodumov, "Study of $B_s^0 \to \mu^+\mu^-$ in CMS", CMS analysis note 2009/xxx. 130, 148, 158

[86] CSA07 Physics Page, see https://twiki.cern.ch/twiki/bin/view/CMS/CSA07Physics. 130

[87] CMSSW CVS repository (CSA07), see http://cmssw.cvs.cern.ch/cgi-bin/cmssw.cgi/CMSSW/Configuration/CSA07Production/data/. 131

[88] "CMS High Level Trigger", CERN-LHCC-2007-021. 134, 135

[89] S. Cucciarelli, M. Konecki, D. Kotlinski and T. Todorov, "Track reconstruction, primary vertex finding and seed generation with the pixel detector", CERN-CMS-NOTE-2006-026. 135

[90] B. Stieger, "Study of the $b \to \mu D^0 X$ Production Cross-section at the CMS Experiment", *ETHZ-IPP Internal Report 2008-14* (2008) 143

[91] M. Reinhard, "Determination of the Muon Reconstruction and Indentification Efficiencies in the CMS Experiment", *ETHZ-IPP Internal Report 2008-12* (2008) 143, 144

[92] P. Valnaer, L. Barbone, N. D. Filippis, T. Speer, O. Buchmuller and F.-P. Schilling, "Impact of CMS Silicon Tracker Misalignment on Track and Vertex Reconstruction", CMS-NOTE-2006-029. 145

[93] S. Bityukov, see http://cmsdoc.cern.ch/∼bityukov/. 163

[94] J. Heinrich *et al.*, "Interval estimation in the presence of nuisance parameters. 1. Bayesian approach", (2004) 163, 164

[95] P. Reznicek, "LHC — very rare B decays program", talk presented at BEACH 2006. 168

Acronyms

2HDM Type-II Two-Higgs-Doublet Model

ACS 400 MHz Acceleration Systems

ADC Analog to Digital Converter

ADT Transverse Damping and Feedback System

BSM Beyond Standard Model

CERN Conseil Européen pour la Recherche Nucléaire

CKM Cabibbo-Kobayashi-Maskawa

CLIC Compact Linear Collider

CMS Compact Muon Solenoid

CMSSW CMS Software

CSC Cathode Strip Chambers

DAC Digital to Analog Converter

DT Drift Tubes

EB Barrel Electromagnetic Calorimeter

ECAL Electromagnetic Calorimeter

EE Endcap Electromagnetic Calorimeter

ES Preshower

FCNC Flavour Changing Neutral Currents

FPGA Field Programmable Gate Array

HB Barrel Hadronic Calorimeter

HCAL Hadronic Calorimeter

HDI High Density Interconnect

HE Endcap Hadronic Calorimeter

HF Forward Hadronic Calorimeter

HLT High Level Trigger

HO Outer Hadronic Calorimeter

ILC International Linear Collider

IP Interaction Point

IR Insertion Region

LHC Large Hadron Collider

LO Leading Order

MC Monte Carlo

MCTF Muon Collider Task Force

MIP Minimum Ionising Particle

MSSM Minimal Supersymmetric Standard Model

MTC Marlon Trigger Chip

NLO Next to Leading Order

NP New Physics

NSD Non-Single Diffractive Inelastic Interactions

PH Pulse Height

PSI Paul Scherrer Institute

PUC Pixel Unit Cell

QCD Quantum Chromodynamics

ROC Read-Out Chip

RPC Resistive Plate Chambers

SC Superconducting Cavities

SLHC Super Large Hadron Collider

SM Standard Model

SSB Spontaneous Symmetry Breaking

TBM Token Bit Manager

TEC Tracker EndCaps

TIB Tracker Inner Barrel

TID Tracker Inner Disks

TNP 'Tag and Probe'

TOB Tracker Outer Barrel

UBL Ultra Black Level

VLHC Very Large Hadron Collider

List of Figures

List of Tables

Acknowledgements

It is only through the expertise and fantastic support of so many colleagues and friends that this work could be completed. I would like to express my deep gratitude towards Urs Langenegger for giving me this opportunity to enter the field of particle physics in such a unique and exciting experiment. Throughout the past three years, his expert guidance, motivation, patience and continuous support have been truly invaluable. I would like to extend a special thanks to Andrey Starodumov, who provided vital input to both major sections of this work, and to Sarah Dambach and Peter Trüb for their valuable contributions. I wish to thank everyone in the group for the excellent team work and their friendship.

Many thanks also to my co-examiner Felicitas Pauss for carefully reading the manuscript and providing many constructive comments.

I would like to thank the entire PSI pixel group for their help and support. In particular, Roland Horisberger for the many instructive discussions regarding the CMS pixel system, Tilman Rohe for his help in the sensor related part of the module qualification and Silvan Streuli for his help in the operation and maintenance of the cooling box. Many thanks also to Kurt Gabathuler, Danek Kotlinski, Hans-Christian Kästli, Willy Bertl, Beat Meier, Stefan König and Wolfram Erdmann. It has been a great pleasure and privilege to have worked within such a knowledgeable and experienced group, with such a friendly and warm atmosphere.

This work also owes a significant debt to Derek Feichtinger, whose excellent support on LHC grid computing is deeply appreciated.

Finally I thank my friends and family for their love and encouragement as always. To my loving partner, companion and greatest ally Milan - what about them apples?

Curriculum Vitae

Personal Details

Name:	Christina Eggel	**Date of Birth:**	4. November 1978
Address:	Höhtalstrasse 23	**Place of Birth:**	Naters (VS)
	5408 Ennetbaden	**Nationality:**	Swiss

Qualifications

Jan 2009 **Dr. sc. ETH Zurich**
Swiss Federal Institute of Technology, Zurich

Oct 2005 **Dipl. Phys. ETH**
Swiss Federal Institute of Technology, Zurich

Jun 1998 **Matura Typus C**
Kollegium "Spiritus Sanctus", Brig

Education

2006 - 2009 **Graduate education in experimental high energy physics**
Swiss Federal Institute of Technology, Zurich
Paul Scherrer Institute, Villigen

Thesis title: "CMS Pixel Module Qualification and Search for $B_s^0 \rightarrow \mu^+\mu^-$"

1998 - 2005 **Physics, specialisation in Particle Physics**
Swiss Federal Institute of Technology, Zurich

Diploma thesis title: "Response Study of the Small Low Energy Particle Monitor for Space Application"

Internships and Work Experience

May 2005	**Paul Scherrer Institute** Villigen, Switzerland	**Semester project** Measurement of the punch through bias grid in the CMS pixel sensors
2000 - 2004	**Paul Scherrer Institute** Villigen, Switzerland	**Group member, Laboratory of Astrophysics** Research within Space Technology group