Studies of radiation damage in silicon sensors and a measurement of the inelastic proton–proton cross-section at 13 TeV

Miguel Arratia
of Christ's College

A dissertation submitted to the University of Cambridge
for the degree of Doctor of Philosophy
September 2016

Studies of radiation damage in silicon sensors and a measurement of the inelastic proton–proton cross-section at 13 TeV

This thesis presents studies of radiation damage in silicon sensors for the new ATLAS tracker at the high-luminosity LHC, calibrations of the LHC luminosity scale, and a measurement of the proton–proton inelastic cross-section at 13 TeV with ATLAS data.

The studies of radiation damage are performed by comparing sensor performance before and after irradiation, and include annealing studies. The measured quantities include: leakage current, depletion depth, inter-strip isolation, and charge collection. Surface and bulk damage is studied by comparing the results of sensors irradiated with protons and neutrons. The observed degradation of performance suggests the current sensor design will endure the radiation damage expected over the lifetime of the experiment at the high-luminosity LHC.

The luminosity is calibrated for the proton–proton, proton–lead, and lead–lead collisions delivered by the LHC during 2013 and 2015. The absolute luminosity scale is derived with the van der Meer method. The systematic uncertainties on these measurements are 1.9%, 2.7%, and 6.0% respectively.

The measurement of the proton–proton inelastic cross-section at 13 TeV is performed with 60.1 μb^{-1} of data and an array of plastic scintillators. The luminosity calibration is one of the dominant sources of systematic uncertainty. The result is 79.3 ± 2.9 mb. The value is 11% higher than at 7 TeV and agrees well with theoretical predictions.

Declaration

This dissertation is the result of my own work, except where explicit reference is made to the work of others, and has not been submitted for another qualification to this or any other university. This dissertation does not exceed the word limit for the respective Degree Committee.

Miguel Arratia

Acknowledgements

I would like to thank Dr. Bart Hommels and Dr. Patricia Ward for their excellent guidance and seemingly unbounded patience. In addition, I thank the "Maestro" Saevar Sigurdsson for sharing some of the infinite tricks he knows.

I also thank my ATLAS colleagues that made enjoyable my experience in this gargantuan collaboration. In particular Witold Kozanecki for teaching me proper rigour and the real meaning of a per-mil level systematic uncertainty; and to Brad Axen, Zachary Marshall, Christian Ohm, Lauren Tompkins and Beate Heinemann for their congenial work style and for all that arises from working with people cleverer than oneself.

I thank the generous support from Cambridge International Trust and CONICYT Becas Chile 72140349.

This work, and everything else in my life, could not have been possible without the love of my beloved wife Jacqueline.

Preface

This thesis describes my research that relates to the ATLAS experiment at the LHC accelerator in CERN. My work spans studies of radiation damage on silicon-strip sensors, luminosity calibrations, and a measurement of the inelastic proton–proton cross-section. Additionally, I have participated in the operation of ATLAS and LHC by taking shifts in the ATLAS control room during physics running and in the CERN main control room during van der Meer scans.

My hardware-related work was guided by my supervisor Dr. Bart Hommels and helped by Saevar Sigurdsson. Chapter 2 describes measurements of large-area prototype silicon-strip sensors that I performed in the HEP group clean-room facility with an existing setup. Chapter 3 describes measurements on irradiated sensors for which I worked on the experimental setup from scratch.

These studies are part of a coordinated effort with several institutes across the globe that aims to study the radiation-hard technology needed for the upgrade of the ATLAS experiment at the high-luminosity LHC. My work contributed to the results published in Refs. [1, 2], which demonstrate that the prototypes tested comply with the required specifications and are resilient enough to endure the radiation damage expected over the lifetime of the experiment.

My work on luminosity started with my "authorship qualification" task that was the calibration of the luminosity for the 2013 proton–lead run. I led this analysis from the beginning, i.e. from the data taking at the CERN control room. I collaborated with a group of three people that produced studies of length scale calibration, non-factorization biases, and luminosity stability. All other studies presented in this thesis are my work. This was the first ATLAS luminosity calibration with nuclear beams. The final systematic uncertainty is 2.7% and is one of the most accurate measurements of its kind at the LHC [3].

Subsequently, I worked on the calibration of the 2015 proton–proton run in collaboration with the ATLAS luminosity group. I studied and implemented some of the major corrections and systematic uncertainties such as those associated with beam-orbit drift, beam-beam effects, and fit modelling [4]. The systematic uncertainty of this calibration is 1.9% [5]; this is one of the most accurate ATLAS results to date.

Finally, I performed the 2015 lead–lead calibration and started as a shifter in the van der Meer runs. The preliminary systematic uncertainty of this calibration is 6%, and is comparable with the best of its kind (by ALICE, 5.8%). This is the first ATLAS result of its kind [6].

My work on the measurement of the inelastic proton–proton cross-section, described in chapter 5, was guided by my co-supervisor Dr. Patricia Ward. I collaborated with an analysis team of six members. The measurements of the trigger efficiency and background estimation are my work, as well as the luminosity calibration described above. The latter two are the two main sources of systematic uncertainty. This measurement is the first performed with 13 TeV data and its total uncertainty of 3.7% represents a significant improvement with respect to a similar ATLAS measurement with 7 TeV data that has a total uncertainty of 11%. This result is published in Ref. [7].

Contents

Chapter 1

The LHC and the ATLAS experiment

1.1 The large hadron collider

The large hadron collider (LHC) [8] is a circular particle accelerator of 26.7 km circumference located at the CERN laboratory near Geneva. It is part of an accelerator complex that is illustrated in figure 1.1. The LHC is divided into eight sections, each consisting of an arc and a straight section of about 530 m. The LHC ring hosts two beam pipes where particle beams circulate in opposite directions and are brought into collision in four of the straight sections where experiments (summarized in table 1.1) are located.

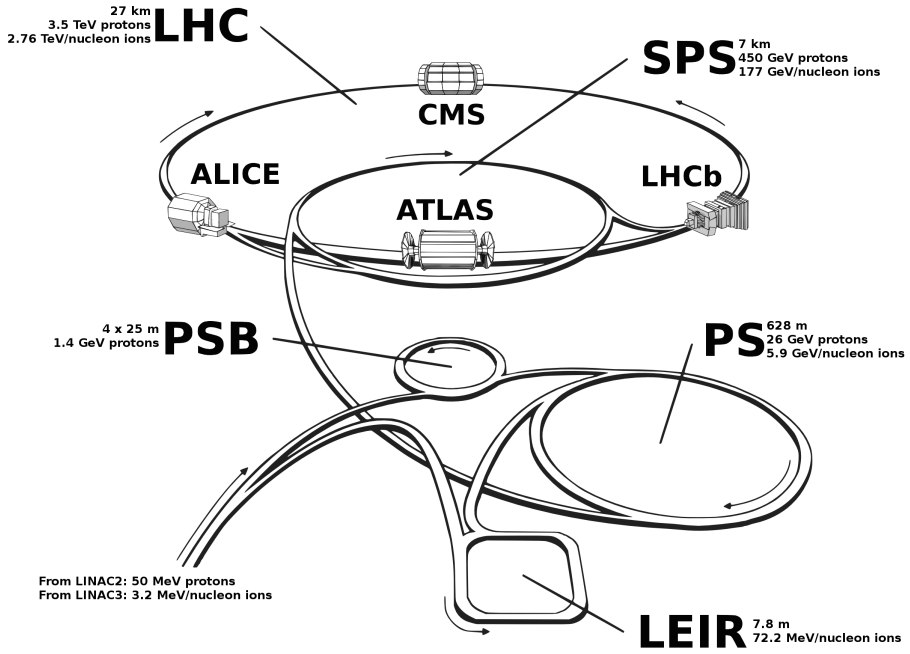

Figure 1.1: Illustration of the CERN accelerator complex and experiments. Source: Ref. [9].

Table 1.1: Summary of the four main experiments at the CERN LHC.

Experiment	Type	Main goal is to study:
ATLAS	General purpose	Higgs and beyond the standard model physics.
CMS	General purpose	Higgs and beyond the standard model physics.
LHCb	Dedicated	CP-violation.
ALICE	Dedicated	Quark-gluon plasma.

The accelerator complex includes a linear accelerator and synchrotrons; these successively accelerate the particle beams to an energy of 450 GeV. The injection to the LHC occurs in two of its straight sections. The LHC can increase the energy of the beams to a maximum of 7 TeV.

The acceleration is achieved in one of the straight sections with electric fields produced by a series of superconducting radio-frequency cavities operating at a frequency of 400 MHz. Given that the revolution frequency of the accelerator is 11.245 kHz ($c/26.7$ km), this produces 35640 radio-frequency "buckets" separated by 2.5 ns.

By design, only one tenth of these buckets are filled with particles (called bunches), therefore the minimum separation between filled bunches is 25 ns. While in principle the LHC can store up to 3564 bunches in each beam, operational limitations during beam dumps constrain this number to a maximum of 2808.

The LHC magnet system consists of more than 9300 superconducting magnets of different types cooled by a cryogenic system based on super-fluid helium at 1.9 K. Dipole magnets are used to bend the beams, quadrupole magnets to focus the beams[1] and multipoles to correct beam orbit distortions.

Each of the LHC arcs contains 152 dipole magnets of 14.3 m length, each capable of a peak field of 8.33 T, limiting the maximum storable beam energy to 7 TeV. Figure 1.2 shows a cross-section of these dipole magnets. Their "two-in-one" design results in equal magnetic fields in opposite directions that bend the counter-rotating particles. This is a limitation when colliding different species (e.g. protons and lead nuclei). This is discussed in section 1.1.2.

[1] While a quadrupole magnetic field focuses in one direction but defocuses in the orthogonal direction, two quadrupole magnets can yield a net focussing effect in both directions if their focusing directions are shifted by 90°. This focusing principle is called strong focusing.

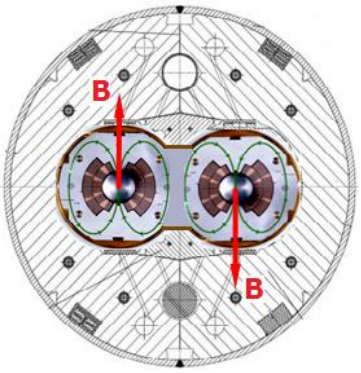

Figure 1.2: Illustration of the cross-section of a LHC dipole magnet. The arrows represent the magnetic field configuration on the beam pipes. Source: Ref. [10].

Quadrupole and multipole magnets are located both in the LHC arcs and in the straight sections. The maximum beam focusing is achieved with "inner triplets" that are dedicated quadrupole magnets close to the experiments; this is done to minimize the overlap of the colliding beams to maximize the luminosity.

The performance of the LHC magnet system and its injection chain can be characterized by the delivered bunch emittance and beta function; these concepts are introduced in the following sections based on a description presented in Ref. [11]. The concept of luminosity is described in more detail in chapter 4.

1.1.1 The concepts of emittance and beta function

Emittance

Emittance is a measure of the volume in position-momentum space of the particles in a bunch. It is an intrinsic quantity that characterizes a bunch because unlike spatial dimensions it does not vary with location in the accelerator[2]. This is a direct consequence of Liouville's theorem that holds in the absence of dissipative forces.

Normally the bunch emittance is defined with only a fraction of its particles; for example, the emittance in one direction can be defined by taking the projection of the three-dimensional particle density in that direction, which typically yields a Gaussian function, and then calculating the area of the phase-space defined with the particles that lie within one standard deviation of the mean position.

[2]The emittance depends on the energy of the beams, so this statement applies when the beams have reached their maximum energy.

The bunch emittance depends on the details of the formation of the bunch, and its transport along the injection chain. It can grow due to non-ideal effects that include particle losses due to intra-beam scattering and imperfections in the injection from one accelerator in the complex to the next. The latter was one of the dominant sources of emittance growth at the LHC during 2012, as reported in Ref. [12].

In electron–positron colliders the process of bremsstrahlung is exploited to reduce and stabilize the transverse[3] emittance of all bunches to a common value. This is impractical in hadron colliders because the bremsstrahlung effect is strongly suppressed; this leads to significant bunch-to-bunch variations of emittance that depend on the specific history of each bunch.

Typically the emittance is multiplied by the Lorentz factor of the beam ($\gamma = 7641$ for 7 TeV proton beams) to define the "normalized emittance", which has the advantage of being invariant with the energy of the beam. The LHC design normalized transverse-emittance is 3.75 $\mu\text{m} \cdot \text{rad}$, but values as small as ≈ 2.2 $\mu\text{m} \cdot \text{rad}$ have been reached.

Emittance is typically measured (at the LHC injection step) via a measurement of the transverse beam size with the wire-scan method (described in chapter 4) and a detailed knowledge of the magnetic field configuration of the accelerator, characterized by the beta function described in the following section.

Beta-function

Ideally the trajectory of each particle in the accelerator would close after each revolution; however, misalignment of the machine and other effects lead to trajectories that do not close but oscillate around the ideal in all coordinates. These are called "betatron oscillations".

The amplitude of the betatron oscillations is given by the product of the beam emittance and the so called beta function, $\beta(s)$, that characterizes the magnetic field (and thus focusing) at a given position in the beam trajectory, s. For complex accelerator systems such as the LHC, the beta function is calculated numerically using dedicated software for accelerator simulation [13].

[3] The bunch emittance in the horizontal or vertical direction in the plane transverse to the direction of the beam, called transverse emittance ϵ_{T}, dictates the luminosity of the collider, as will be discussed in more detail in chapter 4.

The more relevant beta functions are those that characterize the betatron oscillations in the transverse plane: $\beta_x(s)$ and $\beta_y(s)$. Typically these are very similar and are denoted by $\beta_{\mathrm{T}}(s)$. The envelope of the particle trajectories in a given bunch, i.e. the bunch size, is related to the transverse emittance and the transverse beta function as:

$$\sigma_{\mathrm{b}}(s) = \sqrt{\epsilon_{\mathrm{T}}\beta_{\mathrm{T}}(s)}. \tag{1.1}$$

The bunch size is minimized to optimize the luminosity. The beta function at the interaction regions[4] is called β^*. The LHC design[5] β^* is 0.55 m, which together with a normalized emittance of 3.75 $\mu\mathrm{m}\cdot\mathrm{rad}$ leads to a bunch size of 14 μm.

1.1.2 The LHC as a nucleus collider

The main goal of the LHC is to collide protons with a a center of mass energy of 14 TeV (at the time of writing 13 TeV has been the maximum achieved). But, the accelerator complex is perfectly capable of colliding nuclear beams as well [14].

Lead-208 nuclei have been used so far. As with protons they are accelerated in several steps. The nuclear accelerator chain (shown in figure 1.1) is only slightly different than the main one and includes a dedicated linear accelerator called Linac 3, and a synchrotron called low energy ion ring, or LEIR.

In the first-stage acceleration the nuclei are fully ionized and accelerated up to 72 MeV. Subsequently, they are injected into the main accelerator chain with an intensity that is about an order of magnitude less than that of proton beams (about 10^8 ions per bunch) and with larger bunch spacing (about 100 ns) due to limitations of the injection process and to meet requirements from experiments.

In 2015 the LHC delivered lead–lead collisions with a center-of-mass energy of 1 PeV (i.e. 2.5 TeV per nucleon per beam), setting a new record [15] by surpassing by a factor of two the energy achieved in previous LHC runs and by a factor of 25 the energy reached by RHIC at the Brookhaven National Laboratory, USA.

The LHC can also collide beams of different species, for example proton and lead beams. The proton–lead collisions were not initially foreseen, as the two-in-one magnet design

[4] The conservation of emittance leads to an increase in the bunch angular dispersion when the bunch volume is squeezed. This is why the highest bunch squeezing (low β) is achieved only in the interaction regions.

[5] For the high-luminosity experiments ATLAS and CMS.

imposes equal beam rigidity (momentum per proton) leading to different revolution frequencies as the beams have different speeds. This would prevent the delivery of collisions to the experiments because the interaction regions would move. At energies of about 2.5 TeV per nucleon, however, the difference in frequency is small enough that it can be eliminated by distorting the beam orbits within an acceptable range [16].

The beam rigidity restriction leads to an effective energy asymmetry in the collisions. For example, in 2013 the LHC delivered collisions between 4 TeV proton beams and 328 TeV lead beams (4 TeV per proton). With each lead nucleus containing 126 neutrons, the total energy is equivalent to 1.6 TeV per nucleon. The collision of a 4 TeV proton and a 1.6 TeV nucleon is boosted with respect to the laboratory frame by about half unit of rapidity in the proton-going direction.

The resulting center-of-mass energy of 5 TeV is an increase of about 25 times the energy from analogous collisions (proton–gold at RHIC). This is one of the largest energy increases achieved in the history of particle accelerators. This energy provides access for the first time to the production of W and Z bosons and t quarks in this type of reaction [16].

The performance of the LHC delivering proton–proton, proton–lead and lead–lead collisions is summarized in table 1.2.

Table 1.2: Summary of the achieved performance of the LHC delivering proton–proton, proton–lead and lead–lead collisions.

	proton–proton	proton–lead	lead–lead
Beam energy (per nucleon) [TeV]	6.5	4, 1.6	2.5
β^* [m]	0.4	19	0.8
Maximum colliding bunches	2232	40	492
Bunch [10^{11} protons]	≈ 1.2	≈ 0.1	≈ 0.1
Minimum bunch spacing [ns]	25	250	100
Peak luminosity [$\mathrm{cm}^{-2}\mathrm{s}^{-1}$]	10^{34}	3.3×10^{27}	2.7×10^{27}

1.2 The ATLAS detector

ATLAS [17] is a detector capable of measuring a large range of different physics processes. Its dimensions are 46 m long and 25 m high. It covers about 2π in the azimuthal angle[6] and is symmetrical in the forward-backward directions. An illustration of the ATLAS detector highlighting its various components is shown in figure 1.3.

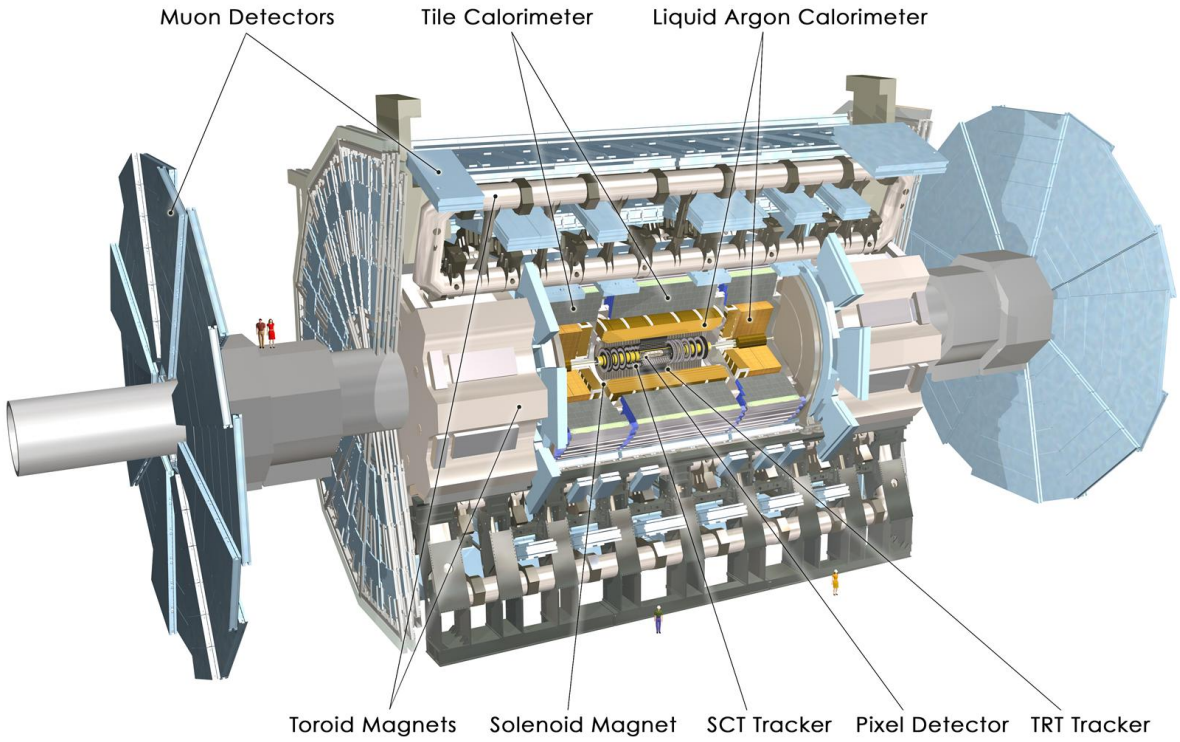

Figure 1.3: Illustration of the ATLAS detector. The main parts of the inner detectors, calorimeters, magnets, and muon spectrometer are highlighted. Source: CERN.

The conceptual design of ATLAS is based on the collider experiment tradition: its inner detector immersed in a solenoidal magnetic field provides measurements of momentum of charged particles and primary vertices; calorimeters provide energy measurements; and its muon system immersed in a toroidal magnetic field provides momentum measurements. Each sub-detector is described in the following sections.

[6] The convention for the coordinate system adopted by the ATLAS collaboration is detailed in Ref. [17]: "ATLAS uses a right-handed coordinate system with its origin at the nominal interaction point in the center of the detector and the z-axis along the beam pipe. The x-axis points from the interaction point (IP) to the center of the LHC ring, and the y-axis points upward. Cylindrical coordinates (r, ϕ) are used in the transverse plane, ϕ being the azimuthal angle around the beam pipe. The pseudorapidity is defined in terms of the polar angle θ as $\eta = -\ln\tan(\theta/2)$."

1.2.1 Inner detector

The main goal of the inner detector is to reconstruct the trajectories of charged particles and measure their charge and momentum. Additionally, reconstruction of vertices is needed to disentangle multiple interactions in a single bunch crossing (about 25 at nominal luminosity), and to identify secondary particles via displaced vertices.

The inner detector occupies a cylindrical volume around the center of the interaction region. It has a diameter of 2.1 m, a length of 6.2 m, and covers $|\theta| > 9.4^\circ$ ($|\eta| < 2.5$). It is a hybrid detector using two main technologies: silicon detectors in the smaller radii and gas detectors in the larger radii. The granularity increases with decreasing radius to provide precise vertex and impact parameter measurements. All these systems are immersed in a 2 T axial magnetic field produced by a superconducting magnet.

As illustrated in figure 1.4, the inner detector is arranged in concentrical cylindrical layers at the central section (barrel) and in disks perpendicular to the beam in the forward section (end-cap). The innermost part ($31 < r < 122$ mm) corresponds to the pixel detector system; it is surrounded ($299 < r < 554$ mm) by the semiconductor tracker (SCT), that is made of micro-strip detectors. The third sub-detector ($554 < r < 1070$ mm) is the transition radiation tracker (TRT) that is based on gas-filled drift tubes.

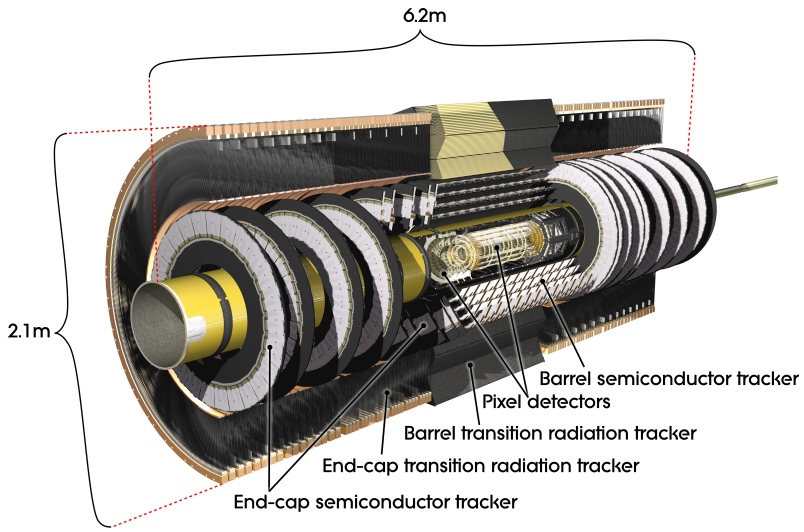

Figure 1.4: Illustration of the ATLAS inner detector. The principal sub-detectors are highlighted. Source: CERN.

The pixel system consists of four barrel layers and three end-cap disks on each side with a total of 1774 modules and approximately 80 million 50×400 μm^2 channels. The

SCT consists of four barrel layers and nine end-cap disks on each side with a total of 4088 modules and approximately 6.2 million 80 μm $\times$ 12 cm channels. The TRT consists of about 3×10^5 drift tubes of 4 mm diameter and 0.4–1.5 m length filled with a gas mixture containing xenon and carbon dioxide.

The resulting signals are encoded using time over threshold for the pixel system and with binary readout for the SCT and TRT. The latter was chosen to reduce data volume and has no significant effect on the tracking performance. On average a charged particle produces 4 hits in the pixel system, 4 in the SCT, and 30 in the TRT.

The silicon part of the inner detector has a resolution of about 15 μm in the bending plane, whereas the TRT has a resolution of about 150 μm. The inner detector vertex resolution, which depends on its number of charged particles, is about 80 (20) μm in the transverse plane and about 150 (40) μm in the longitudinal coordinate for vertices with 10 (80) charged particles [18].

1.2.2 Calorimeters

The ATLAS calorimeter system consists of an electromagnetic calorimeter that measures the energy of electrons and photons, and a hadronic calorimeter that measures that of charged and neutral hadrons. Its layout is illustrated in figure 1.5.

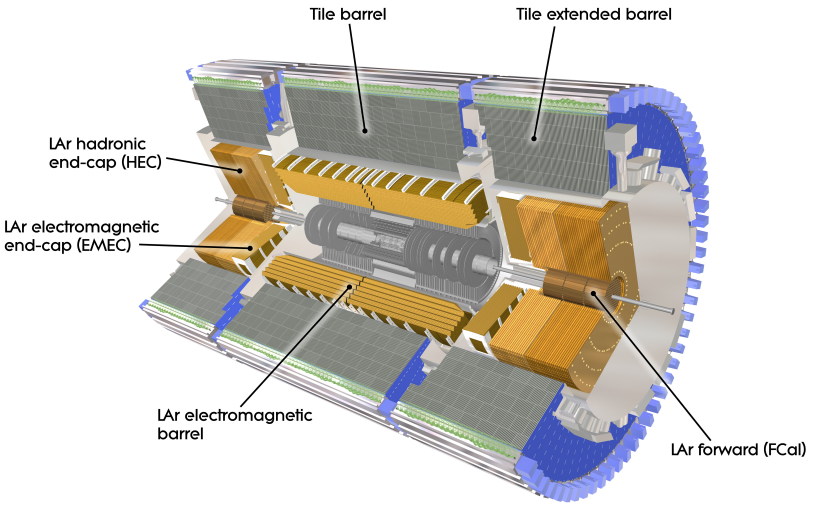

Figure 1.5: Illustration of the ATLAS calorimeter system. The principal sub-detectors are highlighted. Source: CERN.

The electromagnetic calorimeter is located outside the solenoid surrounding the inner detector, covers $|\theta| > 4.7^\circ$ ($|\eta| < 3.2$), and is divided into a barrel and two end-cap sections. It uses liquid argon as a sampling material and lead plates as absorbers. Its total thickness is about 24 radiation lengths in the barrel and about 26 in the end-caps. The constant term of the energy resolution has been measured to be about 1–3% [19].

The hadronic calorimeter surrounds the electromagnetic calorimeter and consists of a barrel and two end-cap sections. The barrel covers $|\theta| > 20.7^\circ$ ($|\eta| < 1.7$) and is divided into three sections. It uses scintillators as active material interleaved with steel absorbers. The end-caps cover $|\theta| > 0.9^\circ$ ($|\eta| < 4.9$) and are divided into two sections, called end-cap and forward. Both use liquid argon as the active material and copper and tungsten as absorbers. Its large angular coverage and total thickness of more than 10 hadronic interaction lengths are designed for high-energy jet measurements. Jet-energy scale uncertainties of order 1% or less have been measured [20].

1.2.3 Magnet system

A strong magnetic field is crucial to provide sufficient bending power for accurate track momentum measurements. ATLAS uses a combination of solenoid and toroidal magnets to achieve good accuracy at different angles. Air core reduces the degradation of momentum resolution due to multiple scattering. This choice of geometry, which is unique in the history of particle physics detectors, is what gives ATLAS (A Toroidal LHC Apparatus) its characteristic shape.

The magnet system, 20 m diameter and 26 m in length, is composed of superconducting coils stabilized with aluminium supports and cooled down to 4.6 K by liquid helium. Figure 1.6 shows the layout of the ATLAS magnet system. It consists of a central solenoid, eight toroidal windings in the barrel region, and two sets of four windings in the end-cap region.

The solenoid magnet provides a 2 T axial magnetic field that is present in the volume of the inner detector. The air-core toroidal magnets provide an average of 0.5 T both at the barrel ($|\eta| < 1.3$) and end-cap ($1.6 < |\eta| < 2.7$) regions of the muon spectrometer, that is described in the following section.

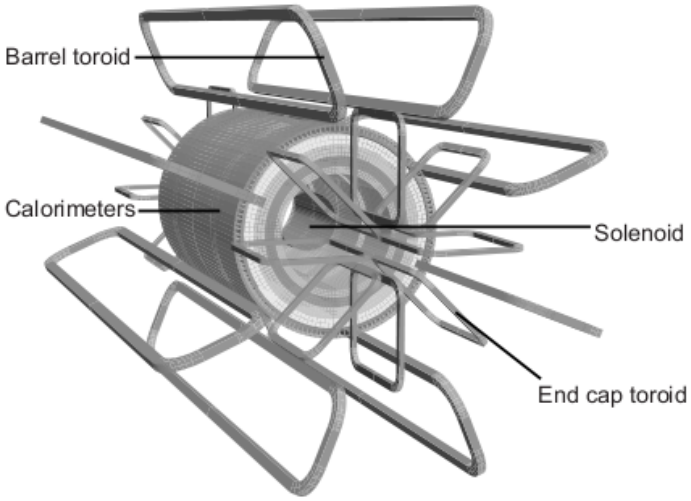

Figure 1.6: Layout of the ATLAS magnet system. Source: Ref. [21].

1.2.4 Muon spectrometer

The muon spectrometer is located at the outer part of the ATLAS detector. With the exception of neutrinos, high-energy muons are typically the only type of particle that penetrate through that region. The main purpose of the muon spectrometer is to measure[7] the trajectory of high-momentum muons, and to trigger events containing muons.

The muon spectrometer covers the region $|\theta| > 7.7^\circ$ ($|\eta| < 2.7$) and follows the geometry of the magnet system, as shown in figure 1.7. The barrel consists of three concentric layers at $5, 7.5$, and 10 m that cover $|\theta| > 40^\circ$ ($|\eta| < 1.0$). The end-cap consists of disk-shaped planes perpendicular to the beam axis at $|z| = 7.4$, 10.5 and 21.5 m that cover $7.7^\circ < |\theta| < 40^\circ$ ($1.0 < |\eta| < 2.7$).

The tracking detection is provided mainly by gas drift tube detectors (MDTs), except in the innermost layer of the forward direction where multi-wire proportional chambers (CSC) are used to better handle the higher particle fluence. The required resolution is achieved by optically monitoring the detectors alignment, with an accuracy of 30–60 μm; the achieved single-hit spatial resolution is 60–80 μm in the bending plane [22]. The time resolution is about 700 ns for the MDTs and 7 ns for the CSCs.

[7] The muon spectrometer can work "stand-alone", providing muon momentum information without measurements from the inner detector.

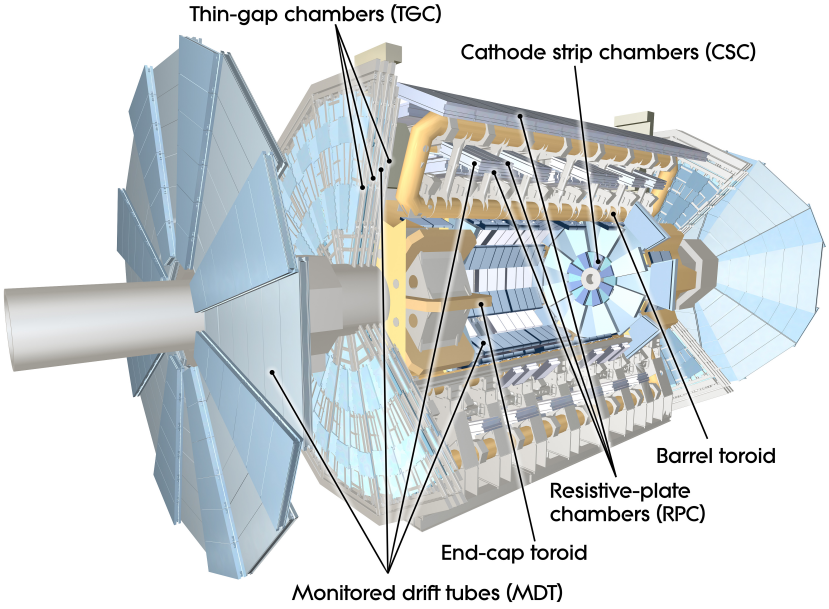

Figure 1.7: Layout of the ATLAS muon spectrometer. The location of the barrel and end-cap toroidal magnets is also shown. Source: CERN.

Due to the response time of the MDTS, separate triggering capability is provided by gas avalanche detectors: resistive plate chambers (RCPs) in the barrel region and thin gap chambers (TGCs) in the end-cap region. Their design position resolution is about 1 cm and their time resolution is about 2 ns. The RCP and TGCs are an integral part of the trigger system that is explained in section 1.2.6.

1.2.5 Forward detectors

The very forward region of ATLAS is instrumented with a variety of sub-detectors that serve different purposes. These include: LUCID and BCM, a Čerenkov and a solid-state detector used to measure luminosity (described in chapter 4); ALFA, a Roman-pot detector used to measure elastic scattering (described in chapter 5); MBTS, an array of scintillators used to trigger on minimum-bias events (described in chapter 5); ZDC, a calorimeter used to measure neutron emission in heavy-ion collisions (described in Ref. [23]).

1.2.6 Trigger and data-acquisition system

The LHC delivers bunch crossings at 40 MHz, which results in a rate of inelastic interactions of about 1 GHz due to the pile-up of multiple proton–proton collisions in the same bunch crossing. The majority of those events are not interesting for the physics programme of the LHC. Even if the signals are real and interesting, not all of the events can be recorded. To handle this problem it is necessary to introduce the concept of trigger system.

The trigger system is a combination of hardware and software used to select events with potentially interesting signatures in real time, and to provide a rejection power of about 10^6 for minimum-bias interactions. The trigger system uses information provided by various sub-detectors. A "menu" of different algorithms process that information and decide if the event should be stored.

The ATLAS trigger system consists of two levels. The first selection known as Level-1 trigger is entirely implemented in hardware. It uses about 7200 composed signals from the muon spectrometer and the calorimeter system. The Level-1 trigger selects signatures, called "objects", including high-p_{T} muons or electrons, photons, jets, etc. The trigger logic can be programmed flexibly, e.g by requiring different p_{T} thresholds or various objects in coincidence. The Level-1 trigger reduces the event rate to about 85 kHz with a latency of about 2 μs. During this time the data of all channels of the detector are stored in pipelines at the front-end electronics.

The second trigger level is named high-level trigger (HLT) and is software based. It consists of a series of algorithms that use the full granularity of the detector within the "region of interest" defined by the Level-1 (i.e. η–ϕ region of objects). The HLT uses reconstruction algorithms that reach similar performance to that of offline analysis and reduces the output rate to about 1 kHz with a latency of about 10 ms.

Typically, if an event is selected by an HLT algorithm, the corresponding data are stored to disk; however, to reduce the output rate of certain algorithms a "prescale" factor is applied. If the prescale factor applied is N, only 1 in N events are randomly selected to be stored. The HLT menu and the corresponding prescale factors are typically changed over time to accommodate the observed output rate. However, they are only changed at the end of "luminosity blocks" that are periods of time of about 1 minute.

1.3 The high-luminosity LHC

Given that the center-of-mass energy achievable in a collider is limited by the diameter of the accelerator ring for a fixed magnetic field strength, plans for upgrades of the LHC have been focused mostly on increases in luminosity. The high-luminosity LHC (HL-LHC) project [24] consists of a series of upgrades to the accelerator that include more powerful focusing magnets and more precise radio-frequency cavities. The HL-LHC design luminosity is 10 times that of the LHC, i.e. 10^{35} cm^{-2}s^{-1}. Its goal is to deliver a total of 3000 fb^{-1} running over a decade starting in 2025.

The main motivation for the HL-LHC is to effectively extend the reach of potential discoveries of new particles to about 30% higher masses, depending on the model, and to study with more statistical precision the properties of the Higgs boson [25]. As statistical uncertainties decrease roughly as the square root of the integrated luminosity, running the LHC at the current luminosity will result in diminishing statistical gains.

The higher particle rate and harsher radiation environment expected at the HL-LHC pose stringent requirements on the detector technology and readout electronics used in the experiments. To cope with this, various ATLAS sub-detectors such as the muon spectrometer, the liquid-argon calorimeters, and the inner detector will be upgraded.

1.4 The inner tracker project

The current ATLAS inner detector will be completely replaced with the "inner tracker" (ITK) [26]. The ITK will be an "all-silicon" tracker, because the gas-tube technology used in the TRT cannot cope with the particle multiplicity expected at a luminosity of 10^{35} cm^{-2}s^{-1}. The ITK current layout covers $|\theta| > 7.7^{\circ}$ ($|\eta| < 2.7$) and consists of layers of pixel and micro-strip sensors.

The goal of the ITK design is to at least preserve the performance of the current inner detector, thus meeting the challenges imposed by the much denser environment at the HL-LHC. To achieve this, the new detectors will have an increased granularity. The expected hit occupancy at the HL-LHC design luminosity is 1% or less. The current prototype design of micro-strip sensors to be used in the ITK is described in chapter 2. One of the key aspects of this design is the stricter requirement of radiation hardness. Studies of radiation damage in silicon sensors are presented in chapter 3.

Chapter 2

Characterization of ATLAS12 prototype sensors

This chapter summarizes the working principle of silicon-strip detectors, as well as some of the technical characteristics present in the design of current and future sensors used in the tracker of the ATLAS experiment. This summary is based on recent literature ([27–31], and references therein) on silicon detectors with an emphasis on tracking applications for experiments at the LHC and future upgrades.

2.1 Silicon-strip sensors

2.1.1 Working principle

Ionizing radiation can be detected by measuring the charge that is liberated as it traverses a medium. An electric field is needed to collect the liberated charge onto electrodes before it recombines. The resulting signal is then amplified and shaped by readout electronics. This principle is illustrated in figure 2.1.

This principle of detection can be realized with the collection of electron-hole pairs generated in a silicon p-n junction under reverse bias. The reverse bias provides the electric field and a volume depleted of the free charges provided by the doping. This depletion volume is effectively the only active volume of the detector, as the free carrier concentration (typically $\approx 10^{12}$ cm^{-3}) overwhelms the charge generated by a minimum-ionizing particle in silicon (≈ 80 electron-hole pairs/μm).

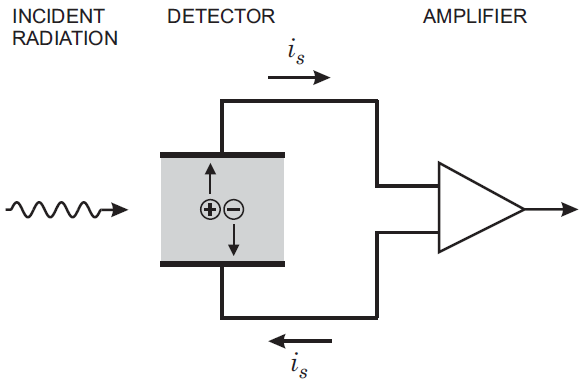

Figure 2.1: A basic principle of detection of ionizing radiation. Source: Ref. [28].

In order to achieve the position resolution needed for particle-tracking applications, a large area diode is segmented into many parallel strip shaped structures that act as separate detectors. Figure 2.2 shows an illustration of a p-in-n detector. The thin (≈ 10 μm wide, ≈ 1 μm deep) p-type strip implants act as electrodes for holes. The doping concentration is much lighter in the thick (≈ 300 μm) n-type bulk in order to create a very asymmetric junction that depletes mostly into the bulk, as follows from the constraint of overall neutrality of the depletion volume.

The typical values for modern detectors are a separation between strips, called pitch, of ≈ 100 μm, and a strip length of ≈ 10 cm. A thin (≈ 1 μm) aluminium metallization layer on top of the strip implants is used to provide contact to the readout electronics. A very thin (≈ 1 μm) oxide layer covers most of the segmented area of the sensor to provide protection from the environment. The backplane of the sensor consists of a thin aluminium metallization layer and a highly doped ($\approx 10^{13}$ cm^{-3}) n-type layer. The aluminium provides an ohmic contact and the highly doped layer prevents the depletion volume reaching the metallization layer.

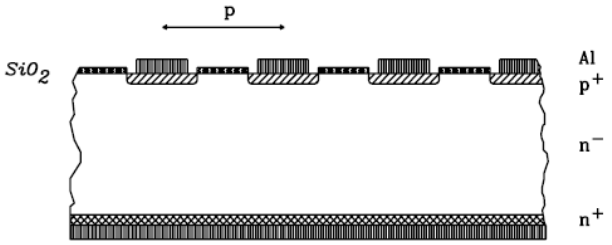

Figure 2.2: Illustration of the cross-section of a silicon strip detector. Source: Ref. [29]

Apart from the p-in-n detector just described, there are various possible combinations of strip and bulk doping. For example a possible detector can use heavily doped n-type strips and n-type bulk (n^+-in-n), or n-type strips in p-type bulk (n-in-p). The next generation of silicon-strip sensors to be used in the ATLAS experiment are of the n-in-p type, for reasons described later.

The following sections describe several of the structures that compose a typical silicon-strip sensor. Emphasis is given to the structures and technological choices made by the ATLAS collaboration for the detectors to be used at the HL-LHC.

2.1.2 Readout coupling

The signal charge can be read out with DC coupling: a direct connection among strip implant, metallization and readout electronics. The problem with DC coupling is that all of the detector leakage current is transmitted to the amplifiers, potentially causing saturation.

The alternative is to use a capacitor and a resistor in a network that prevents the transmission of the DC part of the leakage current into the readout electronics; such an approach is called capacitive coupling or AC coupling. DC and AC coupling are illustrated in figure 2.3. Both the capacitors and resistors can be incorporated in the sensor. Capacitances are built by inserting a thin oxide layer between the metallization layer and the strip implants; bias resistors are part of the bias structure that is explained in section 2.1.4.

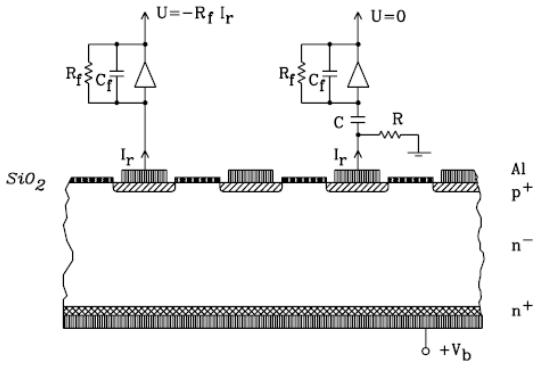

Figure 2.3: DC and AC coupling readout. In AC coupling (right) an extra circuit prevents the DC component of the leakage current to being transmitted to the amplifiers. Source: Ref. [29].

2.1.3 Bias and guard rings

The active area of the sensor is surrounded by the bias ring, an implant with similar doping and width as the strip implants. The bias ring is covered by and connected to an aluminium metallization layer. Together with the backplane, the bias ring provides the main contact for biasing the sensor. Each strip is connected to the bias ring by bias resistors, that are explained in section 2.1.4. In this way it supplies each strip implant with the same bias voltage.

The guard ring surrounds the bias ring and is similar in structure, but is used to shape the electric field at the borders of the sensor. A proper design, that can include multiple guard-rings, avoids concentrations of electric field to prevent micro-discharges.

2.1.4 Strip biasing

The bias resistors are typically implemented using polycrystalline silicon, also known as poly-silicon. This material is used in ATLAS sensors because it has been proven to be resilient to radiation damage. A direct bias to each strip implant is also possible with a direct metallization-implant contact called the "DC pad" or "probe pad". This feature is used for testing purposes because a direct biasing for each strip allows the identification of defects. Figure 2.4 shows an illustration of a bias resistor and probe pad in a silicon-strip sensor.

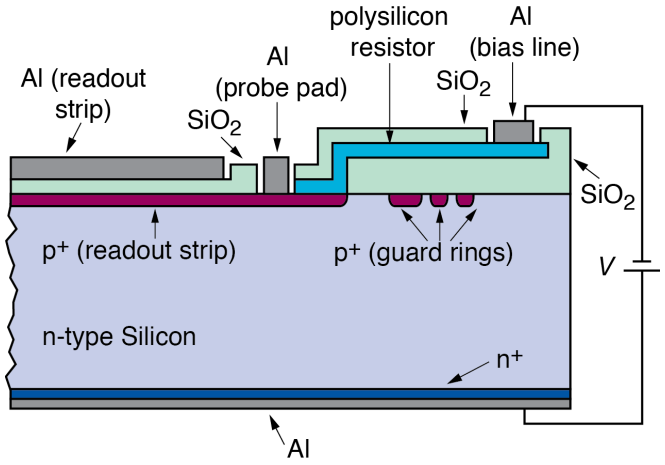

Figure 2.4: Cut through of a silicon-strip sensor, parallel to the strip, illustrating different features including the probe pad, poly-silicon resistor, and guard rings. Source: Ref. [32].

2.1.5 Strip isolation

Because of their lower mobility, holes created in the oxide layer get trapped more often than electrons in mid-gap energy levels created by impurities and defects, leading to a net positive charge in the oxide layer.

The positive charge in the oxide is a problem if the sensor collects electrons (n-type strips) because the electrons attracted to the surface create a conductive channel among strips, providing a short circuit. This is not a problem in the p-in-n sensors used in ATLAS for the LHC era (that collect holes), but are a serious concern for the n-in-p sensors to be used in the HL-LHC era [33].

There are various technological solutions to this problem. The preferred one for ITk is to use lightly doped p-type implants inserted between strips to repel electrons. In that way the accumulation layer that shortens the strips is avoided. The p-type implants can be restricted to certain areas (p-stop method) or be spread across strips (p-spray method). Figure 2.5 illustrates the p-stop approach.

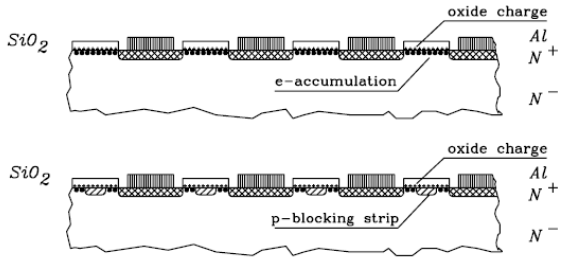

Figure 2.5: The p-stop principle. The electron accumulation layer that is created due to positive charge in the oxide (upper figure) is avoided with p-type implants between strips (lower figure). Source: Ref. [29].

2.1.6 Sensor parameters

In this section the main parameters of silicon-strip sensors are described. The design of the sensors is tuned to achieve the desired values for a particular application. These define the noise level, power consumption, position resolution, speed, etc. Given that some of these are competing requirements, there is a compromise in the choice of parameters depending on the targeted performance. For the LHC and specially for the HL-LHC the sensor design is driven by the (long term) effects of radiation damage.

1. **Full-depletion voltage**.
   The full-depletion voltage, V_{FD}, is the bias voltage needed to extend the depletion volume from the strips to the backplane of the sensor. It is approximately given by:

$$V_{\mathrm{FD}} = \frac{qNd^2}{2\epsilon}, \tag{2.1}$$

   where d is the thickness of the sensor, N is the effective doping concentration (acceptors − donors), q is the electron charge, and ϵ is the permittivity of silicon. Ideally a sensor is operated at or above full-depletion voltage to maximize the active volume. For that reason, and because it lowers power consumption, the sensor design generally aims for low V_{FD}.

   The V_{FD} changes with radiation damage in the silicon bulk that is created by energetic hadrons because this creates acceptor-like defects (i.e. mid-gap energy levels near the valence band) that change the effective doping of the bulk. In p-type bulk this effect increases V_{FD} with increased damage. More dramatically, in n-type bulk the effective doping is initially reduced and eventually the bulk becomes effectively of p-type. This "type inversion" phenomena makes the depletion volume grow from the heavily-doped layer in the backplane rather than from the segmented side. To avoid this, only p-type bulk is considered for sensors for the HL-LHC.

2. **Leakage current**.
   Even in the absence of ionizing radiation, electron-hole pairs that are thermally generated in the depleted volume will produce a current through the detector, called leakage current. It is driven mostly by mid-gap energy levels that act as generation-recombination centres. These are created by impurities or dislocations in the silicon lattice. Thus, radiation damage increases the leakage current.

   The leakage current depends on the bias voltage applied to the sensor. Ideally it is proportional to the depleted volume, and thus proportional to the square-root of the bias voltage, V_{bias}, until it saturates at full-depletion voltage. However this dependence is affected by other effects such as surface currents, currents through the edges of the sensor, or micro-discharges.

   The leakage current increases exponentially with temperature. This implies that the corresponding power dissipated could create a positive feedback, called "thermal runaway". To minimize leakage current and to avoid thermal runaway, the current silicon-strip sensors in ATLAS are operated at a target temperature of $-7°\mathrm{C}$.

3. **Breakdown voltage.**
   When the electric-field strength in the sensor exceeds ≈ 30 V/μm, the generated electron-hole pairs gain enough energy to produce new electron-hole pairs. This leads to a very strong current increase, restricted to a number of channels, potentially saturating the readout electronics of the whole sensor. Moreover, this could cause a thermal runaway. Proper design and testing of bias and guard ring structures are needed to ensure that the planned operation voltage does not lead to breakdown.

4. **Capacitance: coupling, inter-strip and bulk**.
   The metallization layer, thin oxide layer, and strip implant form the classical metal-oxide-semiconductor capacitor. Its capacitance, known as coupling capacitance, is an important figure of merit of the sensor as it dictates the signal transfer efficiency to the amplifiers.

   A large coupling capacitance is obtained with a very thin oxide layer, the limiting factor being the reliability of the production process. Silicon oxide as well as silicon nitride are used in the ATLAS silicon-strip sensors with a thickness of approximately 0.2 μm, yielding a coupling capacitance of ≈ 30 pF/cm.

   The AC coupling readout design depends on the coupling capacitance dominating all other possible capacitances in the sensor, like the inter-strip capacitance that is formed between strip implants, and the bulk capacitance that is formed between the strips and the backplane. This is considered in the sensor design because the width and length of strips, pitch, and bulk depth determine the value of these capacitances.

   The bulk capacitance is inversely proportional to the depletion depth that grows with $\sqrt{V_{\text{bias}}}$ until it saturates after full depletion. Thus, V_{FD} can be obtained in a voltage scan as the voltage at which the bulk capacitance saturates. This method is used in the results presented in section 2.2.4.

5. **Resistance: bias resistor, inter-strip and metal**.
   The resistance of the poly-silicon structure used to bias the strips is called bias resistance, R_{bias}. The bias resistance must be large to minimize the thermal noise it produces and uniform among strips to avoid voltage differences. Typically R_{bias} is ≈ 1 MΩ.

   The bulk resistance that isolates the strips is called inter-strip resistance, R_{inter}. The R_{inter} value must be much larger than R_{bias} in order to ensure that practically

all the leakage current is transmitted to the bias ring, and to maintain proper strip isolation. Radiation damage lowers R_{inter} so it can be a problem if it is comparable with R_{bias}.

Finally, the aluminium layer resistance acts as a series resistor for the readout electronics. For the short peaking times needed for experiments at the LHC ($\approx$ 10 ns), a high series resistance impacts the noise levels, therefore it should be limited. This requirement is satisfied with a thin ($\approx$ 1 μm) aluminium layer.

6. **Charge collection efficiency.**
   The ionization charge liberated by a minimum-ionizing particle traversing silicon is $\approx$ 80 electron-hole pairs/μm; thus $\approx$ 24k in a 300 μm thick detector. The charge collection efficiency is defined as the ratio of the measured charge to the expected charge considering the full active volume of the sensor. For good performance, the ratio between the collected charge and intrinsic noise must be large.

   The collected charge depends on the amount of impurities and defects in the silicon bulk that contribute to charge-trapping, and the drift velocity of the charge carriers. Given that the mobility of electrons is around three times larger than the mobility of holes, the electron collection is less affected by trapping.

   Given that radiation damage increases V_{FD}, effectively reducing the depletion volume at a given V_{bias}, it decreases the collected charge. Moreover, radiation damage also increases charge trapping, leading to a further reduction of the collected charge.

2.2 Electrical properties of ATLAS12 sensors

This section describes the characterization of silicon-strip sensors to be used in the ITk. Several tests are performed on every sensor with the aim to cross-check their quality and uniformity. More detailed characterization results are shown as well. These measurements were performed in the clean-room facility of the HEP group at Cambridge using an existing setup. This work is part of a collaborative effort that has been published in Ref. [1].

2.2.1 The ATLAS12 sensor

The ATLAS ITK collaboration designed a prototype sensor[1] for the HL-LHC with the best radiation tolerant technology known to date [34]. They have a large-area[2] (9.75×9.75 cm^2), which is the maximum possible with 6-inch silicon wafers. High resistivity p-type silicon is used for the bulk with a thickness of ≈ 310 μm, and n-type AC coupled strips. The strips are biased through poly-silicon resistors and isolated with p-stop structures. The sensor is divided into four segments of 1280 strips of 2.39 cm length with a pitch of 74.5 μm. The design also includes a single guard ring.

One of the distinct features of this design, called "ATLAS12" or "A12", is that it uses p-type bulk and incorporates p-stop structures[3]. Current SCT sensors were constructed with n-type bulk that does not need p-stop structures. As discussed in section 2.1.6, this choice was driven by stricter requirements on radiation tolerance.

The substantial reduction in strip length from the current 6 cm and the slight reduction in pitch from the current 80 μm were chosen to lower the hit occupancy to cope with the higher particle-multiplicity expected at the HL-LHC.

2.2.2 Technical specifications

The sensor technical specifications [35] include mechanical properties like thickness, flatness, cut-edge defects; and also electrical properties like leakage current, bias resistance, full-depletion voltage, coupling capacitance, etc. Table 2.1 shows a summary of the technical specifications of the A12 sensor design.

In the early stages of R&D the participating institutes measure every sensor and perform more detailed studies than provided by the contractor with the aim of establishing the uniformity of production, and testing whether it complies with the specifications.

The key goal of the development of this new generation of sensors is to maintain acceptable performance after the severe radiation damage expected over the lifetime of the HL-LHC. Radiation tolerance is studied using miniature sensors (1×1 cm^2) fabricated on the same wafer as the full-size sensors. Measurements of mini-sensors are presented in chapter 3.

[1] These sensors are fabricated by Hamamatsu Photonics.
[2] Larger than the area the sensors currently used (6.4×6.4 cm^2).
[3] A predecessor design [33] known as ATLAS07 differs in details of the p-stop and guard-ring design.

Table 2.1: Summary of the technical specifications of the A12 sensor design. Source Ref. [35].

Thickness	310 ± 25 μm
Bulk type	p
Resistivity	< 4 k$\Omega \cdot$ cm
Outer dimension after dicing	97.54×97.54 mm^2
Inner edge dimension of bias ring	95.70×95.64 mm^2
Strip segments	4
Strip segment length (approximate)	23.86 mm
Strip implant type	n
Strip pitch	74.5 μm
Strip implant width	16 μm
Strip bias resistor	poly-silicon
Strip bias resistance (R_{bias})	1.5 ± 0.5 MΩ
Strip readout coupling	AC
Strip readout metal width	20 μm
Strip AC coupling capacitance	> 20 pF/cm
Strip isolation	$R_{\text{inter}} > 10 \times R_{\text{bias}}$ at 300 V
Strip isolation method	p-stop
Micro-discharge onset voltage	> 600 V
Leakage current	< 2 μA/cm^2 at 600 V
Radiation tolerance	1.2×10^{15} $n_{\text{eq}} \cdot$ cm^{-2}

2.2.3 Instrumentation

This section summarizes the clean-room instrumentation used to perform the measurements of full-size prototype sensors that are described in section 2.2.4.

1. **Probe station:**
   The Alessi Rel-6100 semi-automatic probe station, shown in figure 2.6, consists of a chuck on which the device under test is placed and a support structure that holds probe needles and a microscope. The sensor mounted on a plastic jig is fixed to the chuck through a vacuum system. Its position can be adjusted with the control software that also provides pattern recognition for alignment purposes (with fiducial marks in the sensor). Tungsten needles are used to provide electrical contacts to the sensor. The microscope has a maximum magnification of $\times 50$ and an incorporated camera.

2. **LCR meter:**
   The precision component analyser (called LCR meter) is a device used to measure inductance, resistance and capacitance of a device under test. The LCR meter provides measurements at a frequency adjustable in the range from 20 Hz to 3 MHz. Resonant frequencies of the setup must be avoided, so before any measurement a frequency scan is performed. For certain measurements the frequency is fixed by the technical specifications. The Wayne Kerr LCR meter 6440B is shown in figure 2.6.

3. **Source-measurement units:**
   Source-measurement units (SMUs) are devices capable of providing voltage and measuring current at the same time. The Keithley 487 and Keithley 6517 SMUs are used to measure the small leakage currents of the sensors and to supply the required bias voltages. These SMUs are shown in figure 2.6.

4. **Switching matrix:**
   This device controls the input and output of several other devices. It is controlled by LabVIEW scripts to automate various measurements. This device is shown in the bottom of the right panel of figure 2.6.

5. **Optical coordinate measurement machine:**
   An optical coordinate measurement machine (CMM) allows the measurement of the sensor using a contactless (optical) technique that prevents mechanical damage. It is used for visual inspection, and to measure dimensions and planarity of the sensors. The BATY optical CMM used is shown on the left panel of figure 2.7.

6. **Wire-bonder:**
   An automatic bonding machine is used to place wire bonds that connect the voltage supply and the sensor. That is, 25 μm aluminium wire bonds are used to connect the passivation windows in the sensor surface (connections to the bias ring or backplane) and the metal pads in the plastic jig where the sensor is mounted. The metal pads in the plastic jig are then connected to a bias supply. The wire-bonder machine is shown on the right panel of figure 2.7.

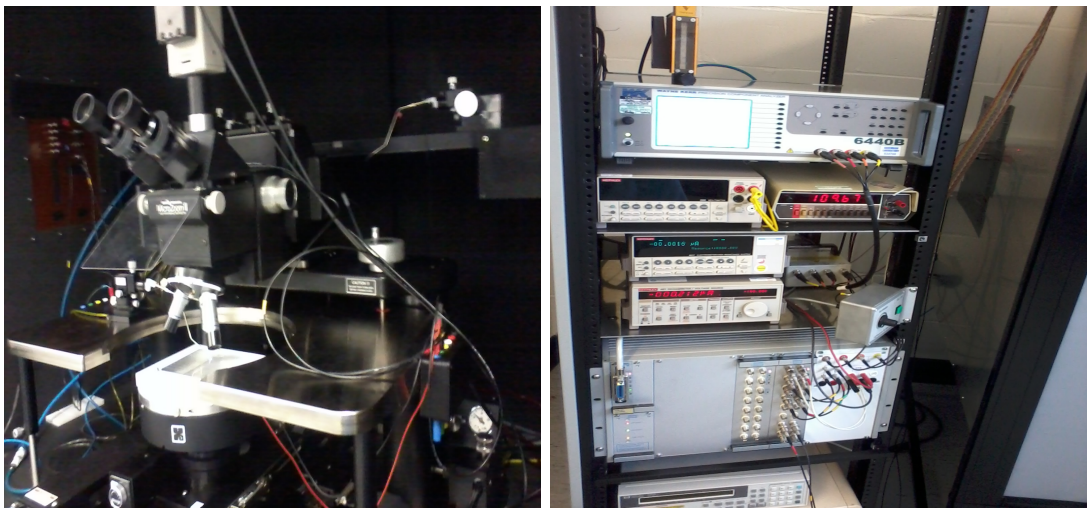

Figure 2.6: Left panel: semi-automatic probe station and sensor in a white plastic jig. Right panel: from top to bottom: LCR meter, SMUs and switching matrix.

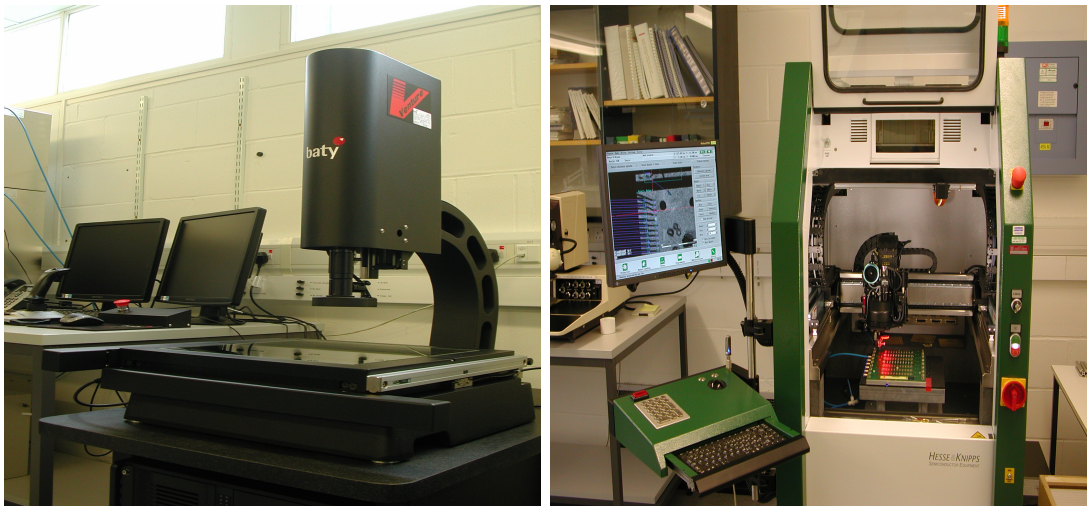

Figure 2.7: Left panel: BATY optical CMM used for visual inspection. Right panel: automatic wire-bonder used to make wire-bonds to connect the voltage supply to the sensor.

2.2.4 Measurements

1. **Visual inspection:**
   Every sensor is inspected visually with the BATY microscope before any other measurement is performed. The aim of this test is to ensure the sensor is free from defects and scratches that could have originated during handling or packaging. The main focus is on the bias and guard rings because defects in this region are prone to trigger sparks when the sensor is biased at high voltages. Figure 2.8 shows a picture of a scratch in this region spotted during visual inspection. This particular feature did not have an impact on the electrical properties of the sensor under test.

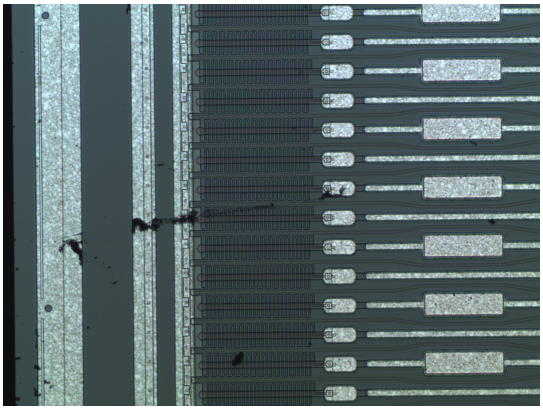

Figure 2.8: Picture of a portion of the sensor a showing scratch spotted during visual inspection.

2. **Leakage current**
   The leakage current is measured at different bias voltages with the Keithley 6517 SMU; this is called an IV scan. This measurement allows a basic check of the sensor quality. The electrical contact to the backplane (negative HV) and bias ring (ground) can be established in two ways. The first alternative uses wire-bonds from the bias ring and backplane bond pads to the metal pads present in the plastic jig where the sensor is mounted. The metal pads are then connected to the SMU. The second alternative uses two tungsten needles connected to the SMU to dig into the bias ring and backplane bond pads. The second alternative is used.

   The IV scan is performed in 10 V steps from 0 to −1000 V with a 10 s delay between steps to avoid transient effects. The measurement is aborted when the sensor presents breakdown, i.e. a sudden increase of leakage current following an increase in bias voltage.

Figure 2.9 shows the IV measurements of 18 sensors, corresponding to the production batch with code VPX12518. None of the sensors presented breakdown so all of them meet the technical specifications. The leakage current at 600 V is about three orders of magnitude lower than the technical specification. The increase of current around 360 V is due to a surface component that appears when the depletion volume reaches the backplane of the sensor. As will be shown in next section, 360 V is close to V_{FD}.

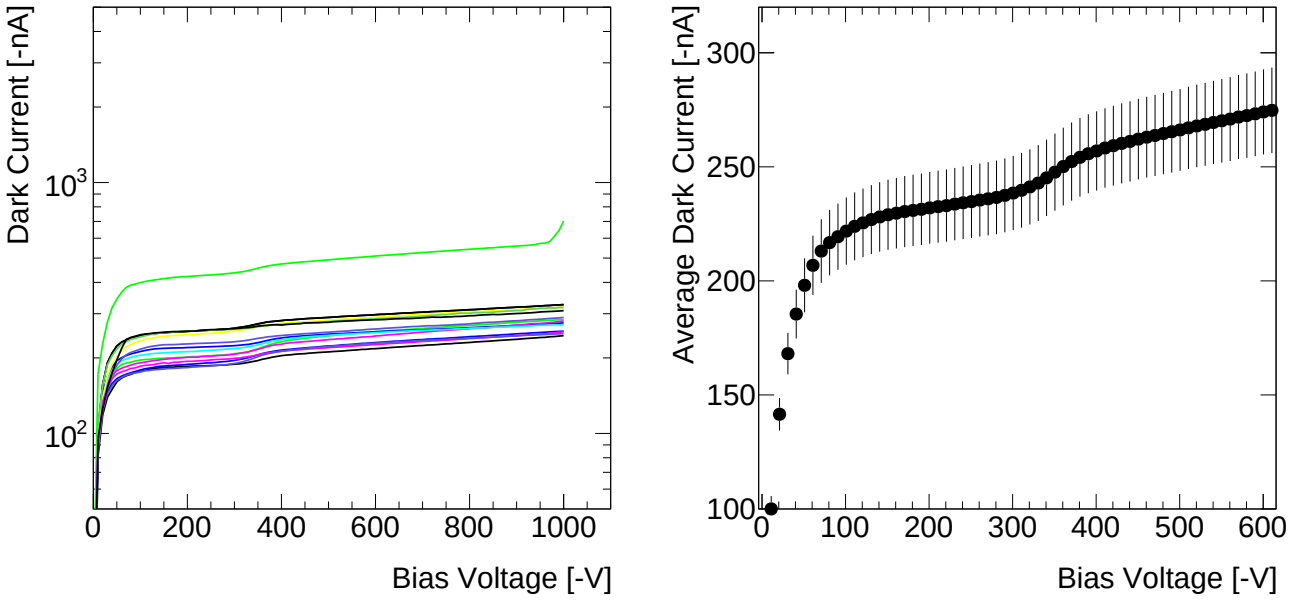

Figure 2.9: Left panel: leakage current as a function of bias voltage of 18 sensors measured at room temperature. Right panel: average of results shown on left panel. The error bars represent the standard deviation of the measurements.

3. **Full-depletion voltage**

   The capacitance between the bias ring and the backplane, C_{bulk}, is measured at 1 kHz in 10 V steps from 0 V up to −750 V or below if the sensor presented breakdown; this is called a CV scan. A delay of 10 s between measurements is applied to avoid transient effects.

   As discussed in section 2.1.6, the voltage dependence of the C_{bulk} can be used to estimate the full-depletion voltage. The data are plotted as $1/(C_{\mathrm{bulk}})^2$ vs V_{bias}, and V_{FD} is extracted from the intersection of linear fits before and after saturation.

   Figure 2.10 shows the CV characteristics of 18 sensors on the left panel, and the average on the right panel. The average V_{FD} extracted is 368 $\pm$ 2 V; the standard

deviation of the V_{FD} extracted independently for every sensor is 7 V. The uniformity is good but the value is higher than the desired ≈ 200 V (or equivalently, silicon resistivity $> 4\ \mathrm{k}\Omega \cdot \mathrm{cm}$). Since the specifications include a clause dependent on the availability of the required silicon wafer resistivity, the sensors meet the technical specifications.

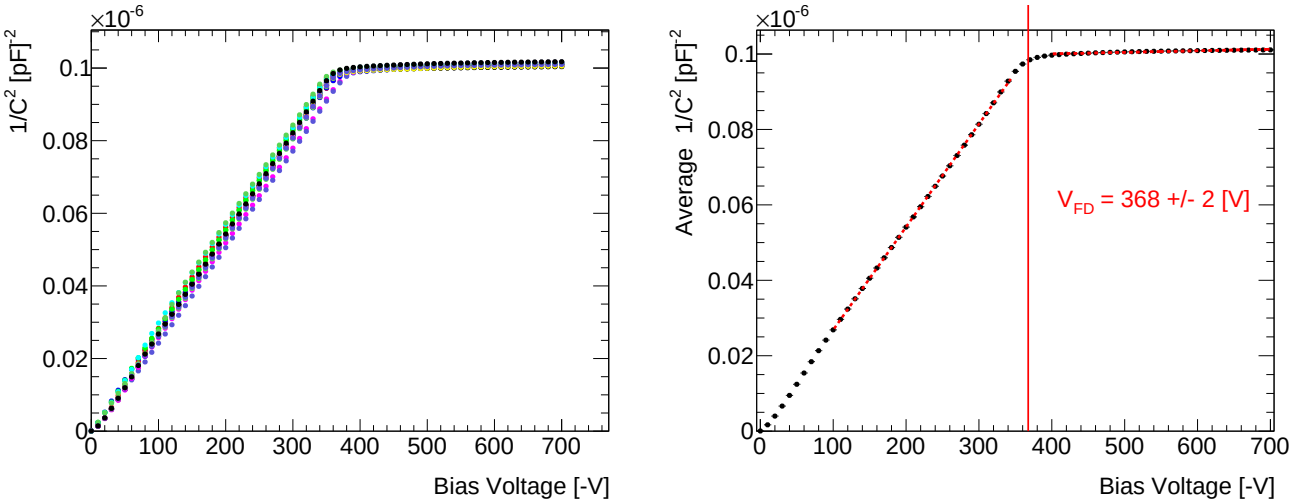

Figure 2.10: Left: CV characteristics of 18 sensors measured at room temperature. Right panel: average of the results shown on the left panel. The result of two linear fits are superimposed, from which the average full-depletion voltage is extracted (shown as a vertical line).

4. **Strip integrity:**
   The aim of the strip integrity test is to check each strip (1280×4 total) for defects in the oxide (pinholes), shorts between strip metals, and for discontinuity in the strip metal layer or implants. The SMU Keithley 487 is used to check for pinholes, a CV meter is used to measure the strip-backplane capacitance, and a switching matrix to control the measurement steps.

   Pinholes are determined by a measurement of current between the strip metal and the backplane with +100 V on the strip, and the backplane on ground potential. If the oxide is faulty a short is expected, so a current-limiting series resistor of 2 MΩ is included. Then, a measurement of the strip-backplane capacitance at 1 kHz is performed with the sensor bias at −150 V to isolate the strips. Shorts between strips are identified with capacitance values consistent with a series connection between two or more strips.

   The technical specifications require that the number of good strips is $> 99\%$ and with no sensor falling below 98%. From a total of 92160 measurements only 19

strips were identified as faulty. Figure 2.11 shows the distribution of faulty strips in the geometry of the sensor. No hot spots, that could have been indicative of a production problem, are observed.

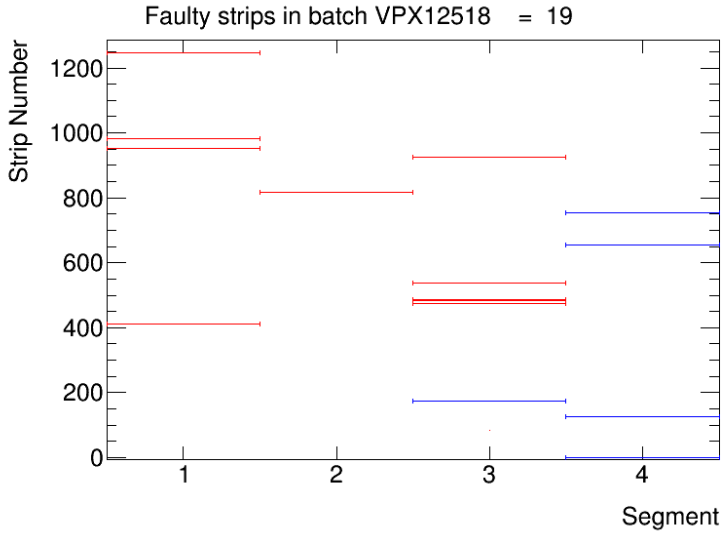

Figure 2.11: Sensor representation showing the location of faulty strips. Faulty strips are 19 out of a total of 92160 measured. Red lines represent pinholes and blue lines defects on strip implant.

5. **Full Strip Test:**
   The bias resistance and coupling capacitance are measured for each strip. The setup is similar to the strip integrity test. This measurement requires a voltage source to bias the sensor, a LCR meter and a switching matrix. The sensor is biased at -150 V to isolate the strips. A probe needle is used to connect the strip bond pad to the input of the LCR meter. The measurement is performed at 1 kHz.

   Figure 2.12 shows the measurements of a representative sensor classified according to segment and strip number. The coupling capacitance measurements show a dependence on the strip number. This has been confirmed to be a feature of the detector and not the setup by inverting the sequence of the measurement and other checks. The results are compatible among all the strip segments. The poly-silicon bias resistance distributions show a few percent spread. The measurements with slightly larger values are caused in most cases by a bad electrical contact between the probe needle and the strip bond pad.

   Figure 2.13 shows the distribution of 92160 strips corresponding to 18 sensors. The standard deviation of the distribution is about a percent. This shows the very

good uniformity of the production. The measured sensors easily meet the technical specifications of $R_{\text{bias}} = 1.5 \pm 0.5$ MΩ and $C_{\text{coupling}} > 20$ pF/cm (the strips are 2.39 cm long).

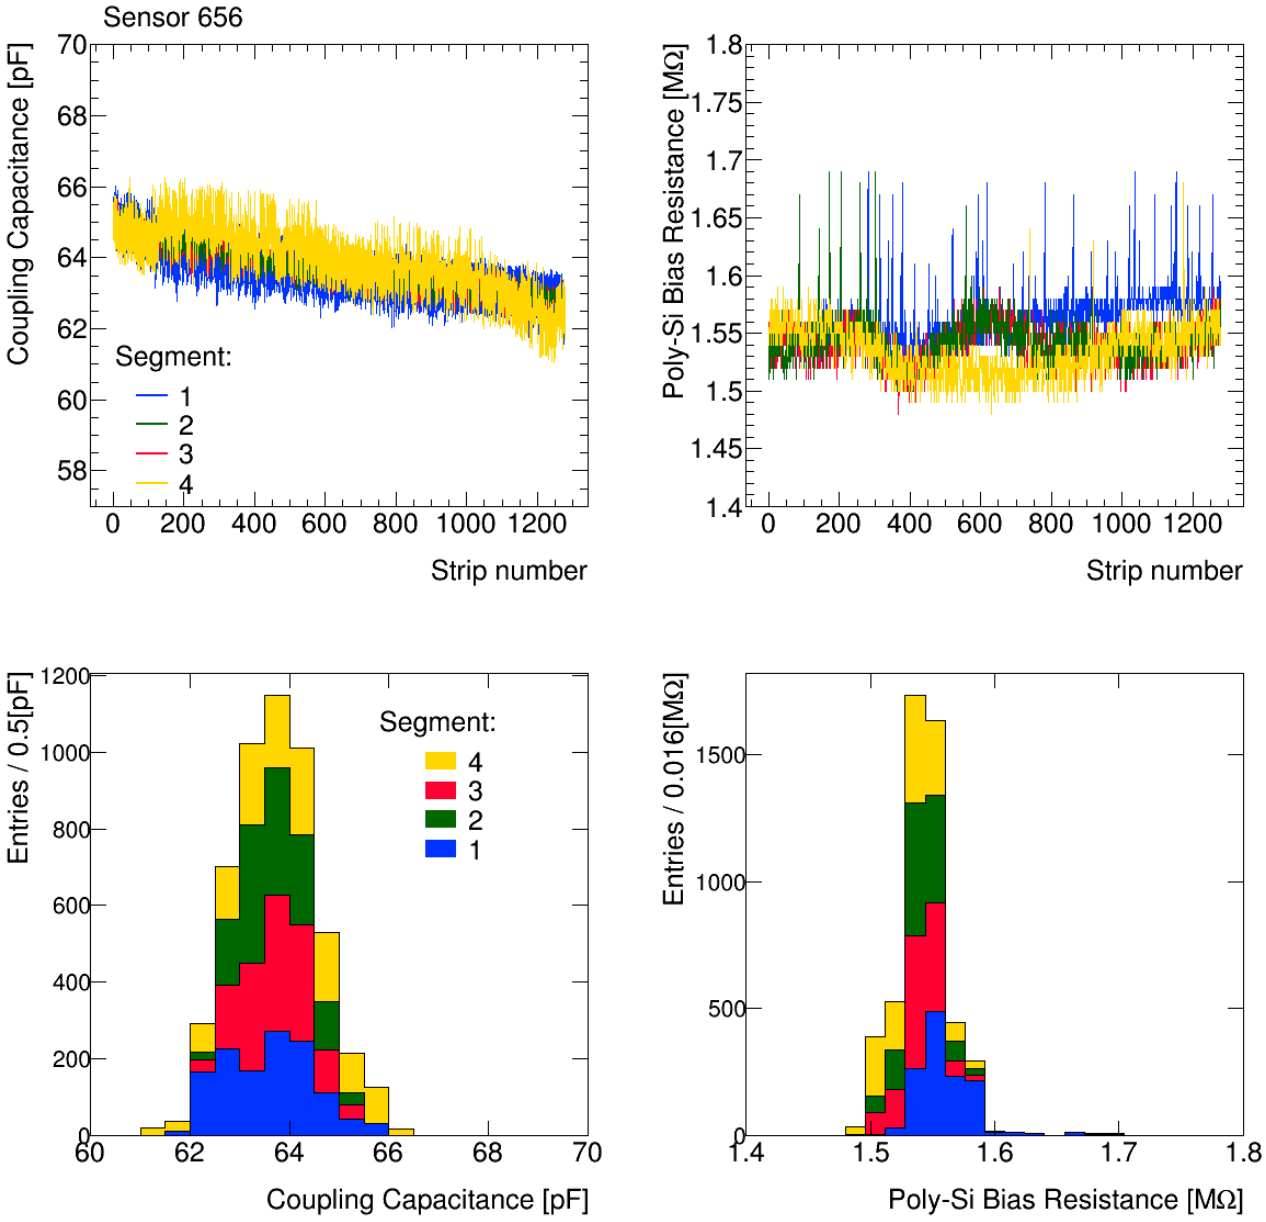

Figure 2.12: Coupling capacitance and bias resistance measurements of all strips of a representative sensor.

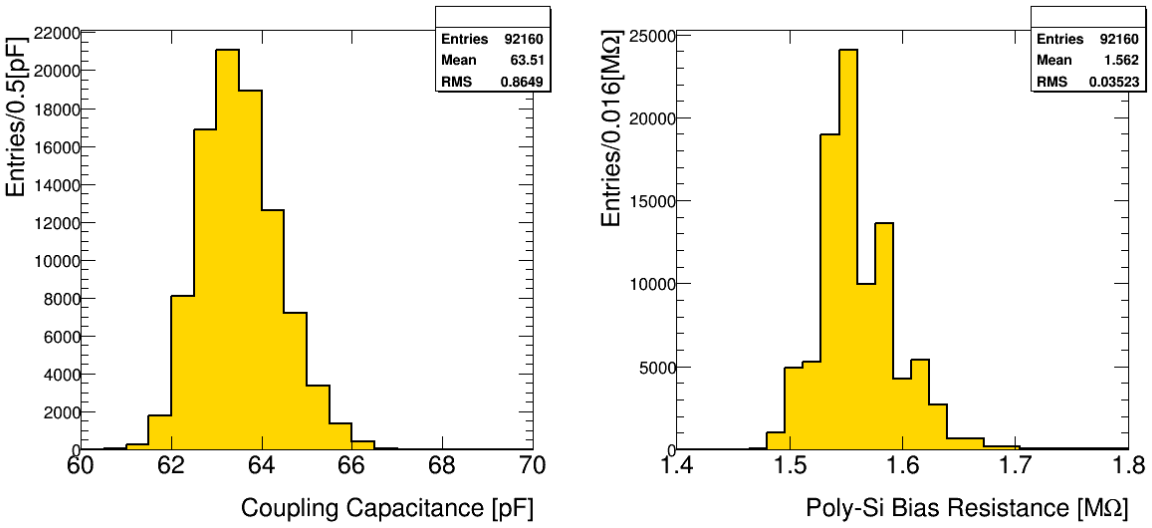

Figure 2.13: Distribution of coupling capacitance (left) and bias resistance (right) measurements of all strips of 18 sensors.

Chapter 3

Studies of radiation damage in silicon-strip sensors

3.1 Introduction

This chapter describes studies of radiation damage in silicon-strip sensors for the ATLAS inner detector upgrade for the high-luminosity LHC. These studies are part of a coordinated effort across several institutes [33] that aim to characterize the degradation of performance of the sensor over the lifetime of the experiment.

For research and development purposes, miniature versions of the main silicon sensors were irradiated and tested. These miniature sensors ($\approx 1 \times 1$ cm^2) were fabricated with the same wafer as the large-area sensors ($\approx 10 \times 10$ cm^2). They have 104 strips with a length of 8 mm. Their irradiation was performed in several facilities with different particles (neutrons, protons, gammas and pions) and different energies.

The total dose of different particles types is scaled to a common unit named "equivalent number of 1 MeV neutron per centimetre square" or $n_{\text{eq}} \cdot \text{cm}^{-2}$ using the "non-ionizing energy loss" model (NIEL) that asserts that damage to the bulk of the sensor only depends on the kinematic energy impacted to the bulk lattice. This assumption has been tested extensively [36].

It has been estimated that the total dose over the silicon-strips sensors over the duration of the experiment at the HL-LHC will be about 5–10 $\times 10^{14}$ $n_{\text{eq}} \cdot \text{cm}^{-2}$, depending on the distance from the beam pipe [37]. The total dose arises from about 10% protons, 30% charged pions, and 60% neutrons.

3.2 Sensor types and fluences

Figure 3.1 shows the sensor-carrying board used for all measurements presented. It contains two 2-mm connectors for high voltage and ground supply. The pitch adapter that is bonded to the sensor is used to connect to the readout chip used in the measurements of charge collection presented in section 3.5. Two traces in the right connect to two pins that are used in the inter-strip resistance measurement shown in section 3.6.

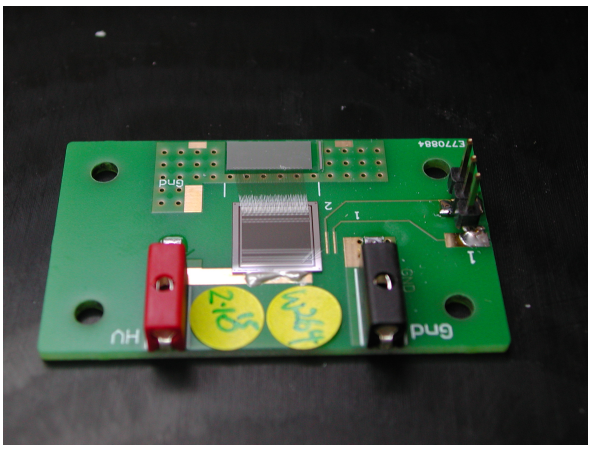

Figure 3.1: Photograph of a sensor-carrying board with sensor included. Wire-bonds connect the ground plug to the bias ring of the sensor; DC pads of each strip to the pitch adapter; and DC pads to the pin connectors. The backplane of the sensor is glued to a metal strip that is biased at large (negative) voltage. The miniature sensor is about 1×1 cm^2.

Given that the backplane of the sensor is glued to the sensor board, it produces a direct contact. An alternative is to use wire-bonds to the top of the sensor that provide an ohmic contact through the edge of the sensor. This yields an unwanted voltage drop across the edge that effectively reduces the applied voltage to the backplane. Given that the edge resistivity changes with radiation damage, this drop is practicably unpredictable, and therefore sensors are biased through the backplane directly.

The sensors studied are prototype sensors of the A12 and A07 design described in chapter 2. These sensors were irradiated with 70 MeV proton beams at the CYRIC cyclotron at Tohoku, Japan; and in the TRIGA nuclear reactor in Ljubljana, Slovenia.

Unless specified otherwise, all sensors are annealed at 60°C for 80 minutes. This is done to standardize measurements in different institutes and neglect the possible accidental annealing at room temperature during transport or irradiation. All measurements presented are performed inside a fridge at $T = -20$°C.

Table 3.1 shows a summary of all sensors studied in this work, their irradiation type, and total dose. As discussed in chapter 2, the A07 and A12 designs are similar; however, they were fabricated with wafers of different resistivity[1]. The comparison between the performance degradation of the proton and neutron-irradiated sensors is a proxy to compare surface and bulk damage (protons produce more surface damage).

Table 3.1: Sensors studied in this work, their irradiation type, dose, and irradiation facility.

Sensor design	Irradiation type	Dose [$n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$]	Irradiation facility
A12	protons, 70 MeV beam	5×10^{14}	CYRIC
A07	protons, 70 MeV beam	5×10^{14}	CYRIC
A07	protons, 70 MeV beam	2×10^{15}	CYRIC
A07	neutrons, reactor	5×10^{14}	TRIGA
A07	neutrons, reactor	1×10^{15}	TRIGA
A12	neutrons, reactor	1×10^{15}	TRIGA
A12	none	—	—

3.3 Leakage current

Setup

The setup for leakage current measurements consists of a printed-circuit board that contains a low-pass filter and pins that can be plugged to the sensor board. The measurements are performed with an SMU controlled by a LabVIEW script. Figure 3.2 shows a photograph of the setup used.

The leakage current is measured as a function of bias voltage in the range[2] from 0 to -700 V with the same protocol used for the large sensors, as described in section 2.2.4. The measured leakage current for the A12 reference sensor is within the noise level of the instruments used (< 1 nA). This is expected because the A12 large-area sensors at room temperature draw a leakage current of $O(100$ nA), as shown in figure 2.9, and the area of the small sensors is about 100 times smaller.

[1] As a consequence, the V_{FD} is about 200 V for A07 sensors [38], and about 370 V for A12 sensors (as shown in chapter 2).

[2] This range is chosen to minimize the chance of micro-discharges of the irradiated sensors.

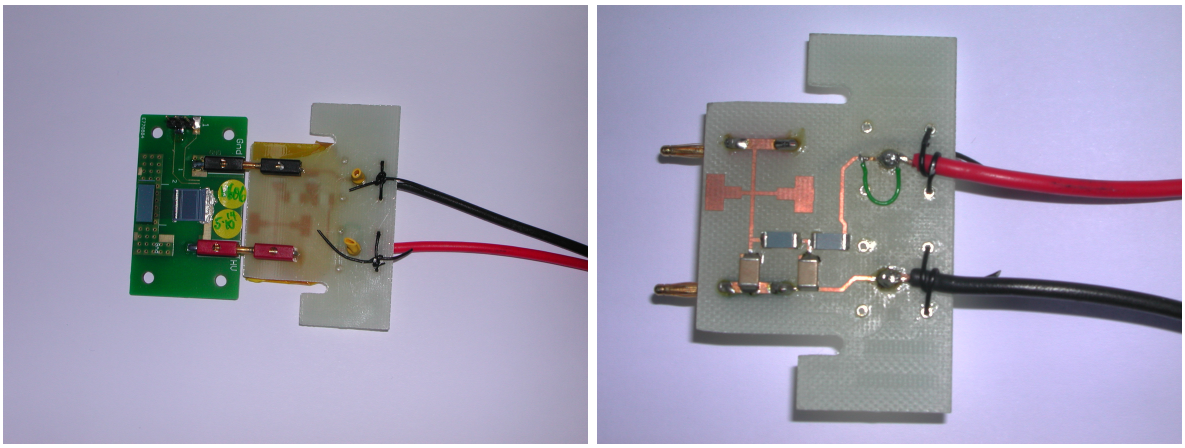

Figure 3.2: Setup for leakage current measurements. The printed-circuit board that contains a low-pass filter and soldered cables that are connected to a source measuring unit. The measurements are performed at $T = -20^{\circ}\mathrm{C}$ in a fridge with controlled humidity.

Figure 3.3 shows the leakage current as a function of voltage for different irradiated sensors. The data show a strong dependence on the bias voltage and total dose, as expected. The higher the dose, the higher the leakage current. At 600 V, the sensor irradiated with a dose of 2×10^{15} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ draws about 28 μA; the A07 (A12) sensor irradiated with 10^{15} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ draws about 21 μA (11 μA); and the sensors irradiated with 5×10^{14} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ draw about 6–8 μA depending on the sensor.

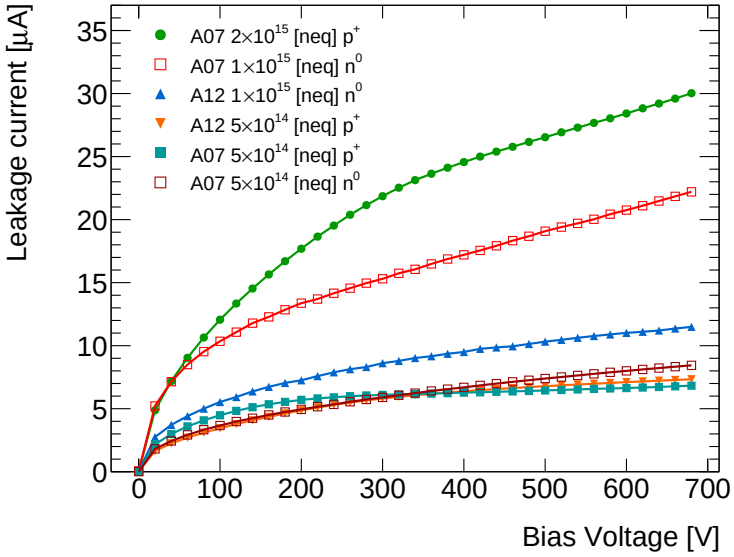

Figure 3.3: Leakage current as a function of bias voltage for various irradiated sensors. The measurements are performed at $T = -20^{\circ}\mathrm{C}$. The lines join the data points to guide the eye. The instrument error is smaller than the marker size.

The measurements of A07 sensors irradiated with protons and neutrons with the same equivalent dose of 5×10^{14} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ agree within 1–2 μA depending on the bias voltage. This supports the NIEL hypothesis[3]. However, the data of the A07 and A12 sensors irradiated with 10^{15} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ differ by about a factor of two. This difference is not understood but it may be related to differences in the total dose (the uncertainty on the dose is about $\pm 10\%$ [2]).

In summary, all sensors show leakage currents of the order of $O(10)$ μA. This is several orders of magnitude larger than observed for unirradiated sensors, but still within the technical specifications.

3.4 Bulk capacitance

Setup

Figure 3.4 shows the setup used to measure the bulk capacitance of miniature sensors. The plug-in board used contains a printed-circuit board with bias connections and a DC-blocking capacitor of about 100 nF. The 4-wires are connected to a LCR meter controlled by LabVIEW script. Measurements are performed at different bias voltages from 0 to -700 V following the same protocol described in section 2.2.4.

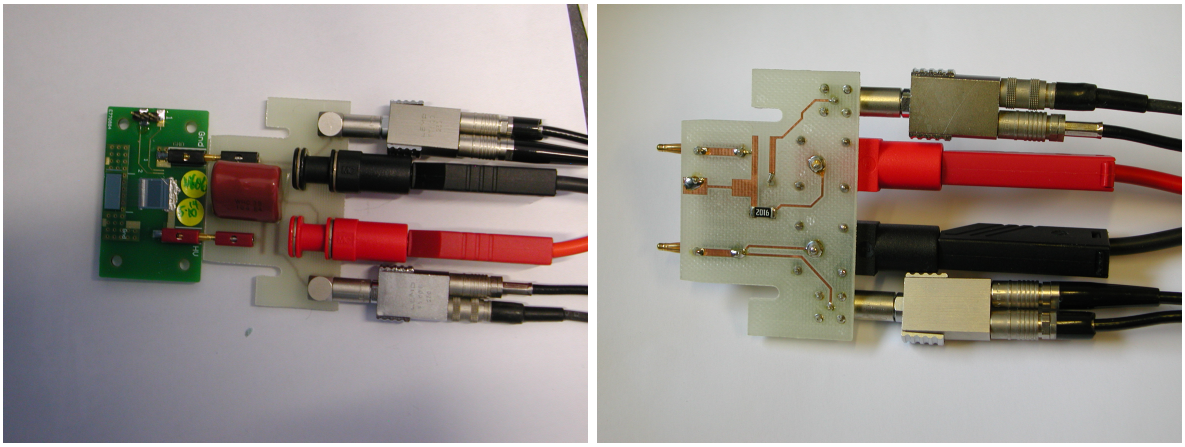

Figure 3.4: Setup for capacitance measurements. The 4-mm connectors provide the bias voltage provided by a source measurement unit. The connection is made through a blocking capacitor. The four wires connected through the LEMO connectors are linked to the LCR.

[3]This assumes that the bulk current, and not surface current, dominates the total leakage current.

Results

Figure 3.5 shows the inverse bulk capacitance squared as a function of bias voltage for proton-irradiated sensors and a reference sensor.

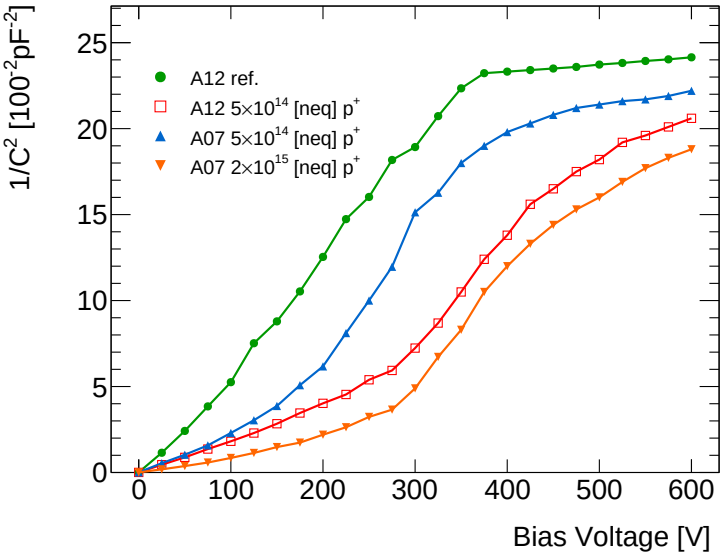

Figure 3.5: Inverse bulk capacitance square, $1/C^2$, as a function of bias voltage for an A12 reference sensor and proton-irradiated sensors. The measurements are performed at $T = -20^{\circ}\mathrm{C}$. The lines join the data points to guide the eye. The instrument error is smaller than the marker size.

The A12 reference sensor shows the expected behaviour; the V_{FD} extracted with the interception of two linear fits is 362 ± 9 V. In contrast, the results for proton-irradiated sensors show a strong non-linearity. This suggest that the assumption of depletion volume growing with the square-root of the bias voltage no longer holds, and thus this measurement does not provide a way to estimate V_{FD}. Alternatively, the RC series model presumed by the LCR might no longer hold given the charged insulation layer present in irradiated sensors. This feature has been reported in other studies of the same sensor design and similar irradiation doses [39].

3.5 Charge collection measurement

Setup

The experimental setup includes a ≈ 100 MBq strontium-90 source[4] and a scintillator crystal coupled to a photomultiplier. The latter are used to trigger events. The data acquisition is performed with the Alibava system [40] that consists of a daughterboard, a motherboard, and a software interface. The daughter-board contains two analogue readout chips capable of handling up to 256 channels at a clock speed of 40 MHz.

The mini sensor is connected to one of the readout chips with wire-bonds via pitch adaptors. The daughterboard is connected to the motherboard via a flat cable. The motherboard provides signal discrimination with the photomultiplier signal as input, and micro-processed control unit connected to a PC via a USB optical cable[5]. A high level software allows the control of the system.

The setup includes a voltage supply to bias the sensor through the high voltage connections in the daughterboard. A 5 V supply is used to power the Alibava motherboard, and the photomultiplier. The setup is placed inside a fridge cooled at $T = -20^{\circ}$C. Figure 3.6 shows a photograph of the setup. Humidity is reduced by circulating nitrogen gas at a rate of about 10 litres per minute. Temperature is monitored with multiple sensors. Studies of temperature monitoring (with an infra-red camera) show no large temperature gradients between the sensor and the environment.

3.5.1 Measurements of reference sensor

Measurements are performed with bias voltage in the range $-$(100–700) V in 100 V steps. For each voltage setting, 10^5 events triggered with the photomultiplier signal[6] are recorded. In addition, for every run a dedicated "pedestal" run is performed; in these runs 10^4 events are recorded randomly to obtain pedestal and noise values. The sequence is controlled by a PC and LabVIEW script.

[4] The strontium-90 decay chain yields 2.2 MeV shared among an electron, antineutrino and stable nuclei.

[5] To provide electrical insulation.

[6] A threshold of 100 mV is about 5 times the noise standard deviation.

Figure 3.6: Photograph of the experimental setup used for the measurement of charge collection of miniature sensors. It includes the Alibava readout bonded to the sensor and the lead enclosure under which the strontium-90 source, scintillator, and photomultiplier tube are located.

The pedestal value is obtained, for each channel, as the simple average of the signal in the randomly-triggered data. It lies in the range 460–470 ADC counts for most channels. This value is near the middle of the dynamic range of the ADC (1024 counts maximum). This feature of the readout system accommodates signals of either polarity.

The measured charge distribution (after pedestal subtraction for each channel) in one typical randomly-triggered sample is shown in figure 3.7. Most of the data for most channels lie in the range 0 ± 20 ADC counts; this corresponds to the noise of the readout system.

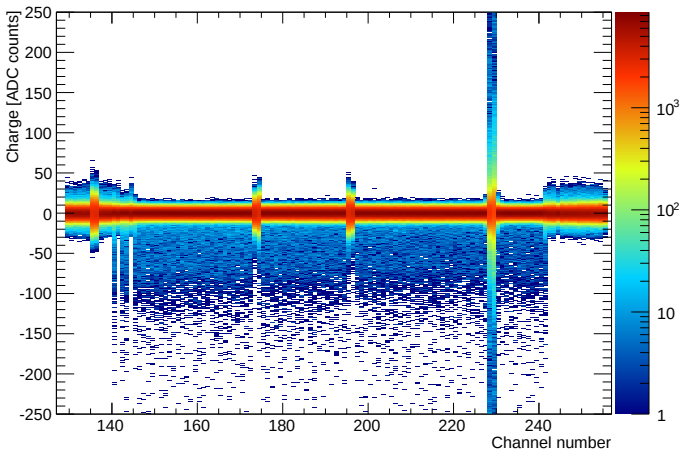

Figure 3.7: Charge distribution as a function of readout channel number measured with an A12 reference sensor in randomly-triggered events. The pedestal values per each channel have been subtracted from the data. The colour axis represents the number of events.

The charge distribution also shows a tail that extends to large negative values. These are events that contain signal, i.e. βs from strontium-90, by coincidence. The A12 sensor collects electrons, so the signal has negative polarity. Some of the first and last channels do not present this tail. This is because some readout channels are not bonded to any strip (the A12 sensor has 104 strips and the readout chip has 128 channels).

Figure 3.8 shows the charge distribution for a typical channel (number 150). The data show a noise distribution and a tail due to signal that is about 3 orders of magnitude lower than the noise peak. A Gaussian fit describes well the data. For each channel, the noise is taken as the σ of the fit. The impact of the negative tail on the fit is estimated by restricting the fitted data to the range 0 ± 15 ADC counts; this does not produce a statistically significant impact on the result.

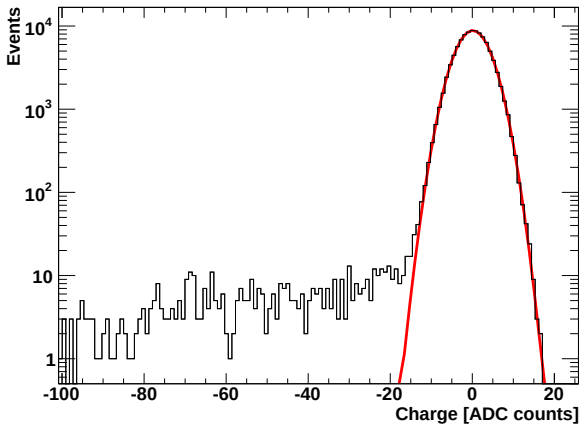

Figure 3.8: Charge distribution for a typical readout channel in randomly-triggered events. The negative tail of the distribution extends beyond the range shown. The result of a Gaussian fit is superimposed on the data as a red line.

The noise for each channel is shown in figure 3.9 for different bias voltages. Most of the channels present a noise level of about 4 ADC counts. The relative statistical uncertainty of each measurement is a few per mil and thus is neglected. The data for different bias voltages differ at the level of 5%. This illustrates the need to take "pedestal data" for each bias voltage.

Some channels present noise at much higher levels (> 6 ADC counts). These come in groups of two or more consecutive channels. This suggest that this noise is due to a short circuit among the corresponding wire-bonds. These noisy channels are not considered in the following analysis.

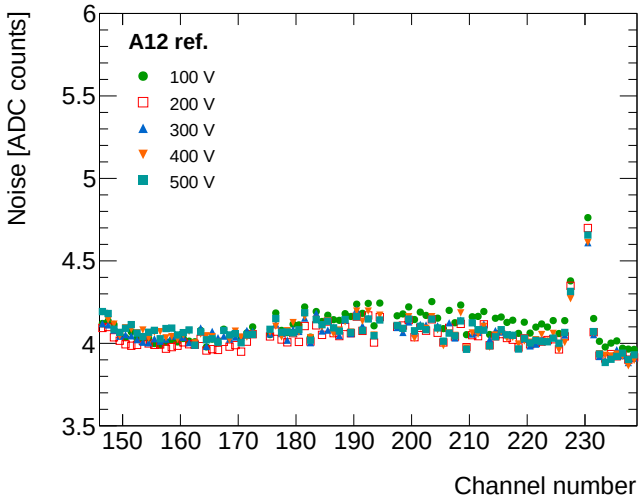

Figure 3.9: Noise for each readout channel for a reference A12 sensor biased at different voltages. The statistical uncertainties of the measurements are smaller than the marker size. Some data points lie above the vertical scale shown.

Figure 3.10 shows the time dependence of the averaged maximum signal (i.e. the charge in the channel with largest signal-to-noise ratio) at different bias voltages. In all cases the distribution peaks at about 30 ns and has a tail that extends to $t \approx 100$ ns. The small rise at $t \approx 0$ is not understood but does not impact the results.

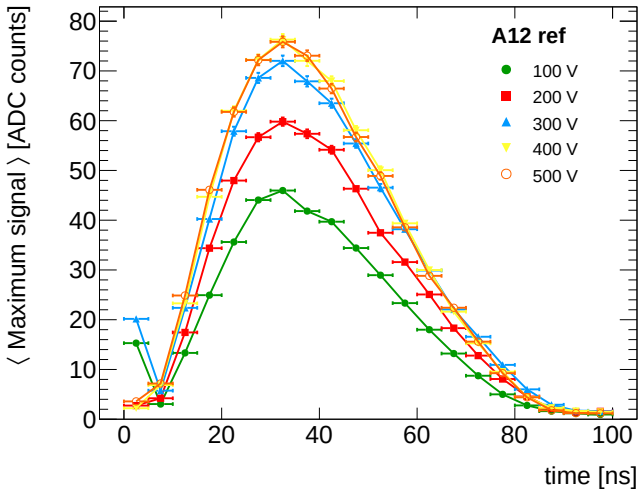

Figure 3.10: Average maximum signal as a function of time for a reference A12 sensor biased at different voltages. The error bars, which in most cases are smaller than the marker size, represent statistical uncertainty only. The lines join different data points to guide the eye.

The arrival of the signal that triggers the event is randomly distributed so the observed distribution corresponds to the pulse shape created by the pre-amplifier of the readout system. The pulse shape is optimized for maximum charge collection[7] as it has a broad peak and relatively slow decay of about 100 ns.

The peak value of the time distribution is strongly affected by the bias voltage. This is due to an increase of the depletion volume with voltage, which increases the measured charge. The peak average charge rises from about 45 ADC counts at 100 V to about 60 ADC counts at 200 V and 70 ADC counts at 300 V. The signal increase from 300 V to 400 V shows an increase that is smaller, only about 5 ADC counts. The distribution at 400 V and 500 V are consistent with statistical uncertainties. This pattern suggests that the sensor has already reached full depletion, as expected from bulk capacitance measurements.

Only events with time in the range 30 ± 5 ns are considered for further analysis. This time window corresponds to approximately the broad peak of the time distribution (by judicious choice of the parameters of the readout chip).

Figure 3.11 shows the hit distribution for an A12 reference sensor biased at different voltages. Here a "hit" corresponds to the channel with the highest signal-to-noise ratio. Most of the channels registered 40–80 events. The slope in the distribution is attributed to the alignment between the source and sensor.

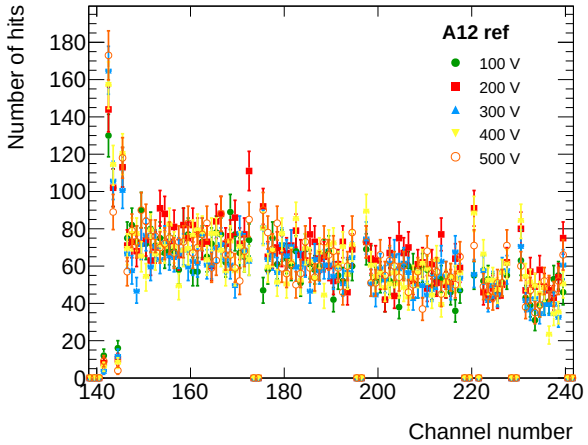

Figure 3.11: Number of recorded events as a function of readout channel number with the largest signal-to-noise ratio. The data correspond to a measurement of a reference A12 sensor biased at different voltages. The error bars represent statistical uncertainty only.

[7] In the real LHC operation the peaking time is set to about 10 ns.

The hit distribution is similar for different bias voltages. This is expected because all channels are biased in the same manner through a common bias ring. The channels with zero events corresponds to masked noisy channels. There are also three channels that present an anomalously small number of events. The charge distributions of these channels in randomly-triggered events do not have the tail attributed to signal events (see figure 3.7), i.e. they are not bonded. Thus, those are events where the largest signal-to-noise ratio was produced by a noise fluctuation. These channels are not considered in further analysis.

For each event, the hit channel is used to define a "seed" for a clustering algorithm. Only events in which the seed-channel charge is larger than 3 σ are considered; here σ is the noise for that channel. Then, the neighbouring strips[8] (considered independently on both sides) are added to the cluster if they registered a charge above 1.8 σ. This procedure is repeated to next-to-neighbouring strips (and so on) until the tested strip is not selected. The selection criteria applied to the seed and neighbouring channels (3 σ and 1.8 σ) have been chosen considering the impact of noise and the available statistics. These cuts are consistently used by the different groups working on this research [2].

Figure 3.12 shows the cluster size distribution for the A12 reference sensor biased at different voltages. When the sensor is biased at 100 V, the distribution peaks at 1 channel and has a tail that reaches ≈ 4 channels. As the voltage increases, the distributions shift to larger values. The distribution at 400 V and 500 V are very similar, as expected.

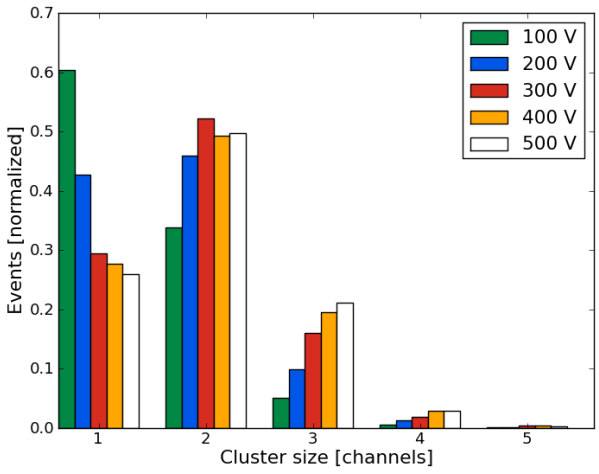

Figure 3.12: Cluster size distribution for a reference A12 sensor biased at different voltages.

[8]If the neighbour strip is a masked channel it is ignored and the next strip is considered.

The voltage dependence of the cluster size distribution is attributed to a threshold effect given that at lower voltages the depletion volume is smaller and the signal produced in the neighbouring channels[9] is not enough to pass the 1.8 σ selection of the clustering algorithm. The larger the depletion volume, the larger the generated charge and thus the chance that the neighbouring channels would pass the selection.

Figure 3.13 shows the average cluster size as a function of bias voltage for the A12 reference sensor. The cluster size grows with increasing bias voltage from about 1.5 channels at 100 V to about 2 channels at 400 V. Then it saturates and stays constant at ≈ 2 channels. This saturation behaviour suggests that the observed cluster size is proportional to the depletion depth.

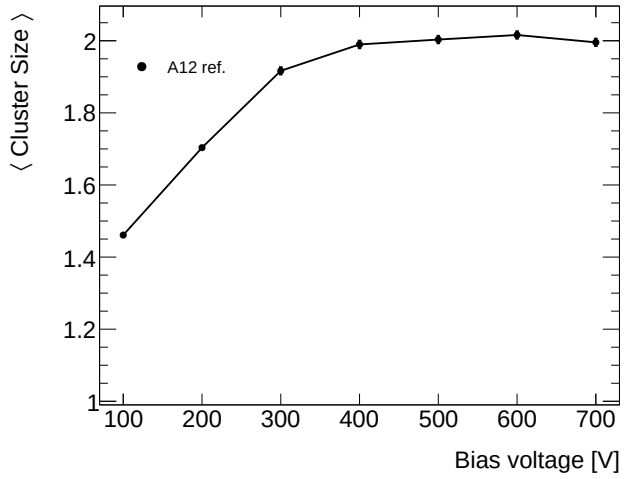

Figure 3.13: Average cluster size as a function of bias voltage for a reference A12 sensor. The error bars, which in some cases are smaller than the marker size, represent statistical uncertainty only. The line joins different data points to guide the eye.

For each event, the sum of the total charge measured in all the strips in the cluster is defined as the cluster charge. Figure 3.14 shows the total cluster charge for the A12 reference sensor biased at different voltages. In all cases the distribution shows a Gaussian-like peak and a tail towards positive values. The tail is expected due to the inherent Landau distribution of energy loss via ionization of minimum-ionizing[10] particles.

[9] The βs from strontium-90 produce signal in multiple channels not by diffusion in the silicon, which is negligible, but rather due to angular dispersion.

[10] Simulations [41] of a setup similar to the one used in this work show that the distribution of energy loss of βs from strontium-90 is similar to a Landau function.

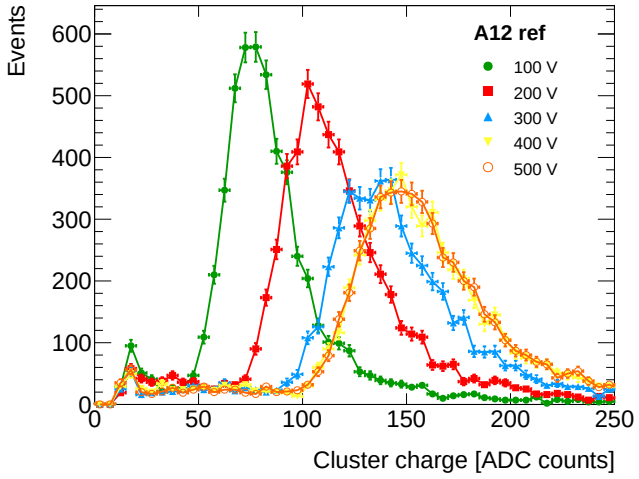

Figure 3.14: Cluster charge distribution for a reference A12 sensor biased at different voltages. The error bars represent statistical uncertainty only. The lines join different data points to guide the eye.

The increase in bias voltage shifts the cluster charge distributions toward higher values and increases their width. Both effects can be attributed to an increase of the depletion depth (that effectively leads to a convolution of Landau distributions). The distributions also show a flat negative tail that is indicative of noise or non minimum ionizing particle (mip) energy loss.

The most probable value of the cluster charge distribution, obtained with a Gaussian fit in a window of ± 30 ADC counts around the maximum value, is taken as a figure of merit of the charge collected by the sensor and is called "most probable charge" or simply "collected charge". Figure 3.15 shows the collected charge as a function of bias voltage for the A12 reference sensor. The data show a clear increase from ≈ 80 to 150 ADC counts from 100 V to 400 V and then it remain constant, as expected.

The collected charge is calibrated from arbitrary ADC counts to a number of electrons by assuming that the maximum collected charge (plateau value) is equivalent to what a minimum-ionizing particle would deposit in the fully-depleted volume of the sensor, according to the Bichsel[11] model [42]. The expected charge for a sensor that has an active[12] depth of 301 μm is 23.1 ke. The simple average of the four measurements at or above 400 V is 147 ± 0.1 ADC counts. Therefore, the conversion factor of 0.157 [ke/ADC counts] is used. The statistical uncertainty of this factor is neglected.

[11] This model predicts a slightly lower charge compared to the Landau model.
[12] Extracted from measurements of capacitance and the area enclosed by the bias ring [34].

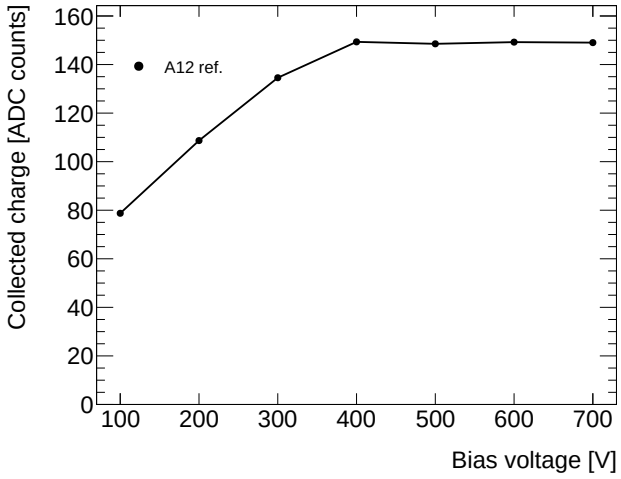

Figure 3.15: Collected charge as a function of bias voltage for a reference A12 sensor. The statistical uncertainty of the data is smaller than the marker size. The lines join different data points to guide the eye.

The scale factor derived from the measurements of the A12 reference sensor is used to calibrate the measurements of irradiated sensors that are shown in the following sections. This is justified because the irradiated sensors are measured with the same readout system and under the same experimental conditions.

3.5.2 Measurements of neutron-irradiated sensors

Figure 3.16 shows the collected charge as a function of bias voltage of different neutron-irradiated sensors[13] compared with the result for the A12 reference sensor. The collected charge from the irradiated sensors is smaller than the reference result by about 5–10 ke depending on the sensor and bias voltage. The curves of the irradiated sensors do not present the discontinuity in the slope that is attributed to the full-depletion of the sensor. An extrapolation using their slope suggests that their V_{FD} is larger[14] than 1000 V.

The results for the A12 and A07 sensors irradiated at the same dose of 10^{15} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ differ by less than 0.5 ke for all bias voltages. Given that these sensors were fabricated using silicon wafers with very different resistivity, this indicates that their performance is mostly dictated by the radiation damage induced to the bulk of the silicon.

[13]The data point at 100 V for the A07 sensor irradiated at 5×10^{14} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ was found to be corrupted and thus is not shown.

[14]This assumes that at full-depletion voltage the irradiated sensors would measure the same charge as unirradiated sensors. This may not be the case due to charge trapping and other effects.

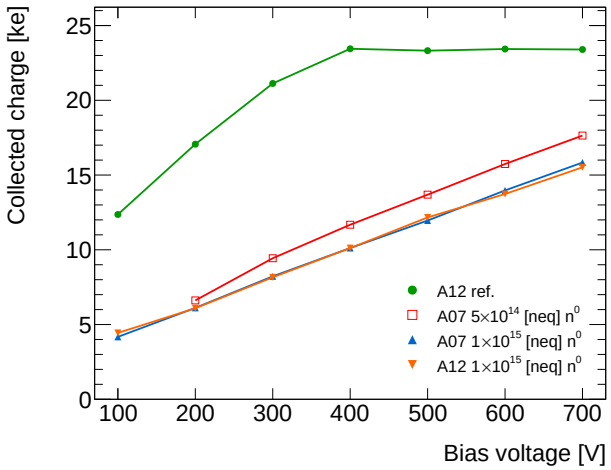

Figure 3.16: Collected charge as a function of bias voltage for different neutron-irradiated sensors. The A12 reference results are shown for comparison. The statistical uncertainty of the data is smaller than the marker size. The lines join different data points to guide the eye.

The A07 sensor irradiated with a dose of 5×10^{14} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ shows higher collected charge than the sensors irradiated at 10^{15} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$, by about 1–3 ke depending on the bias voltage. At 600 V, the collected charges are about 16 ke for the sensors irradiated with 10^{15} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ and about 18 ke for the sensor irradiated at 5×10^{14} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$.

Figure 3.17 shows the bias-voltage dependence of the average cluster size for neutron-irradiated sensors compared with the A12 reference results. The data for the irradiated sensors are significantly smaller than the reference values for all voltages.

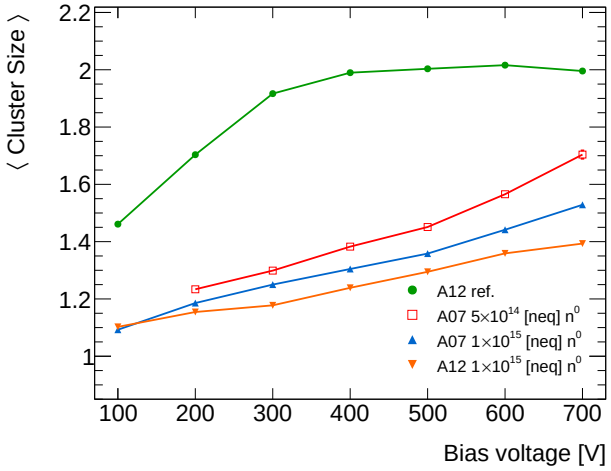

Figure 3.17: Average cluster size as a function of bias voltage for different neutron-irradiated sensors. The A12 reference results are shown for comparison. The statistical uncertainty of the data is smaller than the marker size. The lines join different data points to guide the eye.

For all the irradiated sensors the average cluster size grows with bias voltage in a roughly linear fashion. The slope is different to that observed in the data of the A12 reference sensor before full-depletion. At 600 V, the average cluster size is about 1.4–1.5 channels for the sensors irradiated at 10^{15} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ and about 1.7 channels for the sensor irradiated at 5×10^{14} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$.

Figure 3.18 shows the correlation between the measured average cluster size and collected charge at different bias voltages. The data of all sensors show a very strong positive correlation. The data of the A07 sensors show similar slope, even though the charge and cluster measurements differ significantly at a fixed voltage. This indicates that while the full-depletion voltage increases with received dose, the inherent relation between cluster size and collected charge (both proportional to the depletion volume) seems to remain stable. That is, the sensor charge collection is still dominated by drift, not diffusion or local charge amplification.

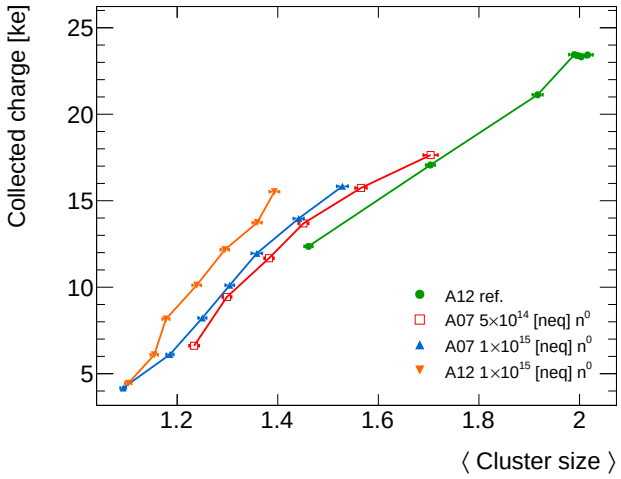

Figure 3.18: Correlation between average cluster size and collected charge measured at different bias voltages (ranging from 100 V to 700 V) and different neutron-irradiated sensors. The A12 reference results are shown for comparison. The lines join successive bias voltages to guide the eye. The error bars represent statistical uncertainty only.

The correlation for the A07 and A12 sensors irradiated with the same dose of 10^{15} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ are different. Given that the collected charge at the same bias voltage is similar (see figure 3.16), this indicates that the relation between collected charge and average cluster size differs. This is probably due to the difference in the geometry and concentration of the p-stops used in the two designs.

3.5.3 Measurements of proton-irradiated sensors

Figure 3.19 hows the collected charge as a function of bias voltage of different proton-irradiated sensors compared with the A12 reference results.

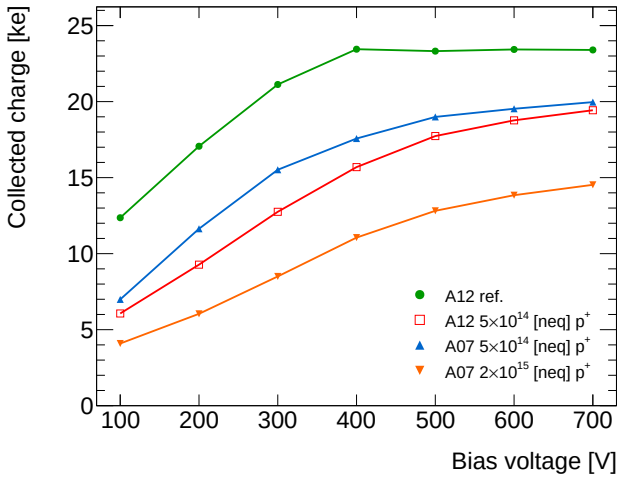

Figure 3.19: Collected charge as a function of bias voltage for different proton-irradiated sensors. The A12 reference results are shown for comparison. The statistical uncertainty of the data is smaller than the marker size. The lines join different data points to guide the eye.

The collected charge in the proton-irradiated sensors is about 3–10 ke lower than the reference sensor depending on the voltage. The A07 and A12 sensors irradiated with the same dose of 5×10^{14} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ differ by about 2–3 ke depending on the bias voltage. This differs from the agreement observed in the results of neutron-irradiated sensors (see figure 3.16). This may be related to uncertainties on the total dose, that in proton beams are more uncertain than for neutron irradiation in nuclear reactors.

The data show a monotonic rise with bias voltage but is much less linear than for the neutron-irradiated sensors. This renders the estimate of the full-depletion voltage (extrapolating the data at higher voltages) more difficult. However, given that there is no sign of saturation in the measured range, 700 V is a lower limit for V_{FD}.

Figure 3.20 shows the average cluster size as a function of bias voltage for proton-irradiated sensors. The data show a monotonic rise with bias voltage and no sign of saturation. The average cluster size is larger than the saturation value of the A12 reference sensor at bias voltages larger than 400–600 V depending on the sensor. This phenomenon is not observed in the data for neutron-irradiated sensors (see figure 3.17). Thus, it is likely to be related to surface damage of the sensor rather than bulk damage.

This is further investigated in the measurements of inter-strip resistance presented in section 3.6.

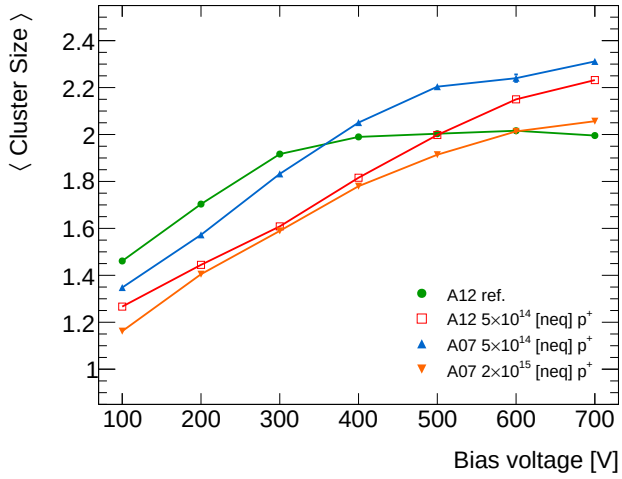

Figure 3.20: Average cluster size as a function of bias voltage for different proton-irradiated sensors. The A12 reference results are shown for comparison. The statistical uncertainty of the data is smaller than the marker size. The lines join different data points to guide the eye.

Figure 3.21 shows the correlation of the average cluster size and the collected charge measured at different bias voltages for proton-irradiated sensors compared with the A12 reference results. The data for all irradiated sensors lie below the curve given by the A12 reference sensor. That is, for a given collected charge the average cluster size is larger. This is the opposite to what is observed in the neutron-irradiated results (see figure 3.18). This is qualitatively consistent with expectations of surface and bulk damage.

The data of the A07 and A12 sensors irradiated with the same dose of 5×10^{14} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ show a very similar trend. This is the same behaviour observed for the neutron-irradiated sensors at 10^{15} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$. This indicates that the apparent differences in the voltage-dependence of the charge and cluster size (see figures 3.20 and 3.19) are due to an overall shift in effective doping. The A07 sensor irradiated with 2×10^{15} $n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ shows a different slope, i.e. for the same collected charge the cluster size is larger. This suggest increasing surface damage with increasing dose, as expected.

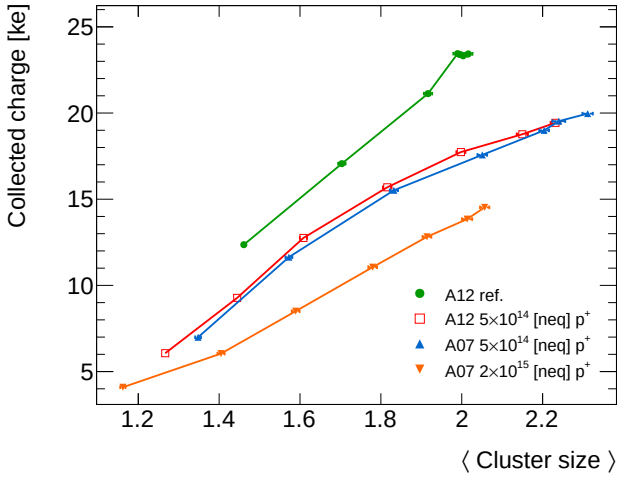

Figure 3.21: Correlation between average cluster size and collected charge measured at different bias voltages (ranging from 100 V to 700 V) and different proton-irradiated sensors. The A12 reference results are shown for comparison. The lines join measurements at successive bias voltages to guide the eye. The error bars represent statistical uncertainty only.

3.5.4 Annealing studies

Measurements of collected charge after controlled annealing time are presented. Only neutron-irradiated sensors are considered. The aim of these measurements is to characterize the time-scale associated with the "beneficial" and "reverse" annealing that have been reported in previous studies of irradiated sensors [43]. The same procedure described in section 3.5.1 is repeated after each controlled annealing session. The accumulated annealing times at 60°C ranged from 10 to 4440 minutes.

Figure 3.22 shows the annealing-time dependence of the collected charge (measured at 600 V) for neutron-irradiated sensors. For total annealing times less than 250 minutes, the data show an increase in all cases of about 10–20% depending on the sensor, i.e. beneficial annealing. At later times a clear decrease of collected charge is observed, i.e. reverse annealing. These results are consistent with measurements performed with sensors irradiated with neutrons and protons at similar fluences [2].

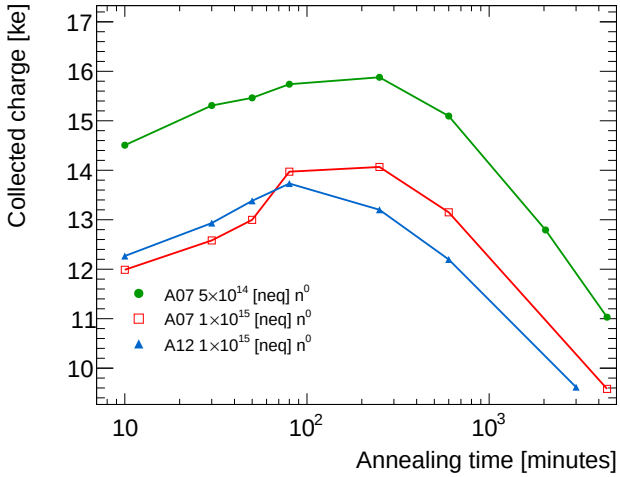

Figure 3.22: Collected charge as a function of annealing time at 60°C for different neutron-irradiated sensors biased at 600 V. The statistical uncertainty of the data is smaller than the marker size. The lines join the data points to guide the eye.

Figure 3.23 shows the annealing-time dependence of the average cluster size (measured at 600 V). The data show a qualitatively similar behaviour to the collected charge result (see figure 3.22). However, the relative increase of the average cluster size is smaller (about 3–5% depending on the sensor). The relative decrease observed is about 10–20%.

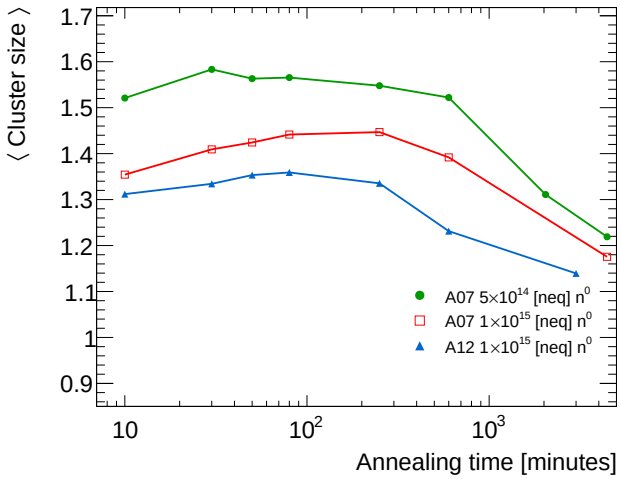

Figure 3.23: Average cluster size as a function of annealing time at 60°C for different neutron-irradiated sensors biased at 600 V. The statistical uncertainty of the data is smaller than the marker size. The lines join the data points to guide the eye.

Figure 3.24 shows the correlation between the average cluster size and the collected charge measured at different annealing times. In all cases the lowest annealing time corresponds to the right end of each curve. In the period of beneficial annealing both the average cluster size and the collected charge increase but afterwards decrease with annealing time. These two observables do not evolve in exactly the same manner with annealing time. This is illustrated by the shape of the correlation curve.

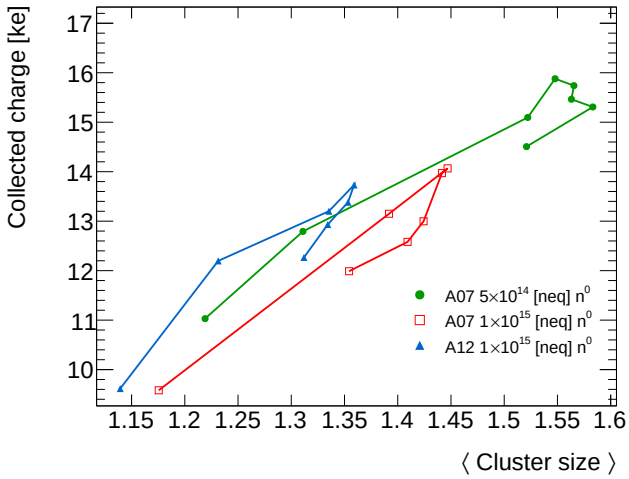

Figure 3.24: Correlation between average cluster size and collected charge measured at 600 V and at different annealing times (ranging in 10–4440 minutes) for different neutron-irradiated sensors. The statistical uncertainty of the data is smaller than the marker size. The lines join the data points of successive annealing times to guide the eye.

3.6 Inter-strip resistance

This section describes measurements of inter-strip resistance, R_{IS}, for irradiated sensors. As described in section 2.2.2, the technical specifications for the A12 strip sensors require that the inter-strip resistance should be at least 10 times the bias resistance (that is, about 10–20 MΩ). This is required to ensure a good strip isolation, which ensures good spatial resolution of the sensor.

Setup

The inter-strip resistance is measured by applying a test voltage to a single strip and measuring the induced voltage in a neighbour strip, from which R_{IS} can be inferred (as

described later). Typically this measurement is performed in a probe station with needles that provide the contact to the strips (as described e.g. in Ref. [44]).

While the probe station described in chapter 2 would be suitable for this measurement, this lacks a cold chuck that would be necessary in order to study irradiated sensors. Therefore, the measurements presented in this work are performed in a fridge with a sensor-carrying board (shown in figure 3.1) that has two traces connected to pins. These traces are bonded to the DC pads of two neighbouring strips, as shown in figure 3.25.

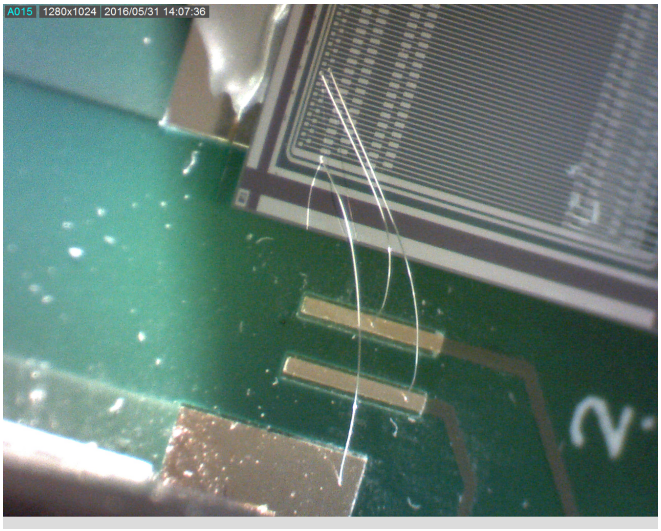

Figure 3.25: Photograph of part of the sensor-carrying board used to measure inter-strip resistance. It shows the wire-bonds that connect the DC pads of two neighbouring strips to the traces in the board. Also visible is the wire-bond that connects the sensor bias ring and the ground pad.

One of the pins in the sensor board is connected to a source measuring unit (control strip) and the other to a high-impedance electrometer (test strip). A separate plug-in board, shown in figure 3.26, is used to bias the sensor and connect the ground of all measurements involved to the same reference point. A 100 kΩ resistor is used to avoid interference between the strip measurements and the leakage current of the sensor. Figure 3.27 illustrates the equivalent circuit of the setup.

Only two strips are considered in the network for simplicity but the inter-strip resistance result is not strongly affected by all the other strips that are connected in parallel. This "two-needle" measurement yields similar results[15] than one that uses two test strips ("three-needle approach"), used e.g. in Ref. [44].

[15]The use of measurements of three strips enhances the sensitivity of the measurement.

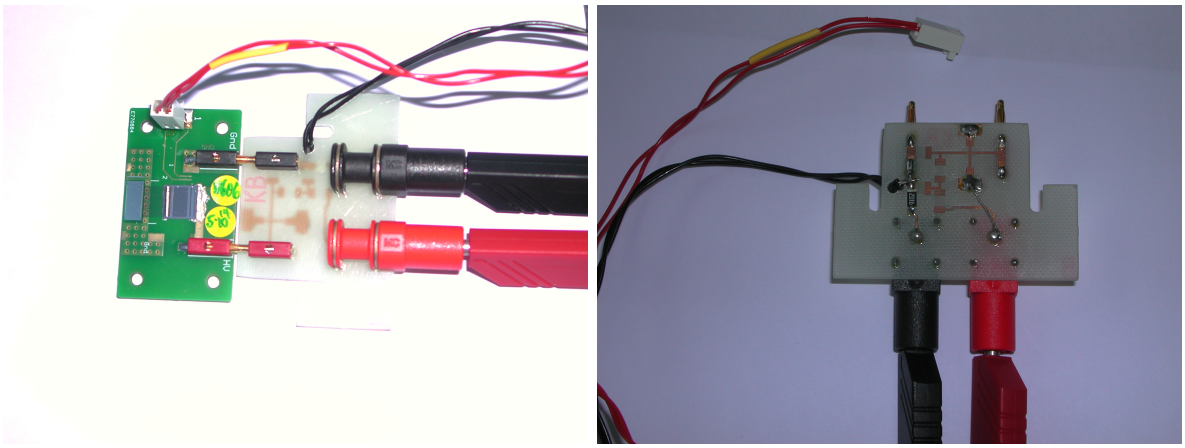

Figure 3.26: Photographs showing the setup used for inter-strip resistance measurements. The red cable is connected to two pins that are connected to traces wire-bonded to two strips. The black cables are the ground cables for the strip measurements; these are connected to the bias ring through a 100 kΩ resistor in the back of the plug-in board (shown in the right panel).

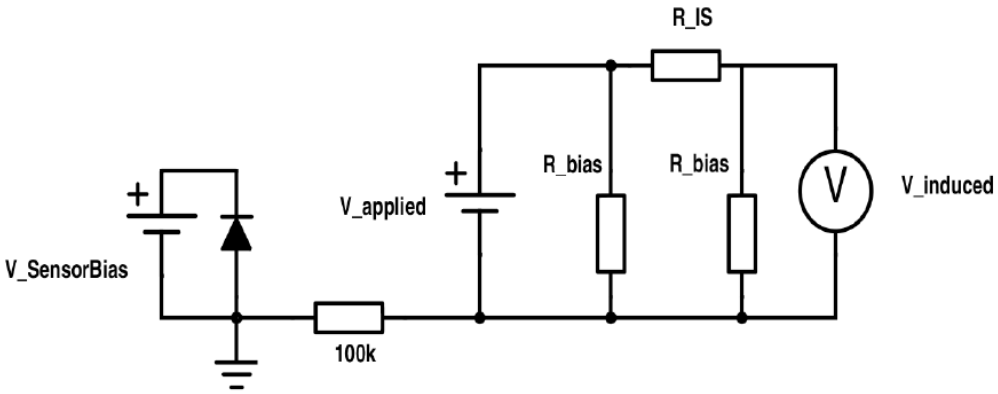

Figure 3.27: Diagram of the electrical network involved in the inter-strip resistance measurement. For simplicity only two strips (two bias resistors) are drawn. The reverse-biased diode represents the sensor leakage current.

The following voltages are defined:

- V_{applied}: corresponds to the voltage applied to the DC pad of the control strip relative to the grounded bias ring.
- V_{induced}: corresponds to the measured voltage that is induced in the test strip, relative to the grounded bias ring.
- V_{bias}: corresponds to the bias voltage applied to the backplane of the sensor, relative to the grounded bias ring.

A source measuring unit is used to supply a V_{applied} in the range $\pm$ 2 V. The resulting voltage V_{induced} is measured with a high-impedance electrometer. The measurement is repeated five times to test the reproducibility of the results. The R_{bias} is determined by measuring the voltage dependence of the current drawn by the test strip[16]. The solution to the circuit network shown in figure 3.27 yields the following relation:

$$\Delta V_{\text{induced}} = \Delta V_{\text{applied}} \left(\frac{R_{\text{bias}}}{R_{\text{bias}} + R_{\text{IS}}} \right). \tag{3.1}$$

The experimental setup presented in this section is validated by measuring a known resistance. Two 1 MΩ resistors ("bias resistors") are soldered to a resistance of 500 MΩ ("inter-strip resistance") and connected to the plug-in board via crocodile clips. A reverse-biased diode is used to model the leakage current of an unirradiated sensor. The measured R_{IS}, obtained with equation 3.1, agrees with that expected.

Results of unirradiated sensors

Figure 3.28 shows the measured current and induced voltage on the control and test strip respectively as a function of applied voltage (on the test strip). The sensor is biased at 400 V. In the applied voltage range of $\pm$ 2 V the current shows a very linear behaviour and reaches about ± 1 μA. This corresponds to $R_{\text{bias}} = 1.9 \pm 0.1$ MΩ, which is consistent with the technical specification of 1.5 ± 0.5 MΩ.

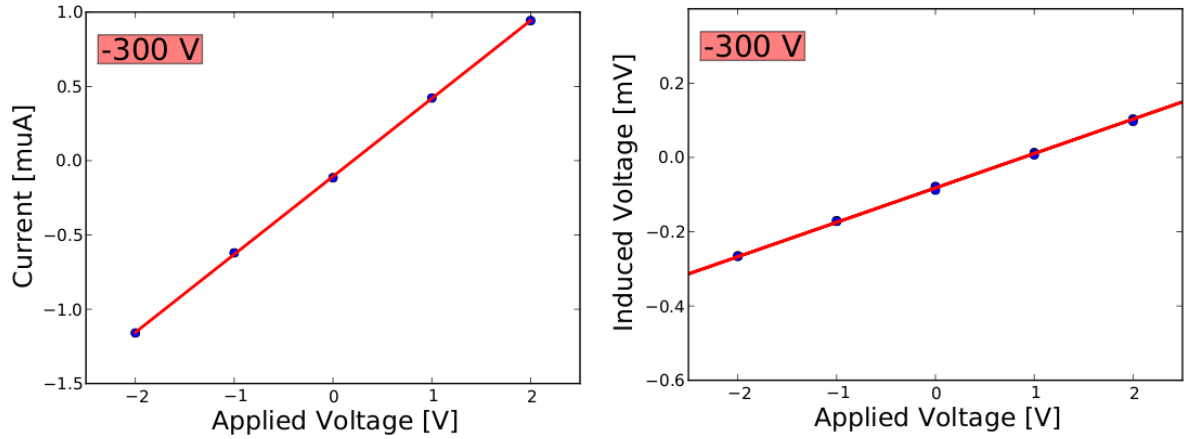

Figure 3.28: Current drawn by the control strip (left panel) and induced voltage on the test strip (right panel) as a function of applied voltage on the control strip. The data correspond to a A12 reference sensor biased at 300 V. The red lines show the results of linear fits.

[16]This is a first-order approximation that assumes that R_{bias} is much smaller than R_{IS}.

The induced voltage on the test strip is well described by a linear fit. The R_{IS} is determined to be about 20 GΩ. This value is close to the input impedance of the electrometer used in the measurement. Therefore, the R_{IS} of the sensor is much larger and cannot be determined with this setup. Similar results were reported in Ref. [44].

Results of irradiated sensors

Figure 3.29 shows the measured induced voltage on the test strip for a proton-irradiated sensor biased at different voltages. In all cases, the data are well described by a linear fit. The offset of the linear fit is much larger than for the unirradiated sensor and it grows with bias voltage. This is attributed to a voltage drop caused by the leakage current of the sensors. The $\Delta V_{\mathrm{induced}}$ observed are much larger than for the reference sensor; this indicates much lower values of R_{IS}. The corresponding values of R_{bias} do not vary strongly with irradiation dose and meet the technical specifications.

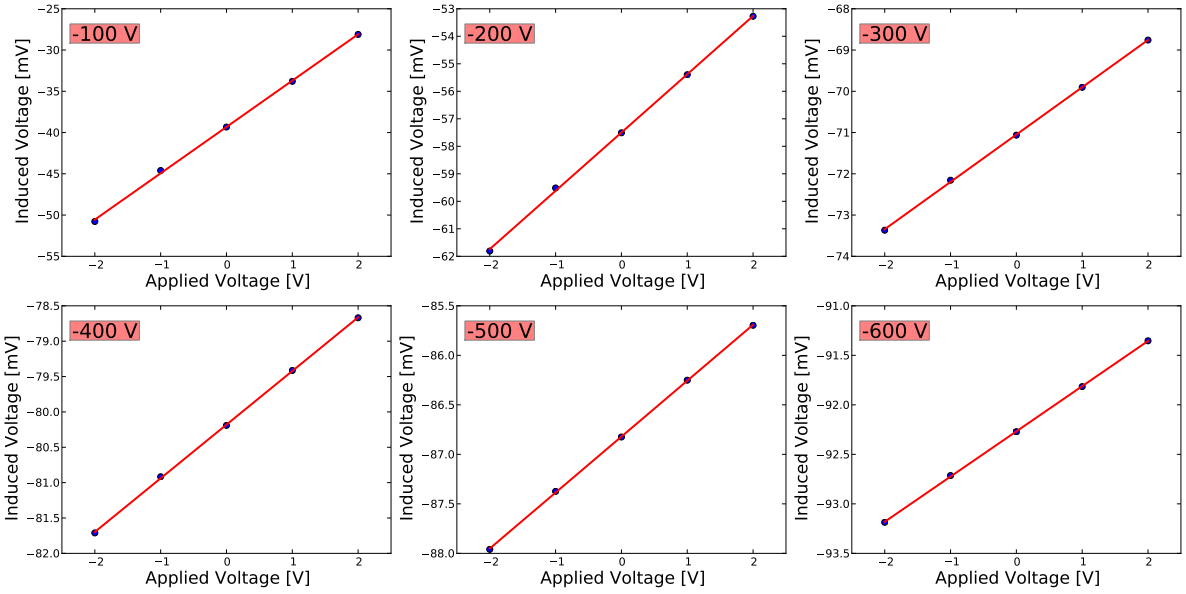

Figure 3.29: Induced voltage on the test strip as a function of applied voltage on the control strip. The data correspond to a proton-irradiated A12 sensor. The red lines show the results of linear fits.

Figure 3.30 shows the R_{IS} as a function of bias voltage for proton and neutron-irradiated sensors. For all sensors, the R_{IS} grows approximately linearly with bias voltage. This can be attributed to a decrease of the resistance caused by the increase of the

depletion volume (that does not saturate for the measured voltage range, as suggested by collected charge measurements).

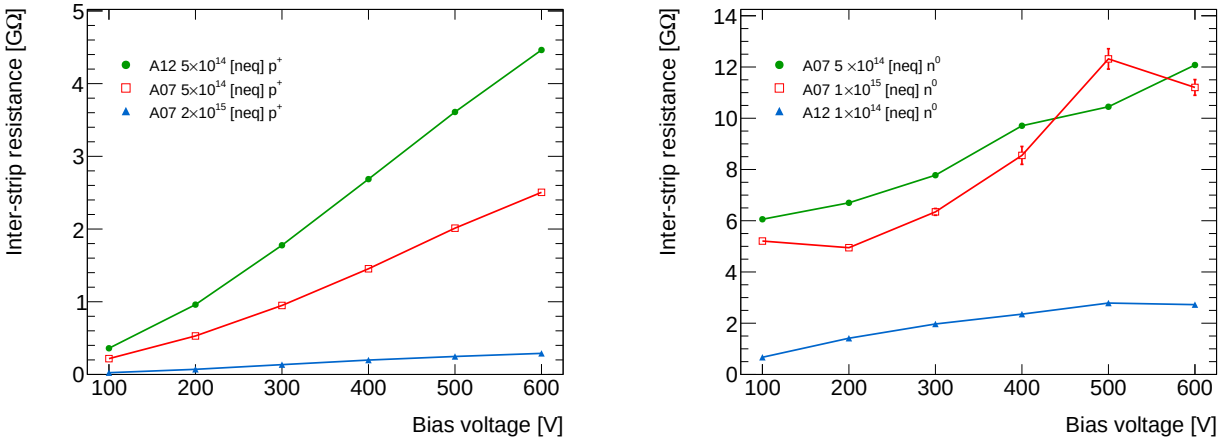

Figure 3.30: Inter-strip resistance as function of bias voltage for proton and neutron-irradiated sensors, shown in the left and right panel respectively. The error bars, which are typically smaller than the marker size, represent the standard deviation of repeated measurements. The lines join the data points to guide the eye.

The measured R_{IS} values depend strongly on the total dose and type of irradiation. The proton-irradiated sensors show much smaller R_{IS} values than the neutron-irradiated sensors with the same equivalent dose of $5 \times 10^{14}\ n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$. Moreover, the A07 and A12 sensor irradiated with protons with $5 \times 10^{14}\ n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$ differ significantly; at 600 V the R_{IS} values are 2.5 GΩ and 4.5 GΩ respectively.

Given that the proton-irradiated A07 and A12 sensors show very similar leakage current (as shown in figure 3.3) and thus bulk-damage, the difference of R_{IS} values may be attributed to the differences in design of strip isolation. The opposite behaviour is observed with the A07 and A12 sensors irradiated with neutrons and total dose of $10^{15}\ n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$. In this case, the A12 sensor R_{IS} values are much lower, reaching about 3 GΩ at 600 V. This is only a quarter of the A07 result. The reason for this is not understood.

The lowest measured R_{IS} corresponds to the sensor irradiated with the highest dose of protons, $2 \times 10^{15}\ n_{\mathrm{eq}} \cdot \mathrm{cm}^{-2}$. This is qualitatively the expected behaviour. While the R_{IS} value at 600 V is 300 MΩ and thus only a small fraction of the R_{IS} value of unirradiated sensors, it still complies with the technical specifications.

Figure 3.31 shows the correlation between the measured inter-strip resistance and average cluster size at different bias voltages for different proton and neutron-irradiated sensors.

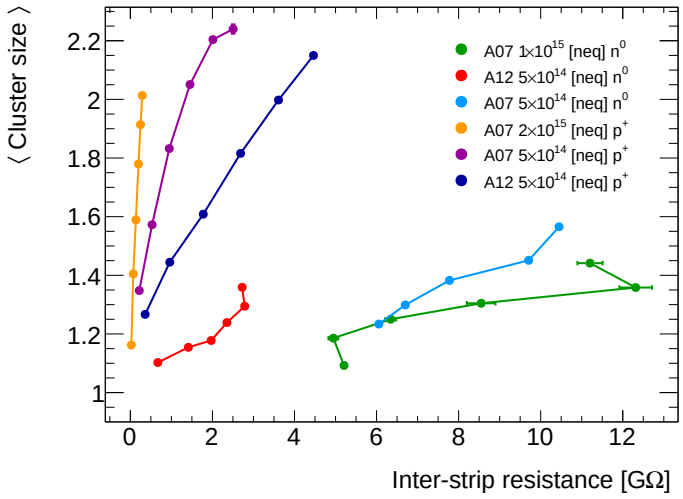

Figure 3.31: Correlation between the measured inter-strip resistance and average cluster size measured at different bias voltages for different irradiated sensors. The lines join data points to guide the eye.

The data of all sensors show a strong positive correlation. The proton-irradiated sensors are clustered on the upper-left side of the plot, while the neutron-irradiated sensors tend to be on the lower-right. This demonstrates that the surface and bulk damage alter these quantities differently. The results of A07 and A12 sensors irradiated with the same equivalent dose show that the differences in strip-isolation design do have a significant impact.

3.7 Summary

The radiation damage in silicon-strip sensors has been studied by measuring bulk and surface dependent quantities such as leakage current, depletion voltage, collected charge, and inter-strip resistance. In all cases the radiation damage altered the nominal design values in various degrees depending on the total irradiation dose and type of irradiation.

The capacitance and collected charge measurements reveal that the full-depletion voltage of the sensors will become larger than the envisaged operational voltage of 600 V. Nevertheless, the collected charge with sensors that were irradiated with doses that

are equivalent to what is expected over the lifetime of the experiment suggest that the signal-to-noise ratio will be 10 or more.

Moreover, the inter-strip resistance measurements suggest that the impact on strip-isolation is within tolerable limits. This suggests that the degradation of tracking performance will be limited. In summary, the results presented in this work and other similar studies [2, 44, 45] show that the prototype silicon-strip sensors that will be used in the new ATLAS inner detector will endure the radiation damage expected at the HL-LHC.

The inter-strip resistance and cluster size, and collected charge measurements may provide valuable input to ongoing and future simulation studies [46–48] that aim to understand the detailed processes caused by non-ionizing radiation in n-in-p sensors, and in particular the effect of the strip-isolation structure used in their design and the interplay between surface and bulk damage.

Chapter 4

Luminosity measurements with ATLAS

This chapter describes luminosity measurements with the ATLAS detector at the LHC. The methodology used follows previous ATLAS analyses of proton–proton collisions[49, 50]. The results presented in this work include the absolute luminosity calibration for the following three runs: the 2013 proton–lead run at 5 TeV; the 2015 proton–proton run at 13 TeV; and the 2015 lead–lead run at 5 TeV. The accuracy aimed at and obtained for each of these measurements varies due to the different physics requirements and details of experimental conditions.

4.1 Introduction

The observed rate of an arbitrary reaction in scattering experiments is given by the product of the corresponding cross-section and the experiment's luminosity. The luminosity depends exclusively on the details of the experimental setup and it reflects its ability to produce reactions.

For collider experiments the expression for the luminosity is involved but for the case of a single bunch pair circulating at speed $v \approx c$ and colliding with no crossing angle is given by[1]:

$$\mathcal{L} = f_{\mathrm{r}} n_1 n_2 \int \hat{\rho}_1(x,y)\hat{\rho}_2(x,y)\mathrm{d}x\mathrm{d}y, \tag{4.1}$$

[1]Note the z-coordinate (along the beam axis) has been integrated out, i.e. this assumes bunch sizes in z that are negligible compared to the time scales used to define luminosity.

where f_r is the revolution frequency; n_1 and n_2 are the number of particles (or ions) in each bunch; and $\hat{\rho}_1, \hat{\rho}_2$ are the normalized particle densities in the transverse plane [51].

The integral in equation 4.1 gives the inverse of the "effective overlap area" of the collisions. If $\hat{\rho}_1$ and $\hat{\rho}_2$ are equal to the product of two Gaussian functions with widths σ_x and σ_y, equation 4.1 reduces to:

$$\mathcal{L} = \frac{f_r n_1 n_2}{4\pi\sigma_x\sigma_y}. \tag{4.2}$$

The LHC ring diameter of 26.7 km sets f_r at 11.2 kHz, the typical n_1, n_2 are close to 1.1×10^{11} protons and nominally $\sigma_x \approx \sigma_y \approx 14$ μm ($\beta^* = 0.55$ m, $\epsilon_{\text{norm}} = 3.75$ $\mu\text{m} \cdot \text{rad}$), which yields a nominal per-bunch luminosity of about 6×10^{30} cm^{-2}s^{-1}.

However, the LHC produces collisions with a crossing angle to minimize unwanted parasitic collisions. The side effect is that the luminosity is reduced by the effective reduction of the overlap of the beams. The reduction factors depends on the crossing angle but for the typical values used at the LHC (nominally 290 μrad) it is 0.83 [52], so the nominal per-bunch luminosity is about 5×10^{30} cm^{-2}s^{-1}.

The LHC can store up to 2808 bunches separated by 25 ns. The total luminosity is the sum of all the per-bunch luminosities. During the 2011 run, there were 1380 bunches stored with a minimum spacing of 50 ns and the LHC broke the world record of collider luminosity with a value of 4.7×10^{32} cm^{-2}s^{-1}, surpassing the 4×10^{32} cm^{-2}s^{-1} record set by the Tevatron a year earlier [53]. In 2016, the LHC reached its nominal luminosity level of 10^{34} cm^{-2}s^{-1} operating with a minimum bunch spacing of 25 ns.

Benefits of high luminosity

The luminosity is one of the most important figures of merit of the performance of an experiment and is normally maximized. A high luminosity would allow an experiment to access previously unmeasured small cross-sections or to achieve higher statistical precision of a given measurement.

One of the main goals of the LHC experiments is to search for and discover new phenomena beyond the standard model. In 2015 the LHC delivered proton–proton collisions with a center-of-mass energy of 13 TeV, reaching almost its design value of 14 TeV. But the discovery potential this energy provides can only be exploited with high luminosity as it enables lower cross-sections to be accessed.

It has been argued [54] that luminosity and center-of-mass energy are related because cross-sections of new phenomena characterized by an energy scale E scale roughly as E^{-2}; this demands that an increase in center-of-mass energy must be matched by a quadratically larger increase in luminosity to achieve similar rates, and thus discovery potential[2]. For future accelerators with a center-of-mass energy of the order of 100 TeV this argument would demand a luminosity increase of a factor of 50. However, more sophisticated estimates suggest that the requirements on the luminosity may be less, depending on the model of new physics that is considered [55].

Drawbacks of high-luminosity

The number of inelastic proton–proton collisions per bunch crossing ("pile-up") follows a Poisson distribution with mean μ given by:

$$\mu = \frac{\mathcal{L}\sigma_{\text{inel}}}{f_{\text{r}}}, \tag{4.3}$$

where σ_{inel} is the proton–proton inelastic cross-section, which at 13 TeV is about 80 mb [7]. So, for the design per-bunch luminosity of 5×10^{30} cm^{-2}s^{-1} the expected pile-up is $\mu \approx 25$. The LHC already surpassed its design value reaching values of $\mu \approx 40$ during 2012.

The LHC design considers the storage of thousands of bunches in the ring to keep the per-bunch luminosity at the 5×10^{30} cm^{-2}s^{-1} level (and thus $\mu \approx 25$), while increasing the total luminosity to the 10^{34} cm^{-2}s^{-1} level. This requires a minimum bunch spacing of 25 ns. This is a benefit only if the detector and electronics used in the experiments can disentangle separate bunches, i.e. run at 40 MHz.

While in principle ATLAS and CMS are designed to handle even the highest luminosity levels that the LHC was expected to deliver, the high pile-up conditions inevitably reduces the performance of the detectors. This affects measurements not limited by statistical uncertainties such as those of electroweak bosons and the inelastic proton–proton cross-section. For these cases, dedicated low-luminosity runs have been performed to better control systematic uncertainties [7, 56].

[2] This was accomplished in the LHC design that contemplated a luminosity about 200 times higher than that of its predecessor the Tevatron, while the center-of-mass energy is only 7 times more.

4.2 Luminosity measurements and calibration

Motivation for precise measurements

The main motivation for luminosity calibrations is to enable measurements of cross-sections that provide an input to theory. In cases where the theoretical calculations are well understood, such as for electroweak processes, precise cross-section measurements are a probe for new phenomena, as these would alter the expectations from known physics[3].

Absolute luminosity measurements are also relevant to estimate background rates for searches for new phenomena because typically these have known cross-sections. The requirements for accuracy in these cases are less strict because normally other uncertainties dominate the sensitivity of the searches. Eventually, the full characterization of new phenomena (if discovered) would require absolute luminosity calibrations.

A luminosity calibration with an accuracy of 1–3% has been achieved at the LHC and now it is deemed necessary for the success of the physics program of all experiments. For example, several results like $t\bar{t}$ production measurements in the dilepton channel have a percent level accuracy that is dominated by the uncertainty of the luminosity scale [57]. In other cases, like the W or Z boson cross-section, the measurements can rival the accuracy of the state-of-the-art theoretical calculations [58, 59].

Moreover, accurate cross-section measurements at the LHC can be used to extract parton density functions that would in turn help to improve theoretical calculations for other processes like Higgs boson production and a number of fundamental parameters such as the W-boson and top-quark masses.

Measurement techniques

The most complicated part of luminosity measurements is the beam overlap integral (see equation 4.1). The transverse density of the beams can be measured, for example, with the "wire-scan" technique that passes a fine wire (diameter $O(1)$ μm) through the beam at high speed ($O(1)$ m/s) and measures the secondary particles produced. By studying the time distribution of the signal (typically measured by scintillators coupled to photomultipliers) the relative transverse density of the beam can be reconstructed.

[3] While ratios of cross-sections bypass the need for luminosity calibrations, absolute measurements provide additional information.

The wire-scan technique enables a measurement of the beam size (width of Gaussian distribution) with micron-level accuracy and it is routinely applied at the LHC during the injection phase for monitoring purposes [60].

The main limitation of the use of the wire-scan technique for luminosity measurements is that the beams are squeezed with quadrupole magnets before each collision takes place in the interaction points of the experiments. A direct measurement with the wire-scan technique close enough to the interaction region is not feasible because it would interfere with the main experiments.

Practically, the wire-scan technique can be used to measure the transverse densities away from the interaction point (tens, or hundreds of meters away). These measurements can be extrapolated to the interaction region with a precise knowledge of the magnetic fields in the accelerator. This procedure should be done independently for both beams, then the beam-overlap integral can be calculated numerically.

Another alternative for absolute luminosity calibrations in colliders is to rely on measurements of a well-understood process with accurate cross-section calculations. This has been the default method of luminosity calibration for all lepton colliders and even lepton–hadron colliders to date. Bhabha scattering ($ee^+ \to ee^+$) has been used successfully in lepton colliders [61, 62], and $ep \to ep\gamma$ at the lepton–hadron collider HERA [63] with an accuracy of $O(1)\%$ or less. In principle the same method can be applied to hadron colliders; however, there is no corresponding process known with such accuracy. For example, the accuracy of calculations of W or Z cross-sections is about 5%, limited by the uncertainty on the parton density functions [64].

The first hadron collider was the CERN intersecting storage rings (ISR); it started operations in 1971 and produced proton collisions with a center-of-mass energy up to 62 GeV. At the time, to achieve a precise luminosity calibration was a priority because one of the high-profile measurements expected at this facility was the inelastic proton–proton cross-section at the new energy frontier[4]. In this context, Simon van der Meer proposed a new technique that has become one of the standards in hadron colliders since then.

[4] As described in more detail in chapter 5, the motivation was to search for a predicted asymptotic behaviour in which all cross-sections would converge to a constant with increasing energy.

4.2.1 The van der Meer method

In 1968, Simon van der Meer proposed a technique to measure the beam overlap integral without having to know the individual beam densities. In an internal CERN document [65], he outlined what is today known as as the "van der Meer technique". He started by stating that for beams colliding with a relative displacement, Δ, the observed mean rate of inelastic interactions should be given by:

$$\mu(\Delta) = C \int \hat{\rho}_1(x-\Delta)\hat{\rho}_2(x)\mathrm{d}x, \tag{4.4}$$

where C is a constant that depends on the efficiency and acceptance of the detector used to measure the rate. His key insight was to propose a "scan", i.e. to displace the beams by different Δ values, to measure the following integral:

$$\int_{-\infty}^{\infty} \mu(\Delta)\mathrm{d}\Delta = C \int_{-\infty}^{\infty}\int_{-\infty}^{\infty} \hat{\rho}_1(x-\Delta)\hat{\rho}_2(x)\mathrm{d}x\mathrm{d}\Delta \tag{4.5}$$

$$\int_{-\infty}^{\infty} \mu(\Delta)\mathrm{d}\Delta = C \int_{-\infty}^{\infty}\left(\int_{-\infty}^{\infty} \hat{\rho}_1(x-\Delta)\mathrm{d}\Delta\right)\hat{\rho}_2(x)\mathrm{d}x = C, \tag{4.6}$$

which is equal to the unknown C. Therefore, the ratio of the rate for head-on collisions ($\Delta = 0$) divided by the integral of the rate over the beam displacements is equal to the beam-overlap integral:

$$\frac{\mu(\Delta=0)}{\int_{-\infty}^{\infty}\mu(\Delta)\mathrm{d}\Delta} = \int \hat{\rho}_1(x)\hat{\rho}_2(x)\mathrm{d}x \equiv \frac{1}{2\sqrt{\pi}\Sigma_{\mathrm{X}}}. \tag{4.7}$$

Here Σ_{X}, called convolved beam size, is defined for convenience. For the trivial case where $\hat{\rho}_1$ and $\hat{\rho}_2$ are the same Gaussian function, the Σ_{X} corresponds to the standard deviation. However, equation 4.7 is general and applies to an arbitrary $\hat{\rho}_1$ and $\hat{\rho}_2$.

While the original proposal dealt with only one dimension[5], the method can be trivially extended to two dimensions. In this case the beam overlap integral is given by:

$$\frac{\mu(\Delta_x=0,\Delta_y=0)}{\int_{-\infty}^{\infty}\mu(\Delta_x,\Delta_y)\mathrm{d}\Delta_x\mathrm{d}\Delta_y} = \int \hat{\rho}_1(x,y)\hat{\rho}_2(x,y)\mathrm{d}x\mathrm{d}y. \tag{4.8}$$

This together with a measurement of the beam currents enables a direct measurement of the luminosity according to equation 4.1.

[5]Simon van der Meer proposed his method while working in the ISR, which used a continuous beam and large crossing angle (14.8°) so only one direction was relevant.

In the beam scans (also known as vdM scans), the beams are moved in a certain number of steps, then data are recorded at each step for enough time to obtain a statistically significant measurement.

In principle, a precise measurement of $\int_{-\infty}^{\infty} \mu(\Delta_x, \Delta_y) \mathrm{d}\Delta_x \mathrm{d}\Delta_y$ would require measurements in a two-dimensional grid in the x–y plane. In practice, due to time constraints, each vdM scan consists of two separate one-dimensional scans. Subsequently, the beam-overlap integral is calculated under the assumption that the beam densities in the transverse plane factorize, i.e. $\hat{\rho}(x, y) = \hat{\rho}_x(x)\hat{\rho}_y(y)$. This leads to:

$$\int \hat{\rho}_1(x, y)\hat{\rho}_2(x, y)\mathrm{d}x\mathrm{d}y = \int \hat{\rho}_{1x}(x)\hat{\rho}_{2x}(x)\mathrm{d}x \int \hat{\rho}_{1y}(y)\hat{\rho}_{2y}(y)\mathrm{d}y. \tag{4.9}$$

The direct measurement of the luminosity with the vdM method is only possible under the very special conditions of the beam scans which are not compatible with normal LHC running. However, the method provides a way to calibrate the detectors used to measure the relative luminosity during normal conditions, as explained in the following section.

Calibration with the vdM method

The key to calibrating the luminosity for normal running conditions is to measure simultaneously the absolute luminosity with the vdM method and the relative luminosity with a dedicated detector or "luminometer". The fraction of the total inelastic rate, μ, that is measured with a given luminometer is given by the "visible rate" μ_{vis}:

$$\mu_{\mathrm{vis}} = \epsilon\mu, \tag{4.10}$$

where ϵ encapsulates the detector acceptance and efficiencies. As μ is given by the product of luminosity and inelastic cross-section, equation 4.10 can be used to define the "visible cross-section" σ_{vis}:

$$\mu_{\mathrm{vis}} = \mathcal{L}(\epsilon\sigma_{\mathrm{inel}}) \equiv \mathcal{L}\sigma_{\mathrm{vis}}. \tag{4.11}$$

A measurement of σ_{vis} can be obtained with a simultaneous measurement of μ_{vis} and $\mathcal{L}$ during the vdM scans. Given that σ_{vis} is a constant, it can be used to translate the visible rate to luminosity during normal physics conditions using equation 4.11. The σ_{vis} for each luminometer is referred as its "calibration constant".

4.2.2 Beam-gas imaging method

The beam-gas imaging method is a novel technique [66] for absolute luminosity calibration used by the LHCb experiment [67, 68]. It is based on a measurement of the individual transverse density of the beams, at the interaction region, obtained by studying the spatial distribution of reconstructed vertices in beam-gas interactions.

LHCb has been the only experiment that has used this method because of its unique capabilities: its high-resolution vertex detector, called VELO; and its gas injection system, called SMOG. The latter is used to inject inert gases into the interaction region.

Beam-gas interactions are measured by triggering on unpaired bunch crossings. One of the key ingredients of the measurement is the determination of the vertex resolution in beam-gas events. For the LHCb VELO, this is 15–100 μm depending on the longitudinal coordinate of the vertex and the number of reconstructed charged particles.

One of the main advantages of this method over the traditional vdM method is that the two-dimensional overlap integral is extracted without the factorization assumption. Indeed, LHCb measurements showed for the first time that this assumption could lead to biases at the level of 3%, depending on the scan.

The systematic uncertainty of this method is comparable to that obtained with the vdM method, i.e. $O(1\%)$, and is dominated by the uncertainty on the vertex resolution in beam-gas events that is about 1%. The combination of the beam-imaging and vdM methods allowed LHCb to achieve a luminosity scale uncertainty of 1.2%. This is the most precise luminosity measurement obtained at the LHC [68].

4.3 Luminosity detectors

ATLAS exploits redundancy of luminometers based on different technologies to control possible instrumental effects. The primary luminometers must be able to distinguish among colliding bunches because averaging out results leads to biases caused by the non-linearity of Poisson statistics (as explained in section 4.4). Another vital requirement for the luminometers is that their readout must be independent of the ATLAS trigger system[6] to avoid dead-times that can bias the luminosity measurement. The two main ATLAS sub-detectors that satisfy these requirement are the BCM and LUCID.

[6]The ATLAS central processing trigger system vetoes further triggers while a trigger is being processed.

BCM

BCM stands for beam conditions monitor and, as the name suggests, it was designed to monitor the LHC beams for safety during operations. Its primary goal is to provide a fast signal of beam loss to trigger an automatic beam dump. It consists of two arms located at $z = \pm 1.85$ m; each arm consists of four 1×1 cm^2 diamond detectors located in the beam pipe support structure at $r = 55$ mm.

Figure 4.1 illustrates the BCM system. The readout of the four detectors in each side is divided in two pairs, named BCM-H (horizontal) and BCM-V (vertical). These are treated as independent luminometers, although their responses are expected to be similar.

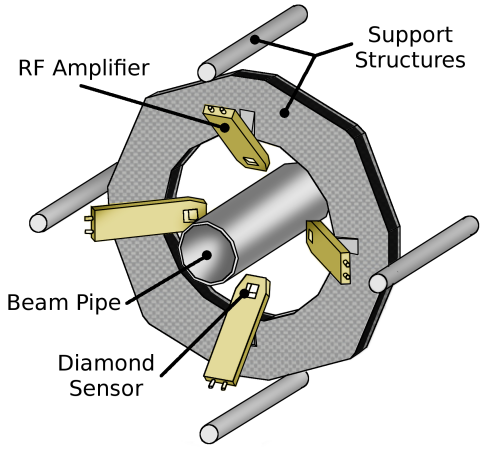

Figure 4.1: Illustration of the BCM system. Source: Ref. [9].

Artificial diamond was chosen as a detector material because its large band gap provides the radiation tolerance needed for continuous operation at $r = 55$ mm without dedicated cooling (in contrast with silicon sensors). Charged particles are detected via measurement of the ionization charge produced in the diamond lattice. It provides single-particle detection capability and sub-ns time resolution.

Beam loss is monitored by measuring secondary particles produced by collisions between the beams and the collimators that protect the ATLAS detector. Secondary particles from beam loss can be distinguished from particles from real collisions by exploiting the timing information of the two BCM arms: a real collision would produce a simultaneous signal in both arms whereas secondary particles would be recorded with a time difference of ≈ 12.3 ns (3.70 m/c).

The small BCM acceptance is not a problem for luminosity monitoring; on the contrary, it helps to prevent saturation up to higher pile-up. The BCM was the primary luminometer during the 2012 proton–proton runs [50].

LUCID

LUCID is a Čerenkov detector specifically designed to measure luminosity. It consists of two arms located at $z = \pm 17$ m, each of which is an array of sixteen aluminium tubes, filled with C_4F_{10} gas, that surrounds the beam-pipe. It covers $0.26^\circ < \theta < 0.53^\circ$ ($5.4 < \eta < 6.0$). Figure 4.2 shows an illustration of the LUCID detector.

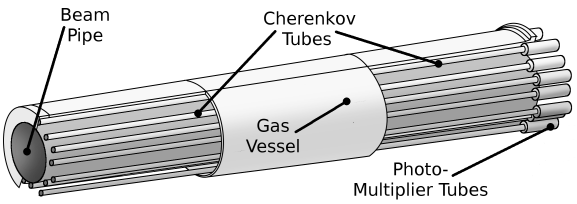

Figure 4.2: Illustration of the LUCID detector during Run 1. Source: Ref. [9].

The principle of operation is to register the light produced by charged particles, with momentum above the Čerenkov threshold, with photomultipliers located at the back of the detector. The main idea behind its operation is to suppress the background of low-momentum particles from beam-gas interactions or interactions with detector material.

The original design of LUCID presented signs of saturation and ageing of its photomultipliers during 2012. The solution adopted was to reduce its efficiency by removing the C_4F_{10} gas; in that configuration the only effective Čerenkov medium was a small quartz window that coupled the aluminium tubes to the photomultipliers. Given that higher instantaneous luminosity was expected for Run 2, LUCID had to be upgraded.

The new LUCID [69] was built following a similar principle but with several modifications that aimed to reduce the acceptance of the detector to avoid saturation. The three main changes relate to the photomultipliers, calibration and readout electronics.

Firstly, all the photomultipliers were replaced with similar ones with smaller photocathode, about a 50% reduction compared with the original. Secondly, an in-situ calibration capability in some photomultipliers was deployed with bismuth-207 sources that provide a mono-energetic electron from internal conversion. Thirdly, the cable

lengths of the readout electronics were reduced from the original 100 m to enable early digitalization of the signal that is needed to cope with the 25 ns bunch spacing used in Run 2.

The upgraded LUCID detector was commissioned during 2015 and it has been actively used for luminosity monitoring during the 2015 runs. It is the main luminometer for the 2015 proton–proton and lead–lead luminosity calibrations presented in this work.

Other luminometers

Several subsystems of the ATLAS experiment are used for relative luminosity monitoring. These are distinguished from the main luminometers BCM and LUCID because they do not provide bunch-by-bunch measurements (for different reasons, as explained in the following sections). However, they provide the redundancy needed to control systematic uncertainties related to the time stability of the response of the main detectors and instrumental effects associated with rate or total dose.

This section describes the inner detector, calorimeter, and medipix systems. This list is not exhaustive as several other sub-detectors have been studied as potential luminometers and others will surely follow.

Inner detector

The principle of operation of the inner detector as a luminometer is based on measurements of vertices, tracks, clusters, or hits; all are proportional to the luminosity. These methods are known as vertex, track, or hit counting. Even though these observables are highly correlated, there are important differences between them.

Vertex counting was the first method developed in ATLAS. It provided excellent linearity at low-μ, but its performance was limited by fake vertices and vertex-merging at high-μ [70]. Track counting provides linearity over a wide dynamic range from the low-luminosity of the vdM runs to the high-luminosity of physics running. It has been used extensively for the luminosity monitoring of the 2012 proton–proton data [50]. Hit counting[7] is independent of the details and weaknesses of the track and vertex reconstruction, but it is more susceptible to noise and cross-talk.

[7] Two hit counting methods are being developed in ATLAS; one using exclusively hits from the first pixel layer and the other using the SCT.

The inner detector is read out at 40 MHz so in principle per-bunch luminosity measurements are possible. However, the luminosity for the runs relevant for this thesis is very low and limits the statistical precision of these kind of measurements.

Calorimeters

The calorimeter system works as a luminometer by measuring a signal proportional to the total number of particles produced, which is proportional to the luminosity. Only a subset of the most precise subsystems are used: the current drawn by the photomultipliers of the hadronic calorimeter, and the current drawn[8] by certain liquid-argon calorimeter cells in the forward region.

These methods provide luminosity measurements integrated over seconds, therefore bunch-by-bunch measurements are not possible. These methods cannot be calibrated with the vdM method because the low luminosity of these runs falls below their sensitivity. But, their time stability and linearity makes them ideal for run-to-run stability measurements. They were used in the 2011 and 2012 proton–proton analysis to correct for time-dependent effects in the response of LUCID and BCM [49, 50].

Medipix

The Medipix system consists of an array of sixteen silicon pixel detectors (≈ 2 cm^2 each) located at different positions around the ATLAS cavern to monitor radiation levels. Its readout electronics (known as Medipix-2) integrates the number of hits due to single particles for a period that is adjustable but typically 5–60 s.

Their single-particle detection capability and large number of channels (256×256 per sensor) makes them attractive for luminosity measurements because they can operate in low and high luminosity regimes. However, the readout integration time prevents bunch-by-bunch measurements. They have been used in the 2012 proton–proton analysis for run-to-run stability monitoring [50].

[8] Current injected to compensate for voltage drop induced by charged particles.

4.4 Event counting algorithms

A luminosity algorithm is a way to extract μ_{vis} from another measured quantity. The main algorithm used with data of the LUCID and BCM detectors is the inclusive event-counting, known as EventOR. In the EventOR algorithm each bunch crossing is counted as an "event" if at least one of the two arms of the detector contains at least a channel with a signal above threshold. The probability to observe at least one inelastic interaction, P, is related via Poisson statistics to μ_{vis} as:

$$P(\mu_{\text{vis}}) = \frac{N_{\text{OR}}}{N_{\text{BC}}} = 1 - e^{-\mu_{\text{vis}}}, \tag{4.12}$$

where N_{BC} is the number of bunch crossings in a given time interval, and N_{OR} is the subset of those that passed the EventOR selection. Equation 4.12 can be trivially inverted to get μ_{vis}:

$$\mu_{\text{vis}} = -\ln\left(1 - \frac{N_{\text{OR}}}{N_{\text{BC}}}\right). \tag{4.13}$$

Evidently, this method fails when $N_{\text{OR}}/N_{\text{BC}}$ approaches unity. This is known as saturation or zero-starvation. The exact value of μ at which this algorithm saturates depends on the acceptance and efficiency of the detector ($\mu_{\text{vis}} = \mu\epsilon$). An ideal detector for the EventOR algorithm has small acceptance but provides statistically-precise measurements.

If the algorithms use the information of only one arm of the luminometers, instead of relying on the two arms, they are known as single-sided algorithms (EventA and EventC). The single-sided algorithms effectively reduce the acceptance of the detector by a factor of ≈ 2, thus delaying the saturation to higher μ values. For example, in 2012 the LUCID-based EventOR algorithm (LucidOR) saturated at $\mu \approx 20$, so the single-sided algorithms (LucidA and LucidC) were used and performed without saturation for pile-up up to $\mu \approx 30$.

Another type of event counting algorithm, known as EventAND, requires at least a hit in each arm. This coincidence requirement reduces certain sources of background and saturates at higher values of μ. The equivalent of equation 4.12 is more complicated and only under a series of approximations, μ_{vis} can be calculated numerically. The EventAND algorithms are not used in the work presented in this thesis because the luminosity levels of the analysed datasets is small enough that the EventOR algorithms can be used without saturation problems.

4.5 Datasets and corresponding accuracy aims

The 2013 proton–lead run

In January and February 2013, the LHC delivered collisions of protons and lead ions for the first time. The center-of-mass energy per nucleon–nucleon collision was 5 TeV. The peak luminosity reached at each run and the cumulative integrated luminosity are shown in figure 4.3. The peak luminosity reached a value of 1.1×10^{29} cm^{-2}s^{-1}, the number of colliding bunches was 254, and the bunch-averaged μ was about 0.1.

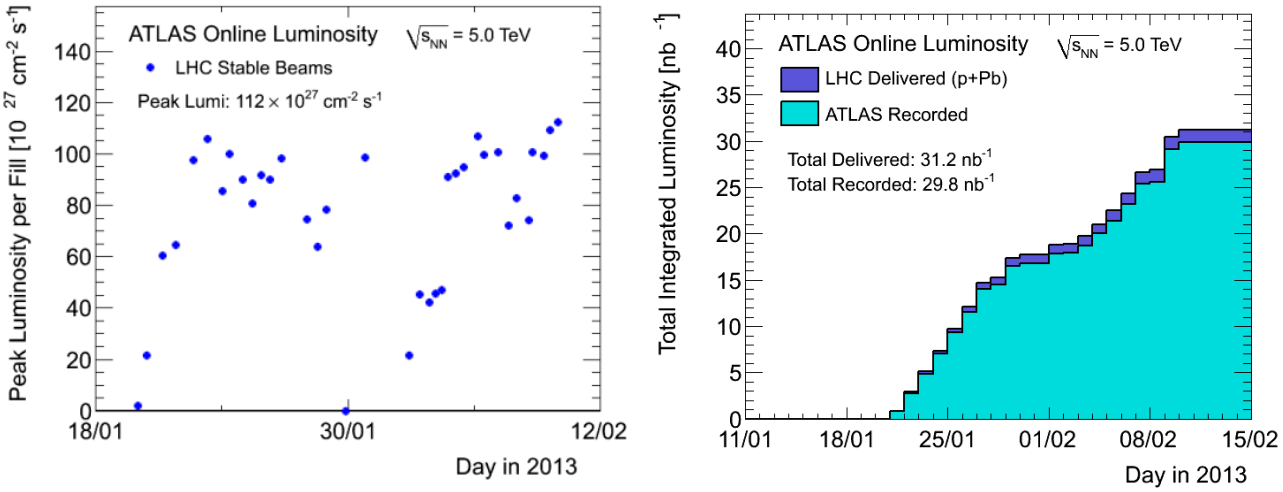

Figure 4.3: Peak luminosity per run (left) and integrated luminosity (right) as a function of time during the 2013 proton–lead runs. Source: Ref. [71].

The ATLAS proton–lead physics program includes measurements of cross-sections of hard-scattering processes that can be used to benchmark phenomena observed in heavy ion collisions (attributed to the quark-gluon plasma), and to study nuclear effects in the proton–nucleus system for the first time at the TeV scale. One of the primary targets is the measurement of electroweak bosons.

The integrated luminosity of the dataset[9] is about 29 nb^{-1}. The corresponding statistical uncertainty of the W and Z boson measurements is about 4% [72]. The uncertainty on the lepton trigger and reconstruction efficiencies are a few percent. Therefore, a luminosity uncertainty smaller than 4% is needed to make it sub-dominant.

[9] The rate of hard-scatterings in proton–nucleus collisions is approximately the rate in proton–proton collisions times the mass number of the nucleus; therefore, the 29 nb^{-1} of proton–lead data are roughly equivalent to 208×29 nb^{-1} = 6.0 pb^{-1} of proton–proton data.

The 2015 proton–proton run

The 2015 proton–proton run at 13 TeV marked the beginning of the LHC Run 2. The center-of-mass energy increase with respect to previous runs provided a stepwise increase in discovery potential. The peak luminosity reached in each run and the cumulative integrated luminosity are shown in figure 4.4. The peak luminosity reached a value of 5.0×10^{33} cm^{-2}s^{-1}, the number of colliding bunches[10] was 2253. The pile-up ($\mu \approx 15$) was significantly smaller than the average of the 2012 proton–proton run.

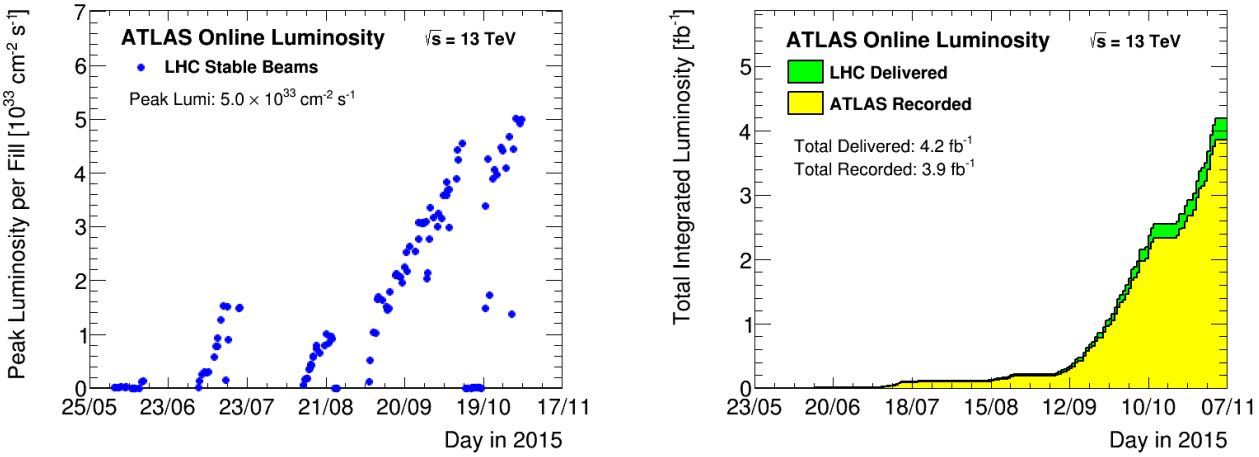

Figure 4.4: Peak luminosity per run (left) and integrated luminosity (right) as a function of time during the 2015 proton–proton runs. Source: Ref. [71].

Given that the integrated luminosity delivered to the experiments was 4 fb^{-1}, several key measurements at the new center-of-mass energy, such as Higgs boson or $t\bar{t}$ production, are limited by statistical uncertainties. One exception is the measurement of the inelastic proton–proton cross section, described in chapter 5. This measurement is the main driver for an accuracy goal of $< 3\%$ to make the luminosity uncertainty sub-dominant.

The 2015 lead–lead run

The lead–lead run in 2015 took place for a period of 3 weeks in November and December 2015. The LHC delivered collisions with an energy of 5 TeV (2.5 TeV per nucleon per beam). The peak luminosity reached at each run and the cumulative integrated luminosity are shown in figure 4.5. The peak luminosity reached a value

[10] This represents a significant increase with respect to the 2012 runs and was enabled by the change of bunch spacing from 50 ns to 25 ns.

of 3.3×10^{27} cm^{-2}s^{-1}, the maximum number of colliding bunches was 492, and the bunch-averaged μ was about 1×10^{-3}.

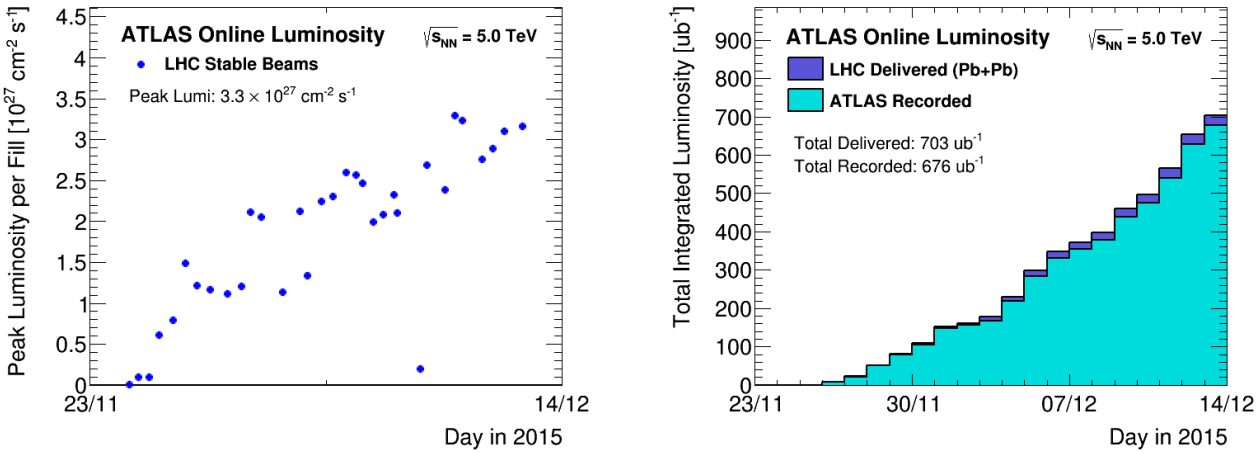

Figure 4.5: Peak luminosity per run (left) and integrated luminosity (right) as a function of time during the 2015 lead–lead runs. Source: Ref. [71].

The main goal of the ATLAS heavy ion physics program is to study global-event properties, and cross-sections do not play a major role. The notable exceptions are measurements of so called ultra-peripheral collisions that provide a proxy to gamma–gamma reactions such as $\gamma\gamma \to \mu\mu$. This process is experimentally clean because muons are reconstructed with percent-level accuracy. Therefore, the luminosity uncertainty is likely to be the dominant systematic uncertainty of the measurement [73]. However, given the novelty of this analysis a relatively large uncertainty in the range 5–10% would be acceptable for preliminary results.

Table 4.1 summarizes the luminosity ranges, physics measurements and corresponding accuracy aims of each of the datasets studied in this work.

Table 4.1: Run information, physics goals and accuracy requirements for the datasets calibrated in this work.

Run	$\sqrt{s}$	Peak luminosity	$\langle\mu\rangle$	Measurements	Accuracy aimed
p-Pb	5 TeV	1.1×10^{29} cm^{-2}s^{-1}	≈ 0.1	W and Z bosons	$< 4\%$
pp	13 TeV	5.0×10^{33} cm^{-2}s^{-1}	≈ 15	σ_{inel}	$< 3\%$.
Pb-Pb	5 TeV	3.3×10^{27} cm^{-2}s^{-1}	≈ 0.001	$\gamma\gamma \to \mu\mu$	$< 10\%$.

vdM datasets

For each data taking period, a series of vdM scans are performed. For all cases, the beams are separated stepwise in either the horizontal or vertical directions by up to $\pm 6\sigma_{\mathrm{b}}$ in one direction while keeping the beams centred in the other. The data are selected with random triggers and recorded at about 3 kHz for 30 s per scan point.

Proton–lead vdM scans

As a collision of a proton with a lead ion is an asymmetric system, there are two different ways the LHC can deliver these collisions, i.e. with protons circulating clockwise or anticlockwise. These are called p+Pb and Pb+p configurations. Both were used during the 2013 proton–lead run to constrain systematic uncertainties. Three sets of scans were performed in January with the p+Pb beam configuration, and two sets of scans in February with the Pb+p beam configuration.

Table 4.2 summarises the main characteristics of the 2013 proton–lead vdM scans. The experimental conditions during the vdM scans were similar to those in the normal running, i.e. same β^* (see table 1.2) and similar bunch pattern (total number of bunches and bunch spacing). This eliminates the need for dedicated "calibration transfer" studies to test the response of the luminometers at both low and high-luminosity[11].

Table 4.2: Summary of the main characteristics of the 2013 proton–lead vdM scans. The values of per-bunch luminosity and μ are given for zero beam separation.

Scan number	1–3	4–5
LHC fill number	3503	3540
Date	28–29 January 2013	8 February 2013
Number of scans	3 pairs of x/y scans	2 pairs of x/y scans
Bunches colliding in ATLAS	254	248
Total number of bunches per beam	272 (p), 338 (Pb)	272 (p), 314 (Pb)
Typical charge per bunch [$10^{11}e$]	0.1 (p), 0.1 (Pb)	0.1 (p), 0.1 (Pb)
$\beta^\star$ [m]	0.8	0.8
Typical luminosity/bunch [$\mathrm{b}^{-1}/\mathrm{s}$]	$1.8 \cdot 10^{-2}$	$1.8 \cdot 10^{-2}$
μ	0.03	0.03

[11] This is a central part of the luminosity calibration of high-luminosity proton–proton data [49].

Proton–proton vdM scans

A total of three vdM scans were performed in two consecutive LHC fills in August 2015. The first two were performed in quick succession; a third was planned in the same fill but a LHC beam dump prevented it. Therefore, the third scan had to be accommodated in another fill shared with the ALICE and LHCb experiments.

Table 4.3 summarises the main characteristics of the 2015 proton–proton vdM scans. The luminosity during the vdM scans was much lower than in physics running. Therefore, dedicated calibration studies are needed. However, this is not true for the calibration of specific runs at low-luminosity such as those used for the measurement of the inelastic proton–proton cross-section.

Table 4.3: Summary of the main characteristics of the 2015 proton–proton vdM scans. The values of per-bunch luminosity and μ are given for zero beam separation.

Scan number	1–2	3
LHC fill number	4266	4269
Date	24–25 August 2015	26 August 2015
Number of scans	2 pairs of x/y scans	1 pair of x/y scans
Bunches colliding in ATLAS	30	8
Total number of bunches per beam	44	51
Typical charge per bunch [$10^{11}e$]	0.9	0.9
$\beta^\star$ [m]	19	19
Typical per-bunch luminosity [$\mathrm{b}^{-1}/\mathrm{s}$]	8.5	8.5
μ	0.7	0.7

Lead–lead vdM scans

The lead–lead vdM scans were performed in December 2015. A total of three vdM scans were performed in the same LHC fill. Table 4.4 summarises the main characteristics of all the 2015 lead–lead vdM scans. As in the proton–lead case, the luminosity during the scans is similar to that in physics running, so no calibration transfer studies are needed.

Table 4.4: Summary of the main characteristics of the 2015 lead–lead vdM scans. The values of per-bunch luminosity and μ are given for zero beam separation.

Scan number	1–3
LHC fill number	4689
Date	3 December 2015
Number of scans	3 pairs of x/y scans
Bunches colliding in ATLAS	432
Total number of bunches per beam	424
Typical charge per bunch [$10^{11}e$]	0.09
$\beta^\star$ [m]	0.8
Typical per-bunch luminosity [$\mathrm{b}^{-1}/\mathrm{s}$]	$2 \cdot 10^{-4}$
μ	$4 \cdot 10^{-3}$

4.6 vdM analysis

For each luminometer and algorithm, the μ_{vis} is extracted bunch-by-bunch at each scan step by exploiting Poisson statistics as described in section 4.4. The measured μ_{vis} as a function of beam separation (a Gaussian-like curve) is fitted to extract the peak value, $\mu_{\mathrm{vis}}^{\mathrm{peak}}$, and the integral of the curve; these are used to calculate Σ according to equation 4.7. Then, the σ_{vis} is calculated as:

$$\sigma_{\mathrm{vis}} = \frac{\mu_{\mathrm{vis}}^{\mathrm{peak}}}{\mathcal{L}} = \mu_{\mathrm{vis}}^{\mathrm{peak}} \frac{4\pi \Sigma_{\mathrm{X}} \Sigma_{\mathrm{Y}}}{f_{\mathrm{r}} n_1 n_2}. \tag{4.14}$$

The main ingredients of the analysis are the measurement of the bunch charge and the beam-overlap integral; these are described in sections 4.6.1 and 4.6.2. The results obtained with no corrections are presented in 4.6.3. A series of corrections are applied to the data to ensure an unbiased measurement; these are described in section 4.6.4. The corrected results are presented in section 4.7. The systematic uncertainties are presented in section 4.8.

4.6.1 Bunch current measurement

The measurement of the per-bunch current is divided into separate measurements of the total beam charge and the relative per-bunch fractions. This is necessary because high accuracy and speed are competing requirements for the detectors and associated readout electronics used for these measurements.

The total beam charge is measured with the LHC DC current transformers (DCCT) [74]. The absolute scale of this device is known at the per-mil level. The relative per-bunch charge is measured with two fast beam-current transformers (FBCT) [75], and with a beam pick-up system (BPTX) [76, 77]. These devices have have the capability to measure each of the 25 ns slots in each beam

The redundancy provided by the FBCT and BPTX measurements is exploited to estimate systematic uncertainties, as described in section 4.8. This redundancy was essential for the analysis presented in this work because the FBCT had operational problems during the January 2013 proton–lead scans. For consistency, the BPTX is used as default for all datasets.

The total charge per beam is shown in figure 4.6 as a function of time during the first two scans in each dataset. The proton–lead scan data show a stark difference between the proton and lead beams. The initial lead-beam charge is about 37×10^{11} protons but decreased by 4% over the duration of a scan (about 25 minutes), whereas the proton-beam charge remained stable at about 38.5×10^{11} protons throughout. This is due to the higher rate at which lead ions are lost from the beam due to intra-beam electromagnetic scattering.

The proton–proton data show no significant beam loss over the duration of the scans. The total beam charge in each beam differs by only a few percent and is about 39×10^{11} protons. The lead–lead data show a significant drop in both beams from the initial total charge of 39×10^{11} protons, but the decay rate (about 2% per scan) is lower than in the proton–lead scans. This may be explained by the difference in emittance of the beams (that influences intra-beam scattering) in the two datasets.

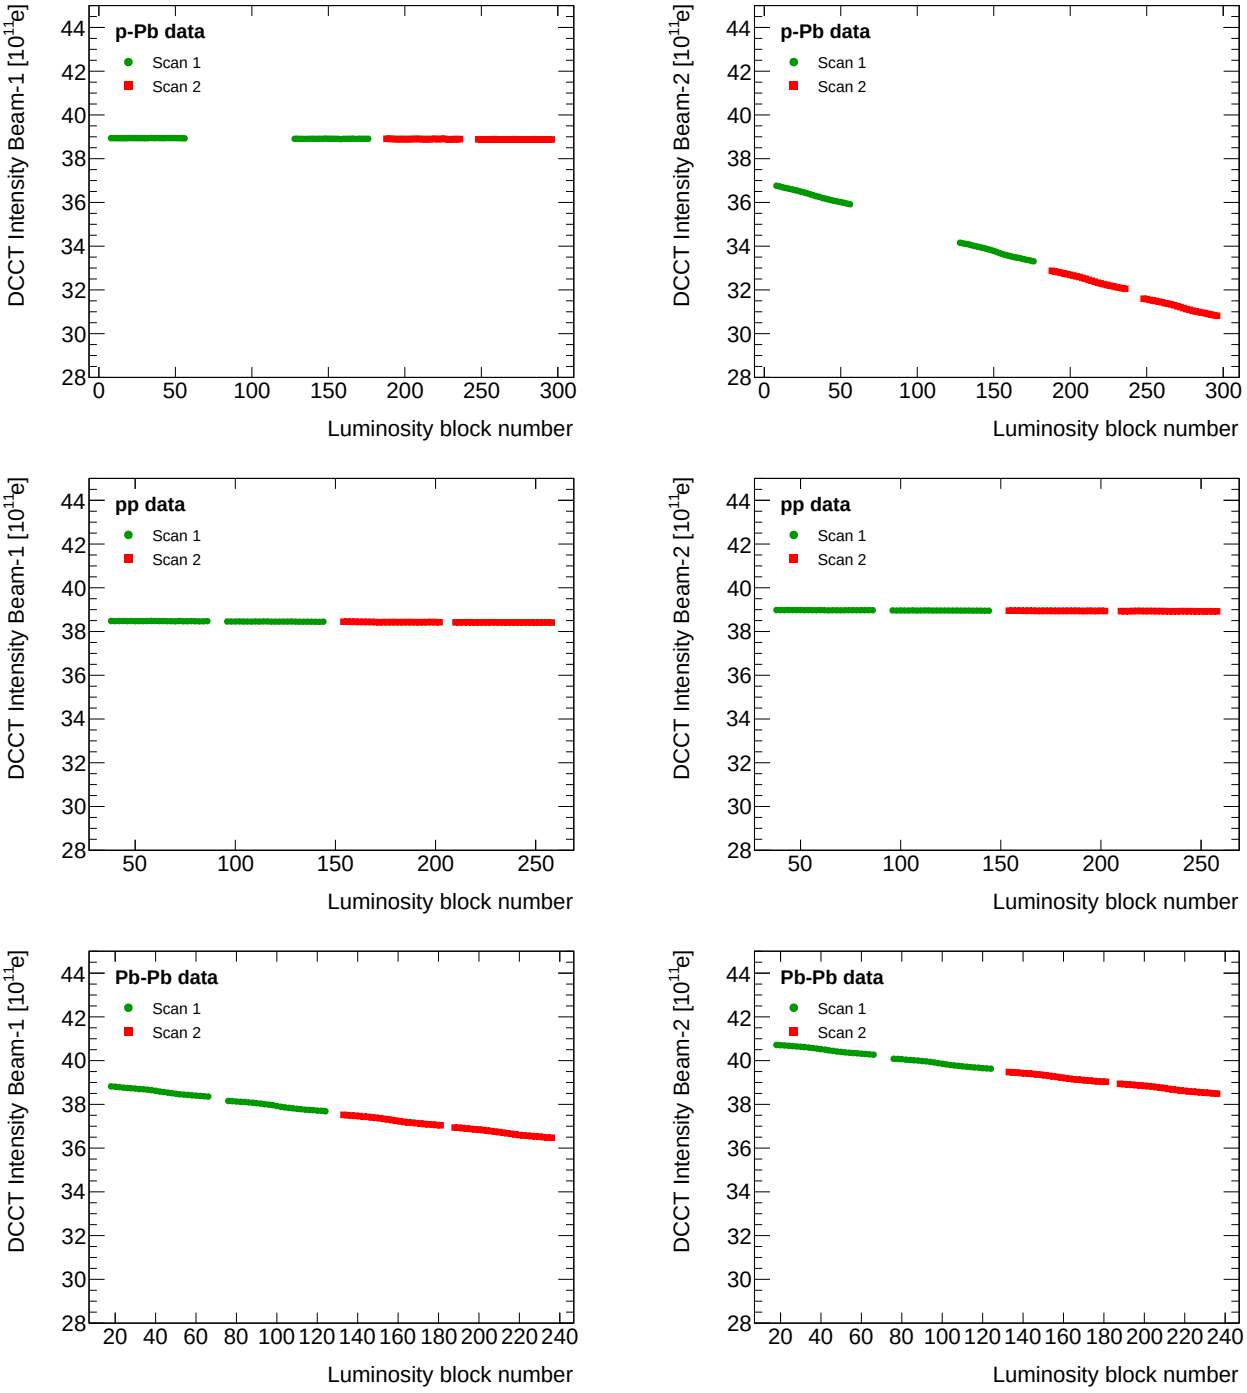

Figure 4.6: Total beam charge as a function of luminosity block number (i.e. time) during the first two sets of proton–lead (top panels), proton–proton (middle panels), and lead–lead scans (bottom panels). The left panels show the beam 1 data, whereas the right panels show the beam 2 data (that corresponds to the lead beam in the proton–lead scans). Each luminosity block lasts about 30 seconds. The data shown correspond to the luminosity blocks during vdM scans only, so the gaps are periods in between scans. In the proton–lead case there is a larger gap due to an aborted scan.

The relative per-bunch measurements provided by the BPTX are normalized to the DCCT scale at each luminosity block. Figure 4.7 presents the nominal per-bunch charge (i.e. at the beginning of the first horizontal scan) for both beams. The proton–proton data show an average per-bunch charge of about 0.9×10^{11} protons in both beams. This value is close to the LHC design value. The variations over the bunches is at the level of 10% and shows no systematic pattern common to all trains.

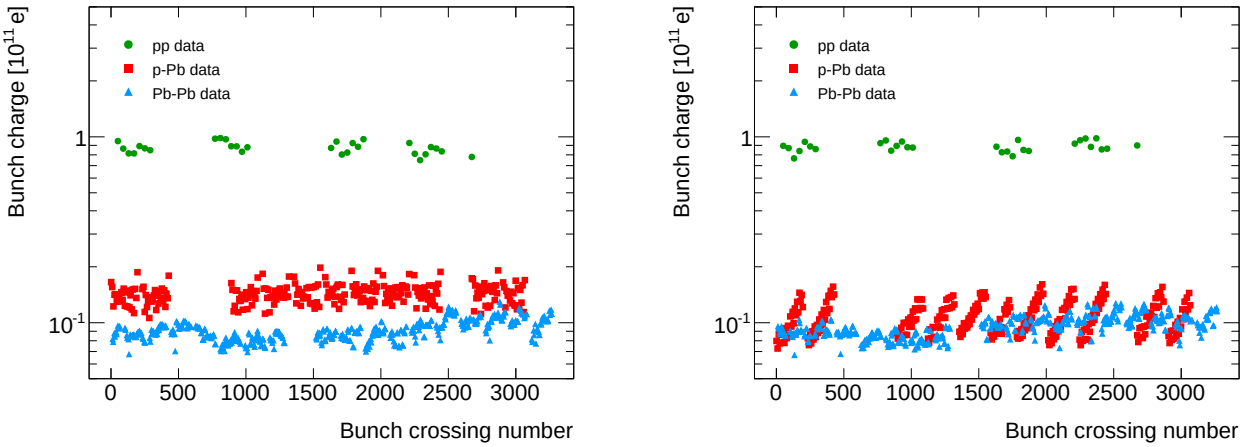

Figure 4.7: Bunch charge in beam 1 (left panel) and beam 2 (right panel) at the beginning of the first horizontal scan in the proton–proton, proton–lead and lead–lead scans. Beam 2 corresponds to the lead beam in the proton–lead scans.

The proton–lead data show a difference between the proton and lead beams. The proton bunches have an average charge of about 0.15×10^{11} protons with a variation of about $\pm 20\%$. The lead bunches have an average charge of about 0.1×10^{11} protons with a systematic variation of about $\pm 30\%$ from the first to the last bunch in the bunch train. This bunch dependence is explained by the LHC fill pattern: the earlier the bunch is injected, the smaller the current due to losses.

The lead–lead data show a per-bunch charge of about 0.09×10^{11} protons in both beams. The bunch spread is at the level of 30% with no systematic variations. The observed lack of systematic pattern, as opposed to the proton–lead case, suggests that the latter was caused by specific details of the fill pattern.

4.6.2 Fits to scan curves

The extracted μ_{vis} as a function of "nominal" beam separation, i.e. the separation specified by the LHC control system, is a Gaussian-like curve that is referred to as a scan curve.

The μ_{vis} at each scan point depends on the corresponding bunch charge product, so the decay of the beam intensities during the duration of the scan (shown in figure 4.6) would distort the shape of the scan curve. To properly take into account the beam intensity decay, the μ_{vis} is normalized at each scan step by the corresponding bunch charge product to define the "specific interaction rate" $\mu_{\text{sp}} = \mu_{\text{vis}}/(n_1 n_2)$.

The data for each bunch crossing of each scan are fit independently to a characteristic function to provide a measurement of $\mu_{\text{sp}}^{\text{peak}}$ and Σ. The horizontal and vertical scans provide in principle an independent measurement of $\mu_{\text{sp}}^{\text{peak}}$, however slightly off-centring of the beams along with other subtle effects (discussed later) lead to differences between the results; thus, the $\mu_{\text{sp}}^{\text{peak}}$ value (to be used in equation 4.14) is taken as the average of the two measurements.

Proton–lead scan curves

Figure 4.8 shows a typical scan curve and fit result of the proton–lead dataset. Neither a Gaussian nor a double-Gaussian function describe the data. To address the non-Gaussian nature of the scan curves, a Gaussian times a polynomial of n-th order (GPn) is used. The lowest order polynomial that is able to describe the data is sixth. The explicit functional form is:

$$f(\Delta) = e^{-\Delta^2/2\sigma^2} \times \left(1 + \sum_{n=2}^{6} \left[C_n \left(\frac{\Delta}{\sigma}\right)^n\right]\right), \tag{4.15}$$

where $\Delta = \delta$ - δ_0 is the difference between the beam separation, δ, and a free parameter, δ_0, that describes the off-centring of the colliding beams. The linear term in the polynomial is set to zero to increase the fit stability.

The GP6 function provides a good description of the scan curves for all bunch crossings and algorithms. The width of the scan curve is about 40 μm. The pulls of the fit, shown in the bottom panel of the figure, show no evidence of a systematic pattern. Figure 4.9 shows the reduced χ^2 of the fits as a function of bunch crossing number. In most cases it lies in the range 1.0 ± 0.5. These results correspond to LucidOR which is the algorithm that provides the smallest statistical uncertainty.

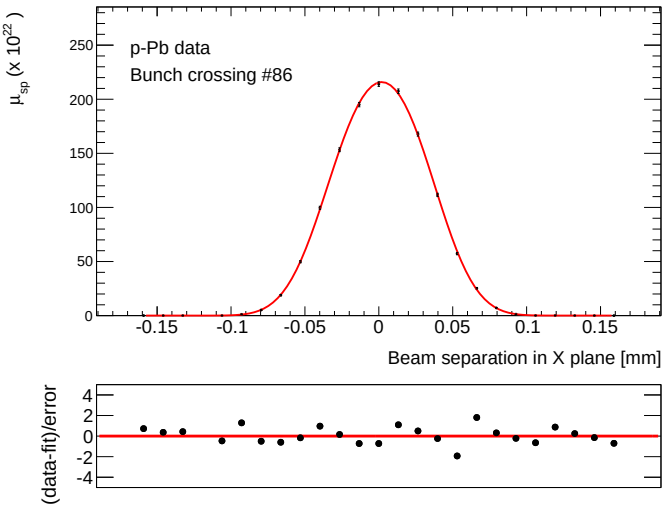

Figure 4.8: Specific interaction rate measured by the LucidOR algorithm as a function of beam separation. The data correspond to bunch crossing number 86 and the first proton–lead scan. The error bars represent statistical uncertainty only. The bottom panel shows the fit pulls.

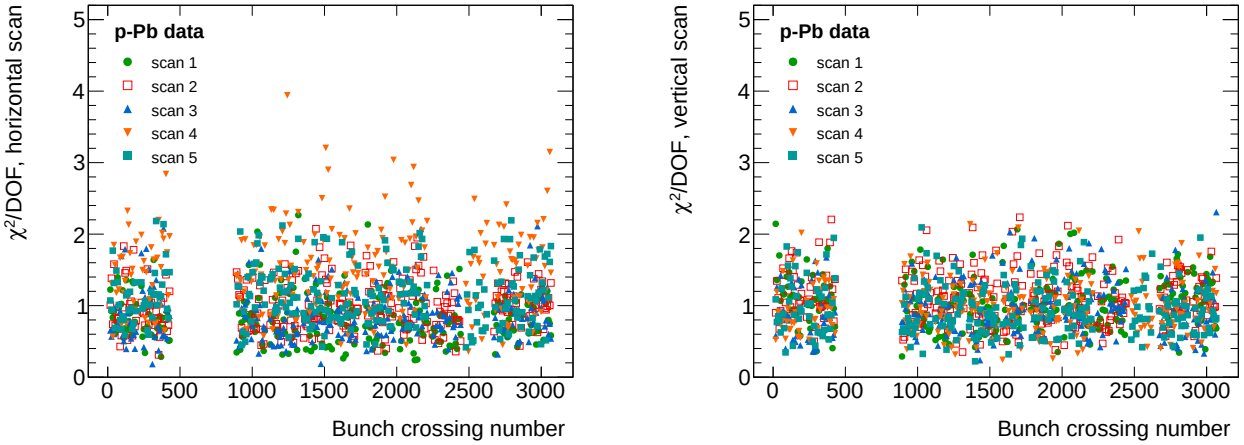

Figure 4.9: Reduced χ^2 of the fit, described in the text, to the scan curves in horizontal (left panel) and vertical (right panel) proton–lead scans. The results correspond to LucidOR.

Proton–proton scan curves

Figure 4.10 shows a typical scan curve and fit result of the proton–proton dataset. The statistical uncertainty of the data near the peak is about 0.4%. The data are fitted successfully with a GP6 function. The pulls of the fit show no evidence of a systematic pattern. The width of the scan curve is about 150 μm which is much larger than for the proton–lead case. This illustrates the corresponding difference in bunch emittance.

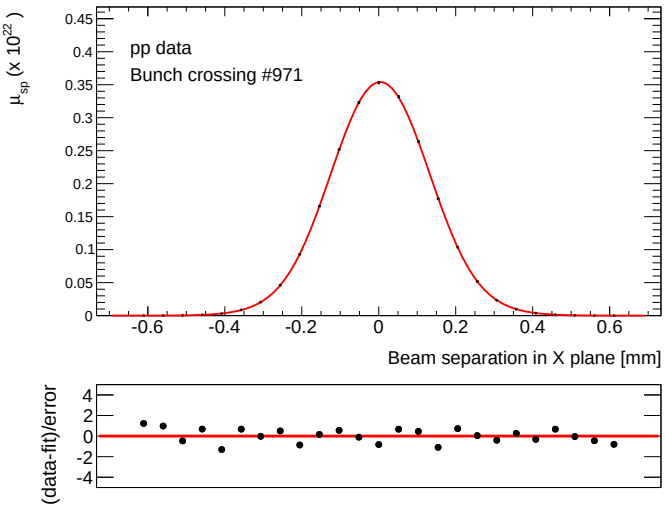

Figure 4.10: Specific interaction rate measured by the LucidOR algorithm as a function of beam separation. The data correspond to bunch crossing number 971 and the first proton–proton scan. The error bars represent statistical uncertainty only. The bottom panel shows the fit pulls.

The reduced χ^2 values for each bunch crossing during the first scan are shown in figure 4.11. The average reduced χ^2 is 1.1 and 1.0 in the horizontal and vertical planes.

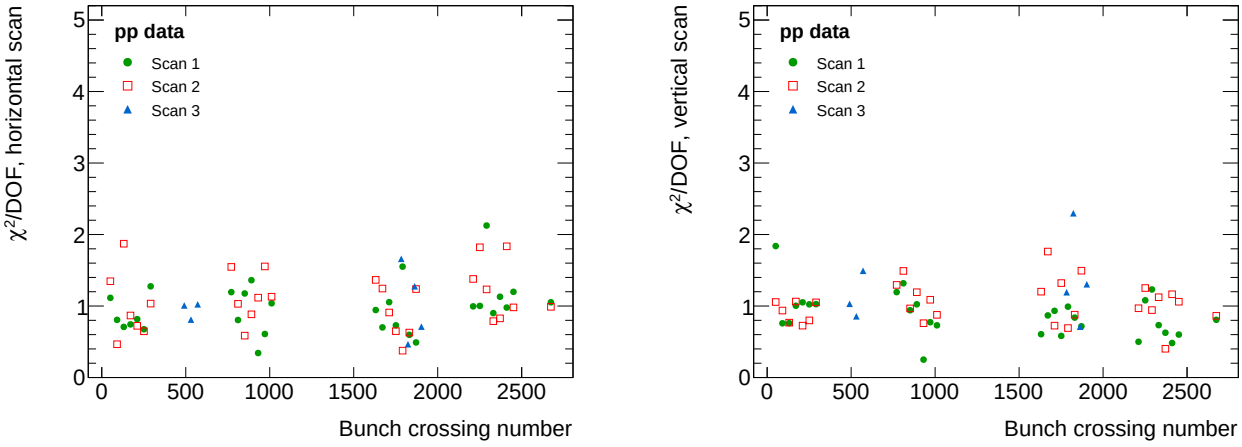

Figure 4.11: Reduced χ^2 of the fit, described in the text, to the scan curves in horizontal (left panel) and vertical (right panel) proton–proton scans. The results correspond to LucidOR.

Figure 4.12 shows the average fit pull as a function of beam separation in scan 1. The results obtained with the GP4 function are also shown for comparison. The GP6 results do not show a significant systematic pattern and all pulls are within $\approx \pm 0.5\ \sigma$. In contrast, the GP4 results show a pronounced systematic pattern and deviations of about $\pm 1\ \sigma$, which corresponds to biases of about $\pm 0.5\%$ near the peak.

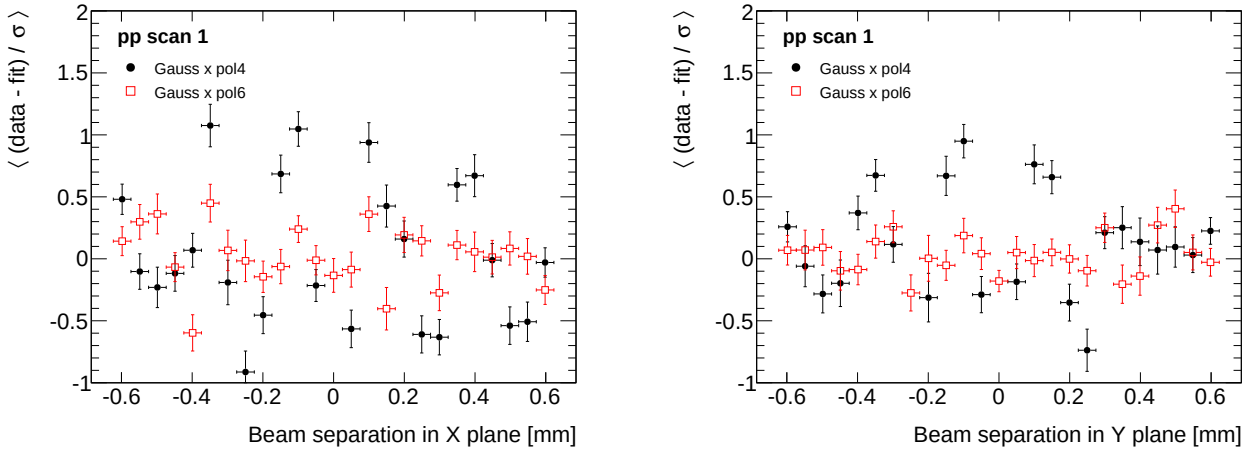

Figure 4.12: Fit pull as a function of beam separation in the horizontal (left panel) and vertical (right panel) planes in the first proton–proton scan. The error bars represent statistical uncertainty only. The results correspond to LucidOR.

Lead–lead scan curves

Figure 4.13 shows a typical scan curve and fit result of the lead–lead dataset. The statistical precision of the data near the peak is about 3%. The default fit function used for this data is a GP4 function. The fit pulls show no evidence of a systematic pattern. The width of the scan curve is about 40 μm, which is similar to the proton–lead data.

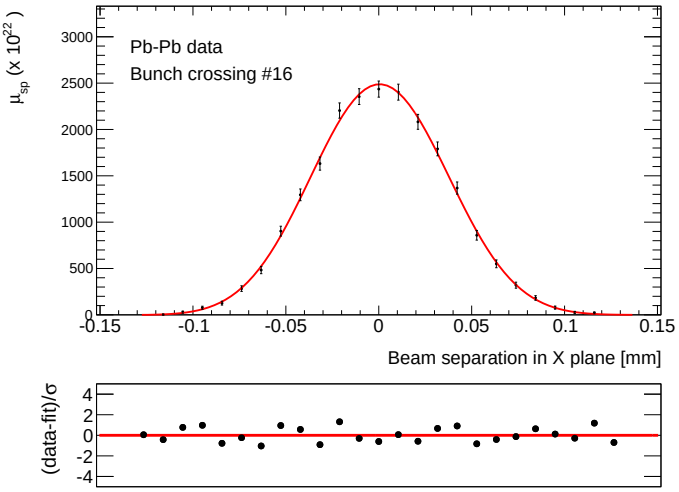

Figure 4.13: Specific interaction rate measured by the LucidOR algorithm as a function of beam separation. The data correspond to bunch crossing number 16 and the first lead–lead scan. The error bars represent statistical uncertainty only. The bottom panel shows the fit pulls.

The reduced χ^2 distributions obtained with different fit functions are shown in figure 4.14. The fit functions shown are barely distinguishable with a study of goodness-of-fit alone as even Gaussian fits yield distributions centred at $\chi^2 \approx 1$.

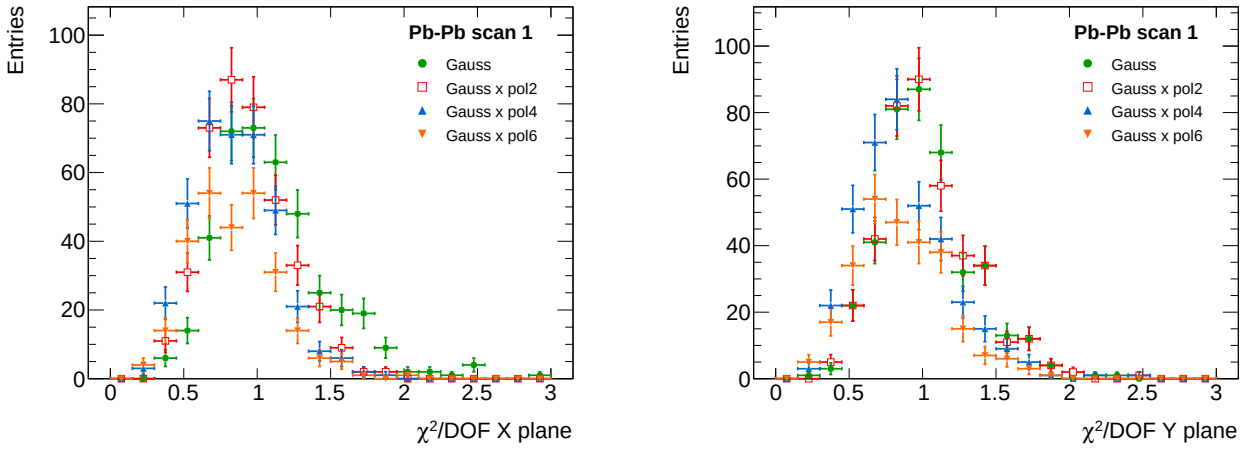

Figure 4.14: Frequency distribution of the reduced χ^2 of fits to scan curves for all bunch crossings in the first lead–lead scan. The error bars represent statistical uncertainty only. The results correspond to LucidOR.

However, a systematic pattern in the pulls of the fit becomes evident when averaging the results of all bunch crossings, as shown in figure 4.15. This average should be consistent with unity because statistical fluctuations should average out.

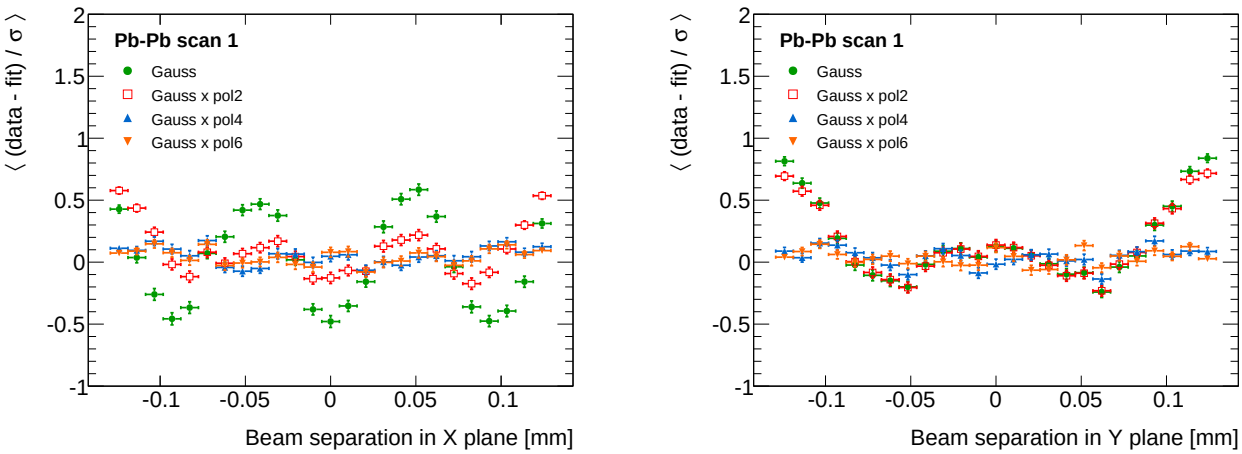

Figure 4.15: Fit pull as a function of beam separation in the horizontal (left panel) and vertical (right panel) planes in the first lead–lead scan. The error bars represent statistical uncertainty only. The results correspond to LucidOR.

The Gaussian fit results show the most evident systematic patterns; the clear modulation observed in the horizontal scan data clearly reflects that the underlying distribution cannot be described by a Gaussian function. The GP2 results are better than the single-Gaussian results in the horizontal plane because the amplitude of the observed modulation near the peak region is reduced by about a factor of 2. However, the data still show a systematic pattern, specially in the tails of the distribution. The results for the vertical plane are almost identical to the Gaussian function results.

The GP4 and GP6 describe well the data in both planes. There is little difference between the GP4 and GP6 results. This indicates that the statistical power of the data is not enough to constrain all parameters of the GP6 function; indeed, a non-negligible fraction of the fits of the GP6 function do not converge. For this reason, the GP4 function is used as the default for the lead–lead data.

4.6.3 Uncorrected results

This section presents the results of the fits to the scan curves described in the previous section, and the corresponding σ_{vis} values. The results shown correspond to LucidOR, which is the algorithm with the highest statistical precision in all datasets.

Proton–lead results

Given the different beam configuration of the January and February scans, these are presented separately. Figure 4.16 shows the Σ and $\mu_{\text{sp}}^{\text{peak}}$ results for the horizontal and vertical scans in the January session. A double band structure is observed in the Σ and $\mu_{\text{sp}}^{\text{peak}}$ measurements in all scans. This may be explained by the difference in bunch emittance of two sets of bunches. The observed Σ_{Y} ($\approx$ 36–40 μm) is significantly larger than Σ_{X} ($\approx$ 30–34 μm) in all bunch crossings and scans. This indicates a larger bunch emittance in the vertical plane, i.e. the beams are not spherical. Most of the $\mu_{\text{sp}}^{\text{peak}}$ measurements are within $(200\text{–}230) \times 10^{-22}$, which is equivalent to a μ_{vis} range 0.02–0.07.

A clear increase of Σ and decrease of $\mu_{\text{sp}}^{\text{peak}}$ is observed with scan number, i.e. with time. This cannot be attributed to the decay of bunch charge because that effect is already taken into account in the definition of $\mu_{\text{sp}}^{\text{peak}}$. Rather, this feature is attributed to emittance growth over time. The impact of emittance growth on the σ_{vis} determination is discussed in section 4.8.

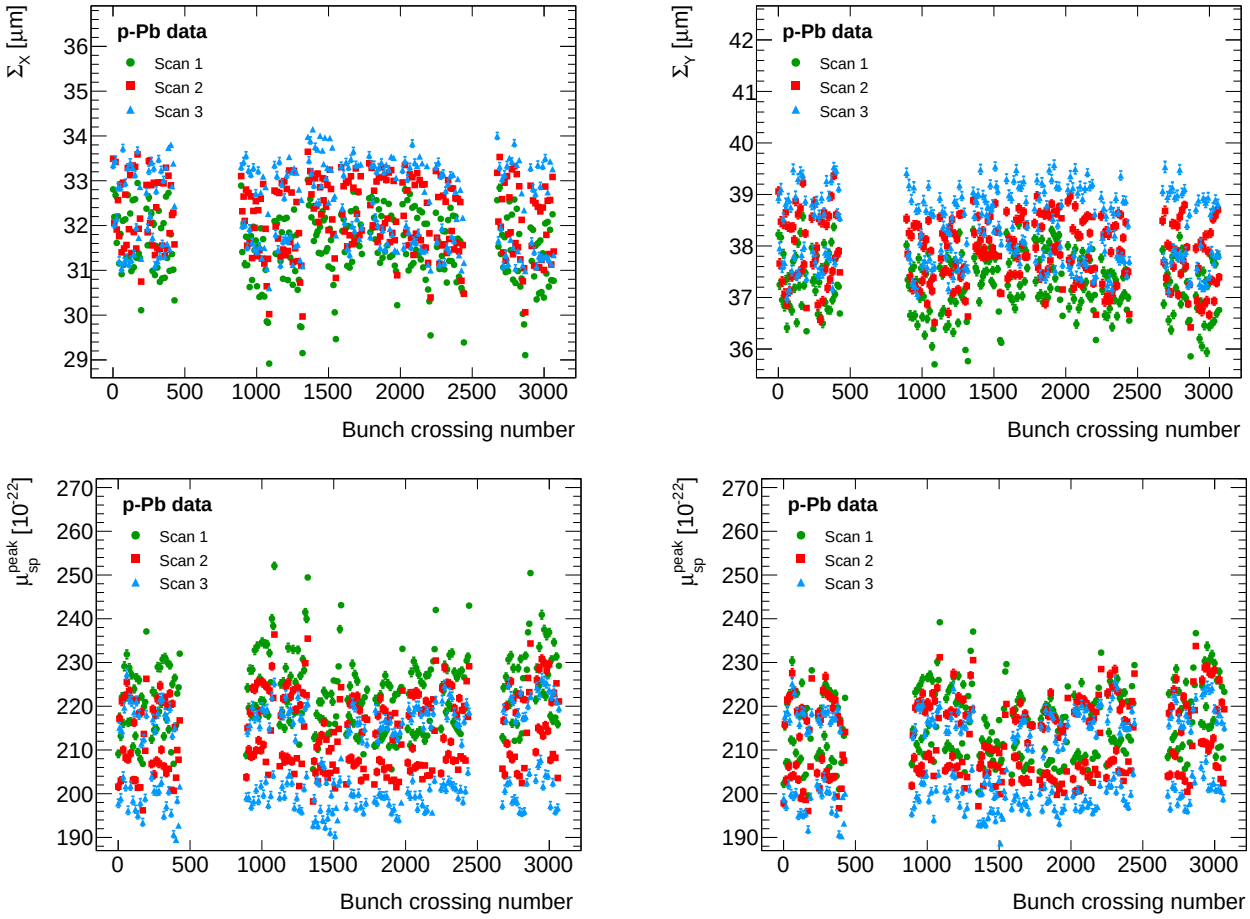

Figure 4.16: Uncorrected Σ and $\mu_{\mathrm{sp}}^{\mathrm{peak}}$ for each bunch crossing in the horizontal (left panels) and vertical (right panels) proton–lead scans of the January session. The error bars, which are typically smaller than the marker size, represent statistical uncertainty only.

Figure 4.17 shows the Σ and $\mu_{\mathrm{sp}}^{\mathrm{peak}}$ results for the horizontal and vertical February scans. The data do not show the double band structure observed in the January scans. However, the data still show a clear systematic pattern: the last bunch crossings in each train have lower Σ and higher $\mu_{\mathrm{sp}}^{\mathrm{peak}}$. This can be explained by emittance growth during the injection phase because the latter bunches were injected earlier to the LHC. This is similar to the pattern observed in the bunch charge measurements (shown in figure 4.7).

The observed Σ_{Y} is in the range $\approx$ 36–38 μm and the Σ_{X} in $\approx$ 32–34 μm (excluding outliers). The $\mu_{\mathrm{sp}}^{\mathrm{peak}}$ values measured in the horizontal and vertical scans are similar and in the range $(200\text{–}220) \times 10^{-22}$. These values are similar to those observed in the January scans. The data show signs of emittance growth but it is less pronounced than

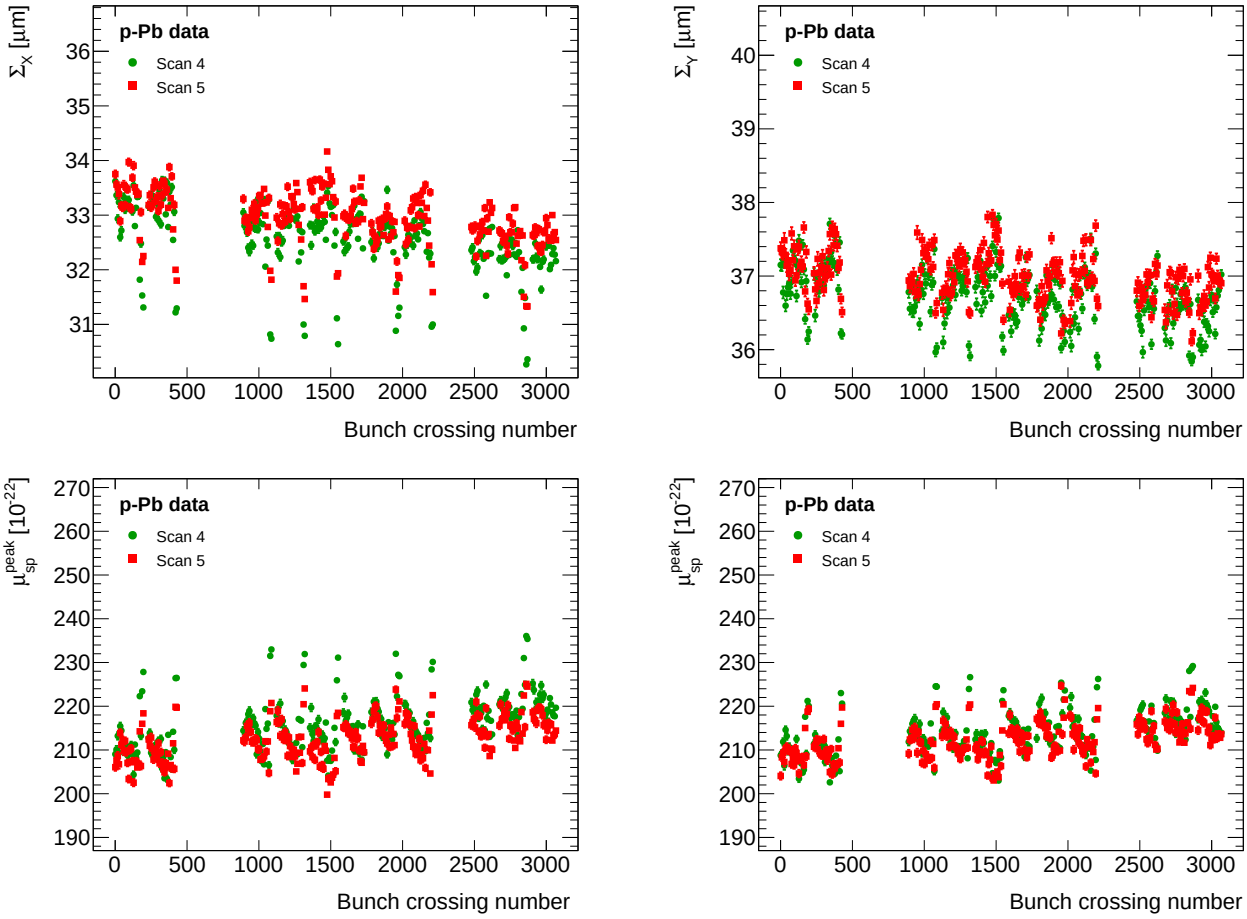

Figure 4.17: Uncorrected Σ and $\mu_{\mathrm{sp}}^{\mathrm{peak}}$ for each bunch crossing in the horizontal (left panels) and vertical (right panels) proton–lead scans of the February session. The error bars, which are typically smaller than the marker size, represent statistical uncertainty only.

for the January scans. This is because the February scans lasted a total of ≈ 90 minutes whereas the January scans lasted ≈ 200 minutes.

Figure 4.18 shows the LucidOR σ_{vis} per bunch crossing in the January and February scans. The different bunch crossings in all scans provide a large number of independent measurements that should be consistent within statistical fluctuations. But, the results of all scans show a systematic bunch dependence: the measured σ_{vis} decreases with position of the bunch in a given train. This is observed in all trains.

The relative difference between the first and the last bunch crossing in each train is about 4% in the January scans and about 3% in the February scans. As will be discussed in section 4.6.4, the origin of this feature is a bias in the bunch charge measurements.

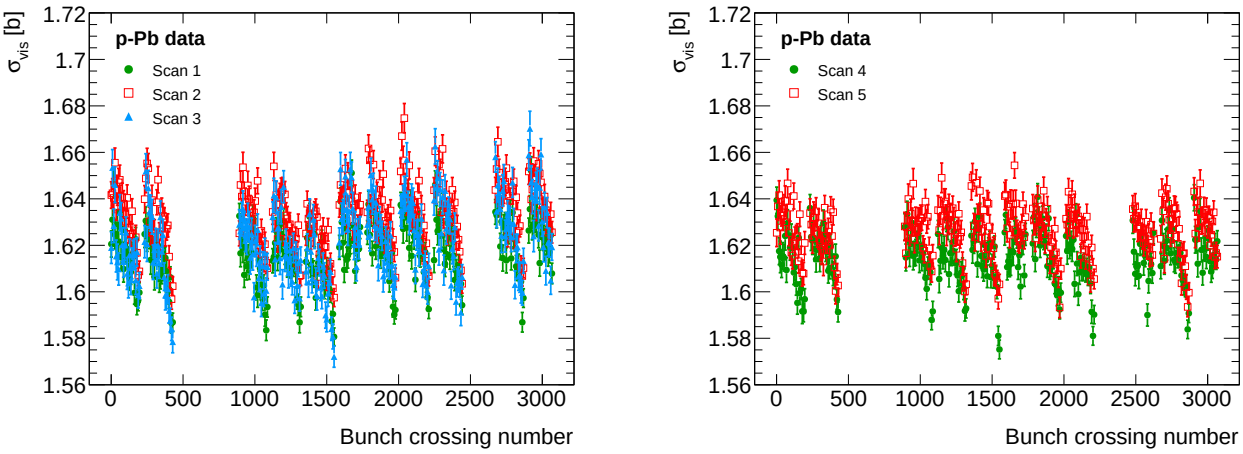

Figure 4.18: Uncorrected LucidOR σ_{vis} for each bunch crossing in the January (left panel) and February (right panel) proton–lead scans. The error bars represent statistical uncertainty only.

Table 4.5 shows the bunch-averaged σ_{vis} in each of the five proton–lead scans for all the algorithms calibrated. These are the inclusive algorithms LucidOR, BcmHOR, BcmVOR; and the the single-sided algorithms LucidA, and LucidC.

Table 4.5: Uncorrected bunch-averaged σ_{vis} (in b) for the calibrated algorithms in the proton–lead scans. The statistical uncertainties of these measurements are negligible.

Scan Number	1	2	3	4	5
LucidOR	1.617	1.633	1.621	1.615	1.626
LucidA	1.291	1.303	1.293	1.073	1.080
LucidC	0.938	0.947	0.941	1.187	1.196
BcmHOR	0.448	0.452	0.450	0.448	0.451
BcmVOR	0.450	0.455	0.452	0.450	0.454

Due to the asymmetric nature of the proton–lead collisions, the single-sided algorithms are not expected to have the same efficiency (and thus σ_{vis}) because the particle multiplicities are different in the proton and lead fragmentation sides. Moreover, the C side of LUCID had fewer number of tubes operating during this run.

The scan-averaged σ_{vis} for the January and the February scans are shown in table 4.6. The EventOR algorithms show little or no dependence on the beam configuration. Given that the proton–lead inelastic cross-section is about 2.1 b [78], the measured σ_{vis} translates to an efficiency of about 80% for LucidOR and about 20% for BcmHOR and BcmVOR.

Table 4.6: Uncorrected σ_{vis} (in b) averaged over the January and February proton–lead scans. The statistical uncertainties of these measurements are negligible.

Scan Number	1–3	4–5
LucidOR	1.624	1.621
LucidA	1.296	1.077
LucidC	0.942	1.192
BcmHOR	0.450	0.450
BcmVOR	0.452	0.452

Proton–proton results

Figure 4.19 shows the Σ and $\mu_{\text{sp}}^{\text{peak}}$ results for the horizontal and vertical proton–proton scans. The data of scan 1 and 2 (performed in the same LHC fill) show reasonable agreement, but also signs of emittance growth in most bunch crossings. The results of scan 3 are significantly different: the Σ is smaller and the $\mu_{\text{sp}}^{\text{peak}}$ is larger. This indicates smaller bunch emittance in the fill of scan 3.

The observed Σ_{X} ($\approx$ 120–132 μm) are significantly larger than the Σ_{Y} ($\approx$ 106–120 μm). This indicates a larger emittance in the horizontal plane. As this is the opposite to what is observed in the proton–lead scans, the degree of sphericity of the beams (similarity of Σ_{X} and Σ_{Y}) must depend on the run conditions.

The bunch-to-bunch variation is clearly non-statistical and is much larger in scan 3. There is no systematic pattern common to all trains. The data of scan 1 and 2 show a clear train dependence where the Σ increases and $\mu_{\text{sp}}^{\text{peak}}$ decreases with train number. This may be due to emittance growth if the injection pattern filled all bunches in the same train consecutively.

Figure 4.20 shows the LucidOR σ_{vis} for each bunch crossing in all scans. The results of the first two scans are in reasonable agreement but differ significantly from the third scan by about 2%. As will be discussed in section 4.6.4, this is mainly due to a bias caused by beam drift during scan 3.

The bunch-to-bunch variation in scans 1 and 2 is clearly not statistical (a constant fit yields a reduced $\chi^2 \approx 3$) but does not follow a systematic pattern common to all trains. In contrast, the results of scan 3 are consistent within uncertainties (reduced $\chi^2 \approx 1.2$).

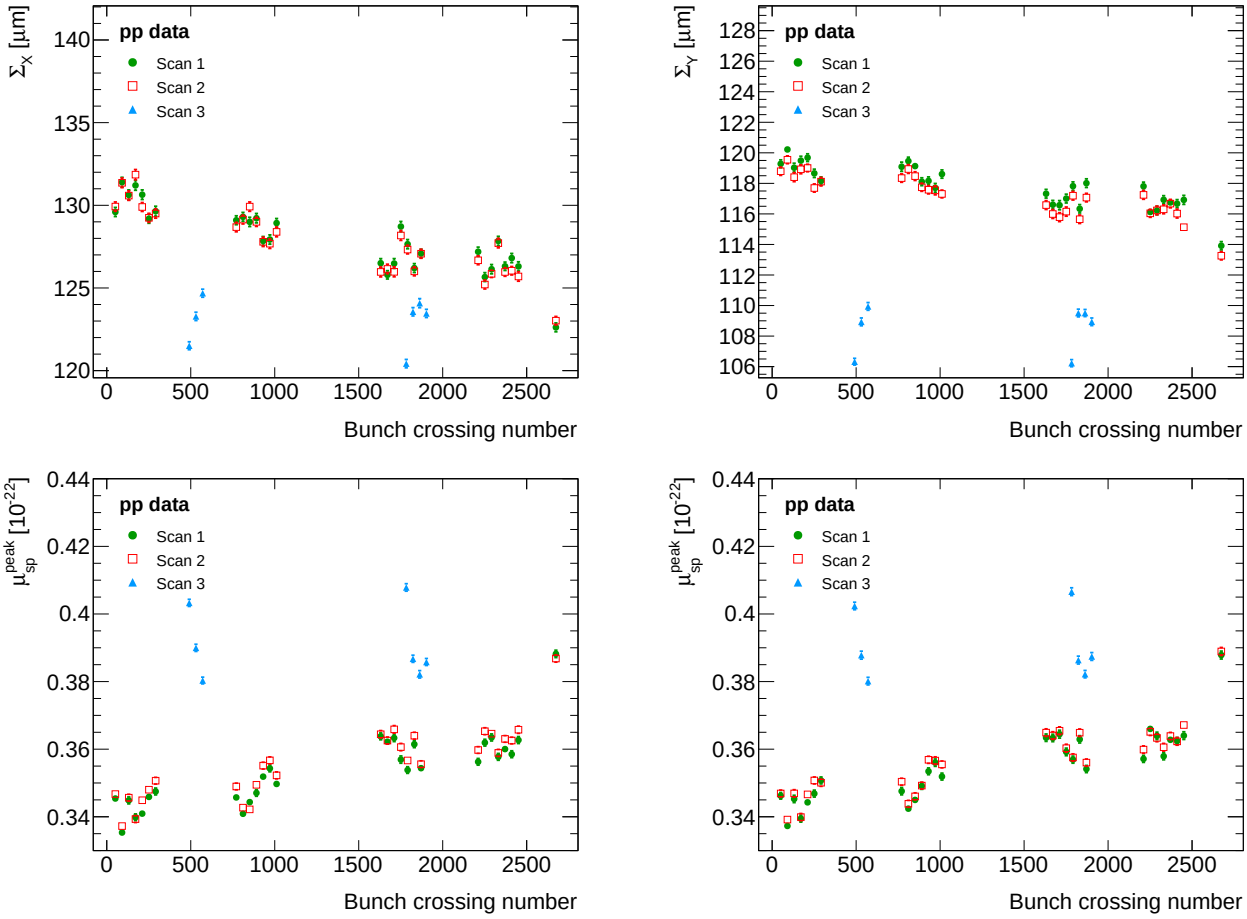

Figure 4.19: Uncorrected Σ and $\mu_{\mathrm{sp}}^{\mathrm{peak}}$ for each bunch crossing in the horizontal (left panels) and vertical (right panels) proton–proton scans. The error bars, which are typically smaller than the marker size, represent statistical uncertainty only.

Table 4.7 shows the bunch-averaged σ_{vis} in each of the three proton–proton scans. Only four inclusive algorithms are calibrated: LucidOR, LucidBiOR[12], BcmVOR, BcmHOR. The scan-averaged σ_{vis} for the LucidOR and LucidBiOR are rather different. As will be discussed in section 4.6.4, this is attributed to background from the bismuth source. The BCM algorithms show good agreement (differ by about 1%).

The proton–proton inelastic cross-section at 13 TeV is about 80 mb (as will be shown in chapter 5); therefore, the measured values of σ_{vis} translate to an efficiency of about 40% for LucidOR and about 25% for BcmHOR and BcmVOR. These differ substantially from the values in the 2013 proton–lead scans. This is due to the LUCID upgrade that lowered its acceptance, and the difference in collision species.

[12]LucidBiOR is similar to LucidOR but uses hits from the photomultipliers with bismuth sources.

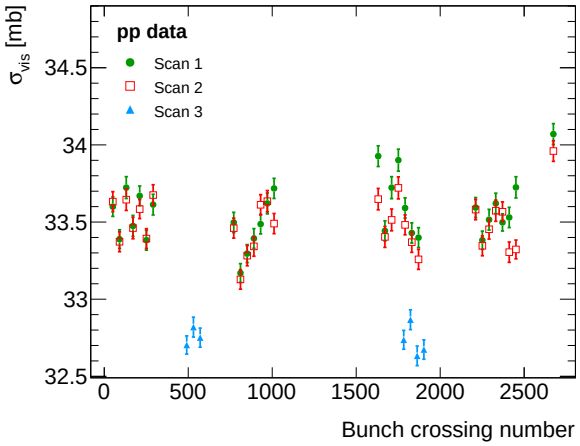

Figure 4.20: Uncorrected LucidOR σ_{vis} for each bunch crossing in all proton–proton scans. The error bars represent statistical uncertainty only.

Table 4.7: Uncorrected bunch-averaged σ_{vis} (in mb) for the calibrated algorithms in the proton–proton scans. The statistical uncertainties of these measurements are negligible.

Scan Number	1	2	3	Average
LucidOR	33.56	33.49	32.74	33.26
LucidBiOR	32.39	32.33	31.62	32.11
BcmVOR	19.36	19.32	18.90	19.19
BcmHOR	19.61	19.57	19.14	19.44

Lead–lead results

Figure 4.21 shows the Σ and $\mu_{\text{sp}}^{\text{peak}}$ results for the horizontal and vertical lead–lead scans. The data show various systematic patterns. First, a train dependence is visible in Σ_{X} and $\mu_{\text{sp}}^{\text{peak}}$ in both planes. The first trains have significantly larger Σ_{X} ($\approx$ 34–37 μm) than the later ones ($\approx$ 31–34 μm). This drives the smaller $\mu_{\text{sp}}^{\text{peak}}$ values observed in both planes. The origin of this pattern is unknown.

Second, a bunch-dependence common to all trains is observed. The difference between the measurement of the first and last bunch crossing is about 10–15%. This pattern can be attributed to emittance growth at the injection phase.

Third, the data also show signs of emittance growth from one scan to the next. The average Σ_{X} increased by about 2% and Σ_{Y} by about 0.6%. The rate of emittance growth is larger than in the proton–lead and proton–proton scans. This is probably due to a larger rate of intra-beam scattering in the lead beams.

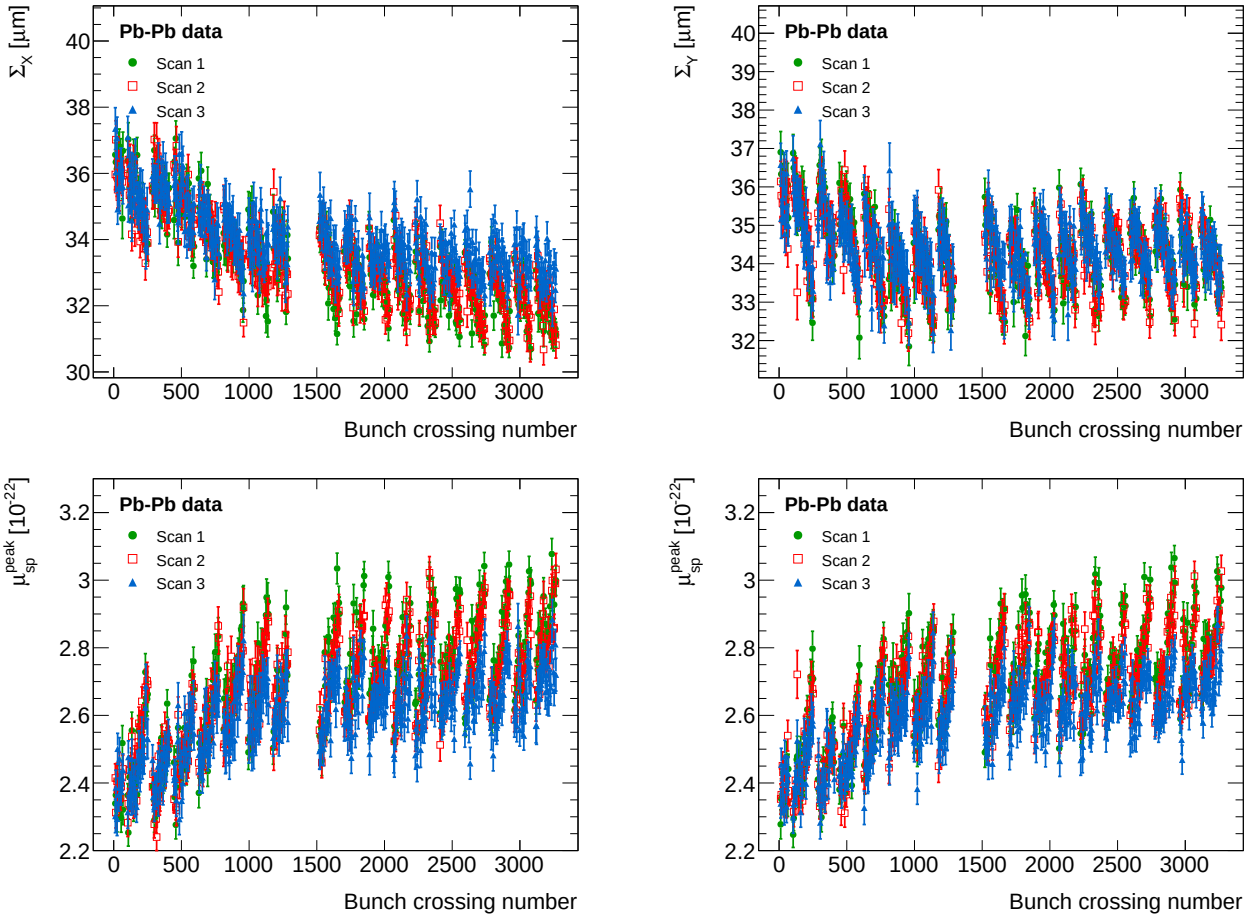

Figure 4.21: Uncorrected Σ and $\mu_{\mathrm{sp}}^{\mathrm{peak}}$ for each bunch crossing in the horizontal (left panels) and vertical (right panels) lead–lead scans. The error bars represent statistical uncertainty only.

Figure 4.22 shows the measured LucidOR σ_{vis} for each bunch crossing in all scans. The statistical uncertainty of each measurement is about 3%. This is significantly larger than in the proton–proton and proton–lead data. A fit with a constant function yields a reduced χ^2 in the range 2.2–2.3 depending on the scan.

Table 4.8 shows the bunch-averaged σ_{vis} values in all lead–lead scans. Only inclusive algorithms are calibrated: LucidOR, LucidBiOR, BcmVOR, BcmHOR. The results of the BCM algorithms are in good agreement. As will be shown in section 4.6.4, the large difference between the LucidOR and LucidBiOR results is attributed to background.

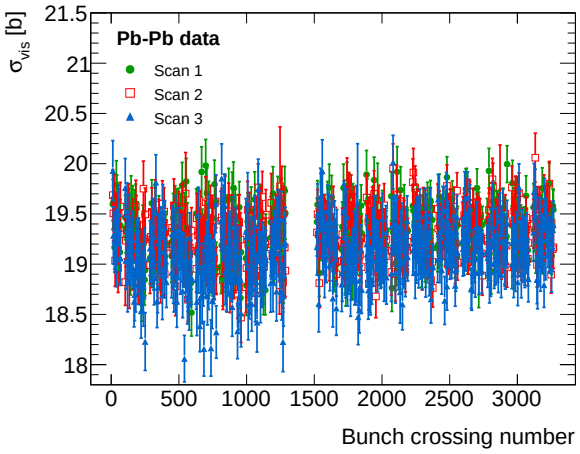

Figure 4.22: Uncorrected LucidOR σ_{vis} for each bunch crossing in all lead–lead scans. The error bars represent statistical uncertainty only.

Table 4.8: Uncorrected bunch-averaged σ_{vis} (in b) for the calibrated algorithms in the lead–lead scans. The statistical uncertainties of these measurements are negligible.

Scan Number	1	2	3	Average
LucidOR	19.39	19.27	19.19	19.28
LucidBiOR	21.54	21.61	21.95	21.70
BcmVOR	12.51	12.46	12.45	12.47
BcmHOR	12.58	12.53	12.51	12.54

The lead–lead inelastic cross-section is about 7.7 b [79]. This is much lower than any of the measured σ_{vis}. This indicates that the bulk of the rate observed by the luminometers is caused by electromagnetic interactions.

4.6.4 Corrections

A number of corrections are applied to ensure an unbiased σ_{vis} measurement. These include corrections to the beam separation, maximum specific rate, and bunch charge. The systematic uncertainties related to each correction are described in section 4.8.

Background subtraction

The main backgrounds are beam-gas interactions, instrumental noise and afterglow. Afterglow is background due to slow neutrons and photons from material activation from

primary collisions that cause hits in the luminometers in later bunch crossings. It is estimated, for each bunch crossing and scan step, as the rate in the preceding (empty) bunch crossing. This estimate also includes background due to instrumental noise.

Beam-gas background is caused by collisions between gas molecules in the LHC beam pipe and the beams. The estimate of the rate of beam-gas interactions is based on the observed rate in unpaired bunch crossings. To avoid double counting, the afterglow and noise background is subtracted in the same way as for the filled bunch crossings (rate in the preceding bunch crossing).

The raw number of events N_{OR} passing a given algorithm criteria is corrected as:

$$N_{\mathrm{OR}}^{\mathrm{subtracted}} = N_{\mathrm{OR}} - N_{\mathrm{afterglow}} - C_1 \times n_1 - C_2 \times n_2, \tag{4.16}$$

where $N_{\mathrm{afterglow}}$ is the afterglow and noise estimate, and C_1 and C_2 are constants that relate rate and bunch current (obtained from unpaired bunch crossings). The corrected number of counts, $N_{\mathrm{OR}}^{\mathrm{subtracted}}$, is then used in equation 4.12 to recalculate μ_{vis}.

Proton–lead results

Figure 4.23 shows the bunch-averaged scan curves of the LucidOR, BcmHOR, LucidA and LucidC algorithms and corresponding background estimates of scan 1. The LucidOR scan curve peaks at $\mu_{\mathrm{vis}} \approx 4 \times 10^{-2}$ and has tails that flatten out in the outermost points at a level of about $\mu_{\mathrm{vis}} \approx 10^{-5}$. This background-like feature is consistent with the corresponding beam-gas estimate that is independent of beam separation (as expected). The afterglow estimate shows a Gaussian-like core that resembles the scan curves (as expected) but has tails that flatten out at $\mu_{\mathrm{vis}} \approx 10^{-6}$. The flattening out, which is more pronounced than for the scan curve, can be attributed to the dominance of noise that is about $\mu_{\mathrm{vis}} \approx 10^{-6}$. The afterglow peak is about $\mu_{\mathrm{vis}} \approx 3 \times 10^{-5}$ and corresponds to about 0.1% of the peak of the scan curve.

The BcmHOR data show a similar pattern to the LucidOR one, but with a peak value about 4 times smaller. The beam-gas estimate is about $\mu_{\mathrm{vis}} \approx 8 \times 10^{-6}$ and describes well the tails of the scan curve. The afterglow estimate suggests that the noise of this algorithm is about $\mu_{\mathrm{vis}} \approx 3 \times 10^{-6}$. The afterglow background is a sub-dominant contribution except at the peak of the scan curve where it is slightly larger than the beam-gas estimate and about 0.1% of the peak of the scan curve. The BcmVOR data (not shown in the figure 4.23) are consistent with the BcmHOR, as expected.

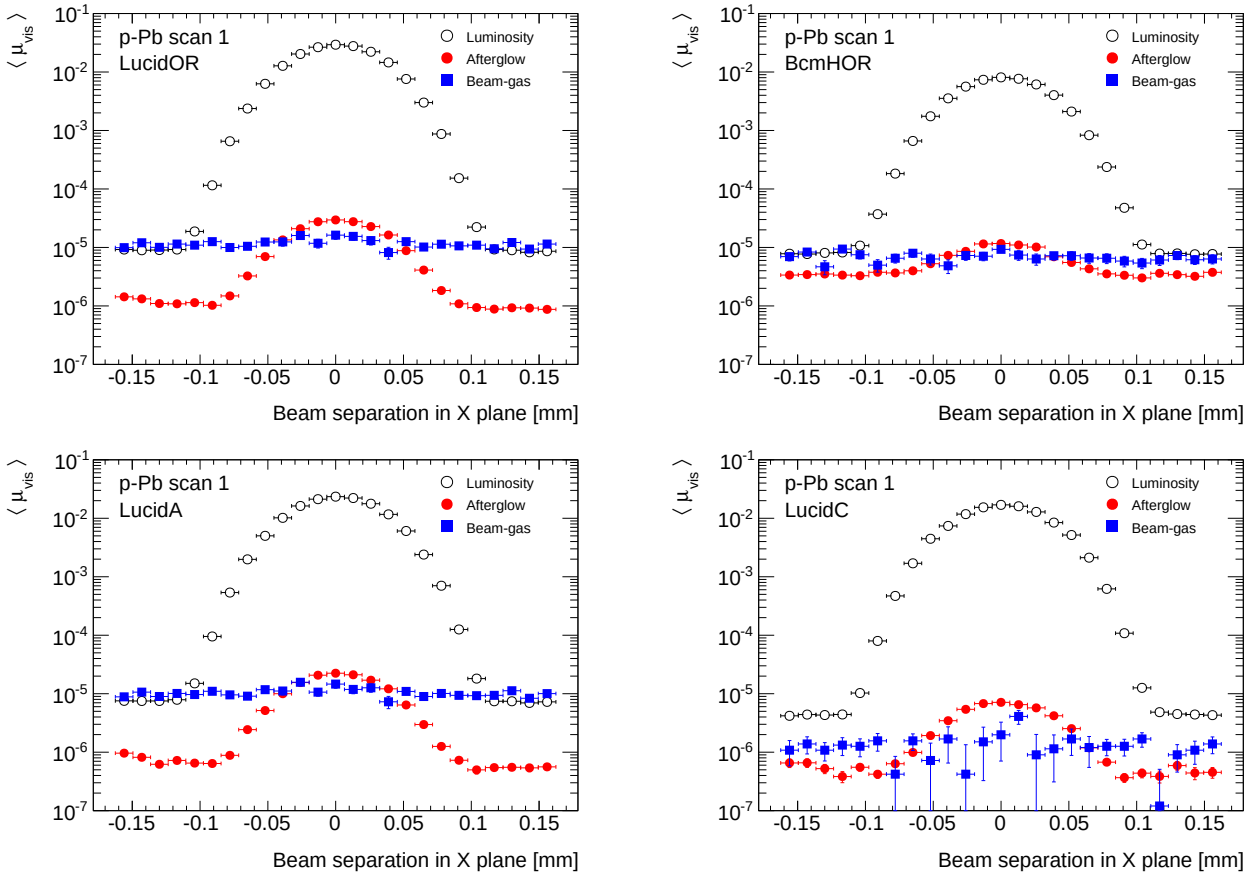

Figure 4.23: Bunch-averaged μ_{vis} and estimates of afterglow and beam-gas interactions as a function of beam separation in the first proton–lead scan. The algorithms shown are: LucidOR (upper left), BcmHOR (upper right), LucidA (lower left) and LucidC (lower right). The error bars represent statistical uncertainty only.

The LucidA and LucidC data present several differences that can be attributed to the asymmetry of the proton–lead collisions. In the January scans the lead beam travelled from C to A side; thus, the background in the LucidA is expected to be larger. The background estimate does not seem to be completely satisfactory for the LucidC case. The reasons for this are not understood. The results of February scans (not shown in the figure 4.23) show the opposite behaviour.

The impact of the background subtraction on the σ_{vis} determination is summarized in table 4.9. The largest impact is for the BCM algorithms that decreased with the correction by up to 0.3%.

Table 4.9: Impact of the background subtraction on the bunch-averaged σ_{vis} of the proton–lead scans. The range of values covers the results of all scans.

Algorithm	$\Delta\sigma_{\text{vis}}/\sigma_{\text{vis}}$
LucidOR	$-(0.1\text{–}0.2)\%$
LucidA	$-(0.1\text{–}0.2)\%$
LucidC	$-(0.1\text{–}0.2)\%$
BcmHOR	$-(0.2\text{–}0.3)\%$
BcmVOR	$-(0.2\text{–}0.3)\%$

Proton–proton results

Figure 4.24 shows the bunch-averaged scan curves of the LucidOR, LucidBiOR, BcmHOR and BcmVOR algorithms and the corresponding background estimates for scan 1. The LucidOR data peak at $\mu_{\text{vis}} \approx 0.2$ and reach a minimum of $\mu_{\text{vis}} \approx 2 \times 10^{-5}$. The tails do not show signs of flattening out. The beam-gas estimate is a constant $\mu_{\text{vis}} \approx 2 \times 10^{-5}$; thus, it dominates the last points in the tail of the scan curve. The afterglow estimate shows a Gaussian-like core that peaks at $\mu_{\text{vis}} \approx 10^{-5}$ and a noise level of $\mu_{\text{vis}} \approx 10^{-6}$.

The LucidBiOR scan curve peak value is similar to the LucidOR one but presents a much more pronounced flattening of the tails at a level of $\mu_{\text{vis}} \approx 4 \times 10^{-4}$. This is well described by the afterglow estimate that seems dominated by noise. This noise is much higher than the one observed for LucidOR. This can be attributed to the activity of the bismuth source that is recorded by the photomultipliers used by LucidBiOR.

The data of the BcmHOR and BcmVOR algorithms are very similar, as expected. The luminosity curves reach $\mu_{\text{vis}} \approx 10^{-1}$ and show a flattening out at the last two outer points in the tails. These are consistent with a noise level of about $\mu_{\text{vis}} \approx 2 \times 10^{-4}$. This is significantly larger than the one observed in the proton–lead data (see figure 4.23). The difference may be attributed to radiation-induced damage during the high-luminosity proton–proton runs in 2015.

Table 4.10 summarizes the impact of the background subtraction on the σ_{vis} determination. The largest impact is for the LucidBiOR results that decreased with the correction by about 0.3%.

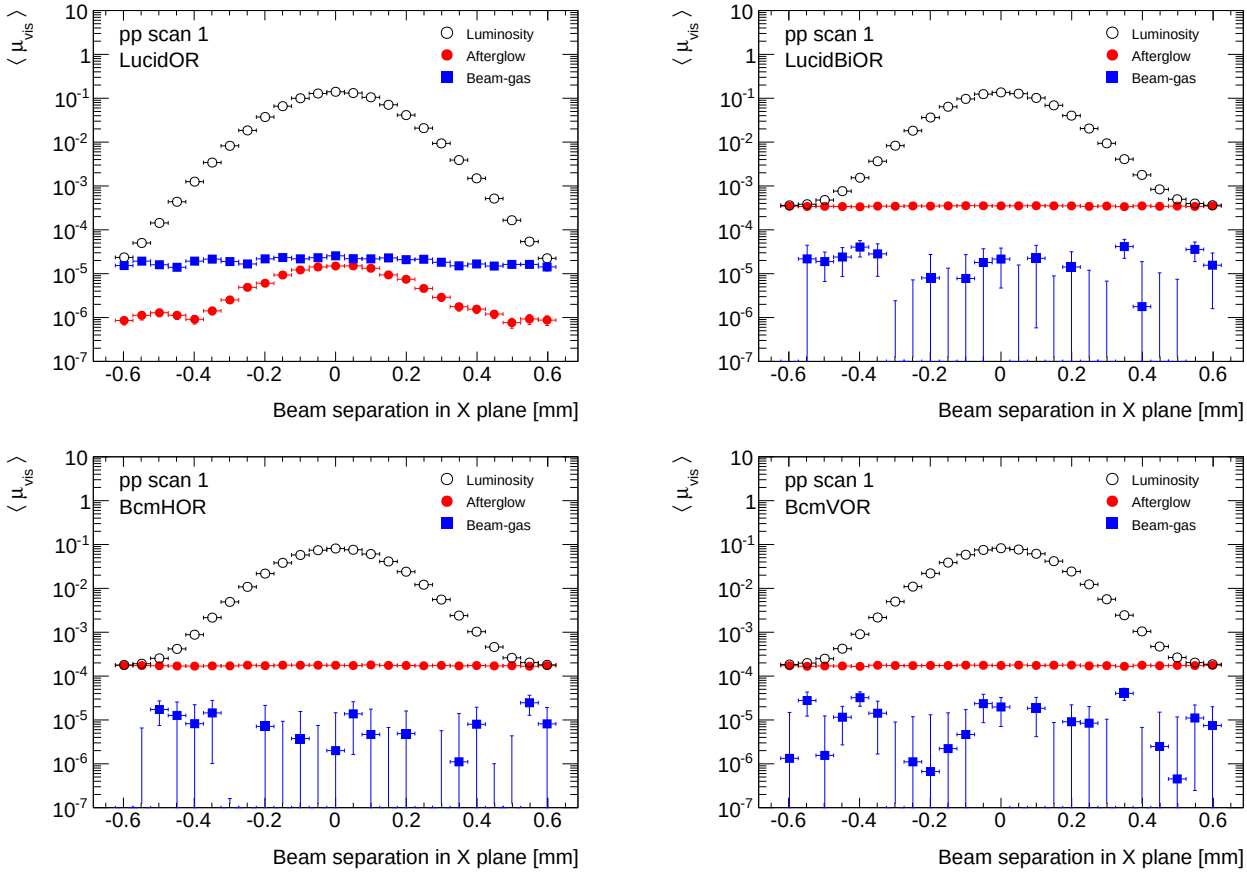

Figure 4.24: Bunch-averaged μ_{vis} and estimates of afterglow and beam-gas interactions as a function of beam separation in the first proton–proton scan. The algorithms shown are: LucidOR (upper left), LucidBiOR (upper right), BcmHOR (lower left) and BcmVOR (lower right). The error bars represent statistical uncertainty only.

Table 4.10: Impact of the background subtraction on the bunch-averaged σ_{vis} of the proton–proton scans. The range of values covers the results of all scans.

Algorithm	$\Delta\sigma_{\text{vis}}/\sigma_{\text{vis}}$
LucidOR	$< 0.1\%$
LucidBiOR	$-(0.2\text{–}0.3)\%$
BcmHOR	$-(0.1\text{–}0.2)\%$
BcmVOR	$-(0.1\text{–}0.2)\%$

Lead–lead results

Figure 4.25 shows the bunch-averaged scan curves of the LucidOR, LucidBiOR, BcmHOR and BcmVOR algorithms and the corresponding background estimates of scan 1.

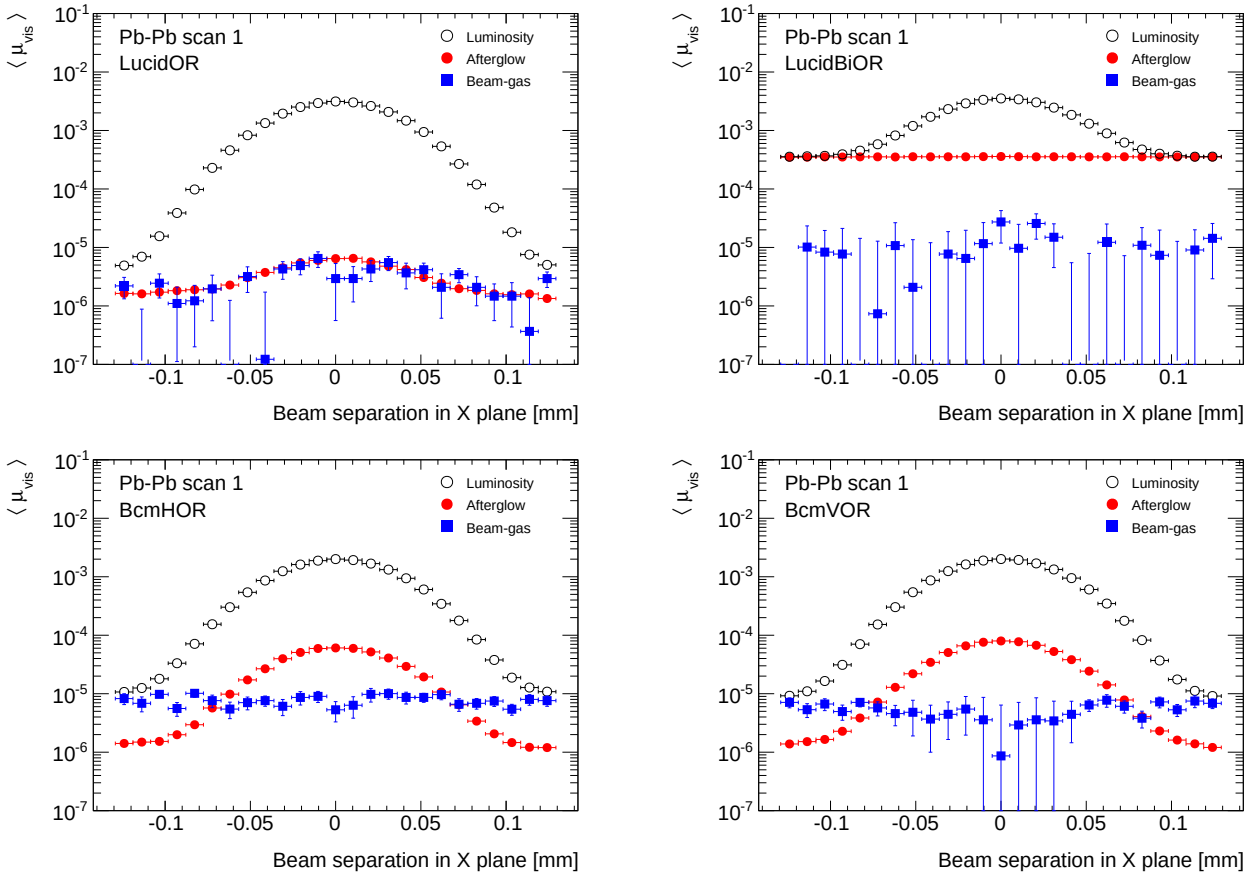

Figure 4.25: Bunch-averaged μ_{vis} and estimates of afterglow and beam-gas interactions as a function of beam separation in the first lead–lead scan. The algorithms shown are: LucidOR (upper left), LucidBiOR (upper right), BcmHOR (lower left) and BcmVOR (lower right). The error bars represent statistical uncertainty only.

The LucidOR scan curve peaks at $\mu_{\text{vis}} \approx 3 \times 10^{-3}$ and the tails reach $\mu_{\text{vis}} \approx 5 \times 10^{-6}$. The afterglow estimate peaks at about 0.3% of the scan curve and shows a noise level of $\mu_{\text{vis}} \approx 2 \times 10^{-6}$. The beam-gas estimate is similar to the noise level. The LucidBiOR data show a similar peak value to LucidOR but a much pronounced flattening out that affects about 8 points of the scan curve. The afterglow estimate is dominated by the bismuth activity that is about 10% of the scan curve peak.

The data of the BcmHOR and BcmVOR algorithms are similar. The scan curves peak at $\mu_{\text{vis}} \approx 2 \times 10^{-3}$ and the tails reach $\mu_{\text{vis}} \approx 10^{-5}$. The beam-gas estimate is about $\mu_{\text{vis}} \approx 10^{-5}$ and describes this feature. The afterglow estimate peaks at about 3% of the scan curve peak and shows a noise of $\mu_{\text{vis}} \approx 10^{-6}$.

Table 4.11 summarizes the impact of the background subtraction on the σ_{vis} determination. The largest impact is for the LucidBiOR (-10%).

Table 4.11: Impact of the background subtraction on the bunch-averaged σ_{vis} of the lead–lead scans. The range of values covers the results for all scans.

Algorithm	$\Delta\sigma_{\text{vis}}/\sigma_{\text{vis}}$
LucidOR	$-(0.4\text{–}0.5)\%$
LucidBiOR	$-(10\text{–}11)\%$
BcmHOR	$-(3.2\text{–}3.4)\%$
BcmVOR	$-(4.2\text{–}4.4)\%$

Beam–beam deflections

Beam–beam deflections refer to distortions of the beam orbits caused by the mutual electromagnetic repulsion between colliding bunches. This distorts the shape of the scan curves and thus biases the measurement of the overlap integral.

The beam–beam deflections are calculated analytically using the results described in Ref. [80]. This calculation assumes that the transverse densities of the beams are described by Gaussian functions. The inputs for the calculation are: n_1, n_2, β_x, β_y, Σ_{X}, and Σ_{Y}. The calculation is performed for each bunch crossing.

Figure 4.26 shows the bunch-averaged beam deflection for the first proton–proton scan. The deflection reaches ± 2 μm at about ± 200 μm beam separation. It tends to increase the absolute separation of the beams, as expected. The estimates for different scans are similar, which shows that it does not depend strongly on the bunch currents.

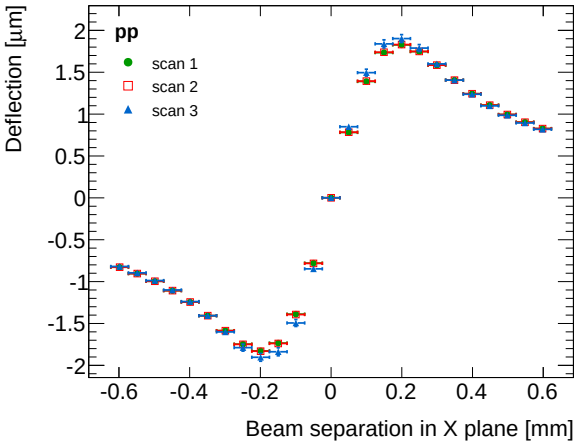

Figure 4.26: Bunch-averaged calculation of beam–beam deflection as a function of beam separation during proton–proton scans. The error bars represent the standard deviation over all per-bunch calculations.

The beam–beam deflection correction effectively stretches the scan curves (increases Σ). The shape of the correction is similar in all datasets analysed, but the peak-to-peak amplitude is different: it is ≈ 0.1 μm for the proton–lead scans, ≈ 4 μm for the proton–proton scans, and less than 0.1 μm for the lead–lead scans.

The beam separation is corrected at each step by the estimated deflection, D:

$$\delta(s) \to \delta(s) + D(s). \tag{4.17}$$

Then the analysis presented in section 4.6 is repeated with the corrected beam separation. In principle there is also a deflection in the out-of-scan plane (i.e. vertical deflection during the horizontal scan and vice versa). However, this effect is only a small fraction (percent level) of the in-plane correction so it is neglected.

Table 4.12 summarizes the input for the beam–beam calculation, the resulting peak-to-peak amplitude and the impact on the bunch-averaged σ_{vis}. The largest correction corresponds to the proton–proton scans where σ_{vis} increases by about 2%. This correction affects all algorithms in the same manner.

Table 4.12: Typical bunch parameters used as input for the calculation of beam–beam deflections; estimated deflection; and impact on the bunch-averaged σ_{vis} for each dataset analysed.

	proton–lead	proton–proton	lead–lead
Approx. Σ_{X} [μm]	30	130	30
Approx. Σ_{Y} [μm]	40	120	35
Approx. n_1 [10^{11} protons]	1	1	0.1
Approx. n_2 [10^{11} protons]	0.1	1	0.1
Peak-to-peak deflection [μm]	≈ 0.1	≈ 4	≈ 0.06
$\Delta\sigma_{\text{vis}}/\sigma_{\text{vis}}$	+(0.2–0.3)%	+(1.9–2.1)%	< 0.1%

Non-linear response of BPTX (FBCT)

The bunch charge relative measurements are corrected for a possible non-linear response of the instrument used to measure them (BPTX is used by default) using a method described in Ref. [81]. This method assumes that the non-linear response can be approximated by a linear polynomial near the mean bunch current. Given that the measurements are relative, a constant offset fully determines the slope of the polynomial.

The measurements of the bunches in the two beams are independent, so two possible offsets (b_1 and b_2) are considered. These offsets change the per-bunch relative fractions but do not change the absolute total charge because that is normalized to the DCCT scale. In order to determine the offsets, the following χ^2 function is minimized:

$$\chi^2(b_1, b_2) = \sum_{i}^{\text{bunches}} \left(\sigma^i_{\text{vis}}(b_1, b_2) - \bar{\sigma}\right)^2 / \sigma_i^2, \tag{4.18}$$

where $\sigma^i_{\text{vis}}(b_1, b_2)$ is the visible cross-section for the i-th bunch crossing calculated with the offsets b_1 and b_2; $\bar{\sigma}$ is a free parameter of the fit; and σ_i is the uncertainty on σ^i_{vis}. That is, the minimization finds the offsets b_1 and b_2 that make the σ_{vis} values more consistent among the different bunch crossings.

Proton–lead

The fitted offsets found in the proton–lead data are in the range $-0.3 < b_1 < 0$ and $-0.3 < b_2 < 0.2$ (10^9 protons). The statistical uncertainties on these parameters are negligible. These values are a few percent of the typical bunch charge.

Figure 4.27 shows the σ_{vis} for each bunch crossing with and without the offset correction in scans 3 and 5. The correction improves significantly the bunch-to-bunch consistency in both scans. A constant fit to the data yields a reduced χ^2 of 2.6–3.9 before the correction and 1.5–1.9 after the correction, depending on the scan.

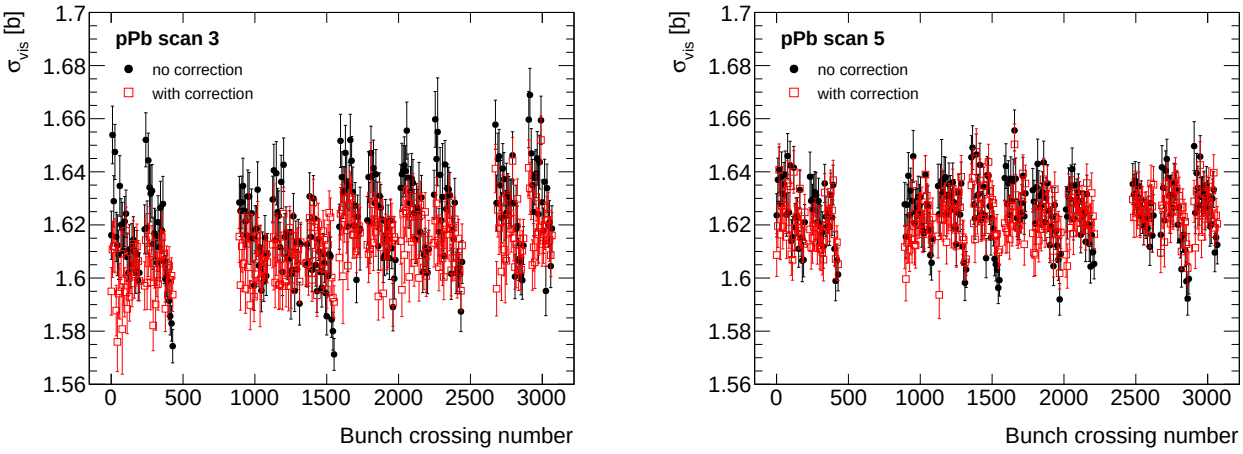

Figure 4.27: LucidOR σ_{vis} for each bunch crossing with and without offset correction in proton–lead scans 3 (left panel) and 5 (right panel). The error bars represent statistical uncertainty only.

While the partial elimination of the bunch-dependent systematic pattern validates the hypothesis upon which the method is based, it is not the only evidence. The correction also reduces the correlation between σ_{vis} and $n_1 n_2$ that is shown in figure 4.28. This correlation is not an input of the offset fit correction.

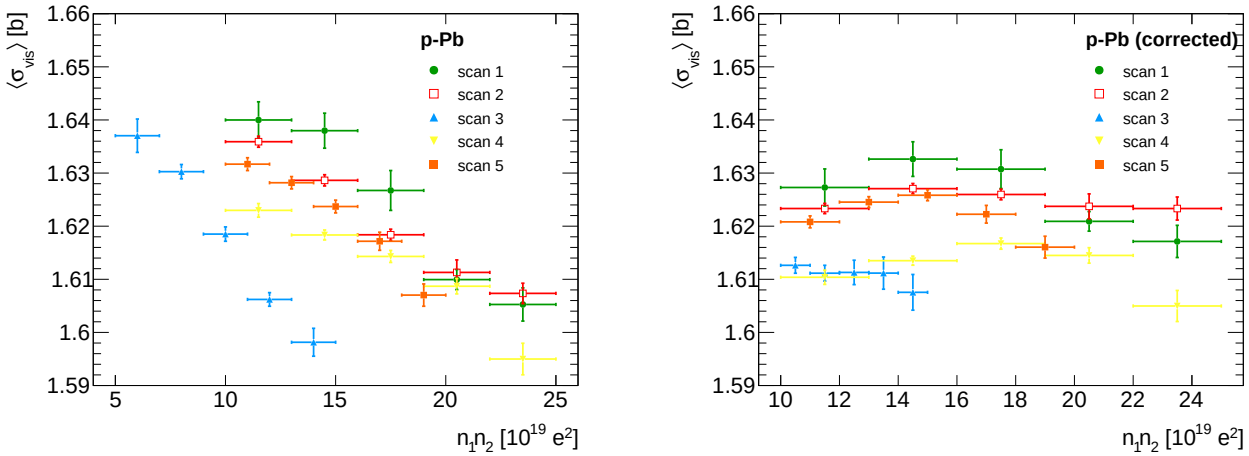

Figure 4.28: Bunch-averaged LucidOR σ_{vis} as a function of $n_1 n_2$ with (right panel) and without (left panel) offset correction in proton–lead scans. The error bars represent statistical uncertainty only.

The maximum deviation between the bunch-averaged σ_{vis} with and without offset correction is -0.6% (scan 3) in the January scans and -0.2% (scan 5) in the February scans. These results are an order of magnitude larger than the ones reported in any previous analysis, e.g. for the 2012 proton–proton analysis the effect is 0.02% [81]. This may be explained by the much smaller bunch charge in the 2013 proton–lead scans.

Proton–proton

The fitted offsets in the proton–proton scans are in the range $-3.5 < b_1 < -2.0$ and $0 < b_2 < 1.6$ (10^9 protons) depending on the scan. These correspond to about 2–3% of the typical bunch charge. Figure 4.29 shows the σ_{vis} for each bunch crossing with and without the offset correction in scans 1 and 3. A constant fit to the scan 1 (scan 3) data yields a reduced χ^2 of 9.3 (1.6) before and 6.6 (1.5) after the correction. This suggests the systematic bunch variation observed in scans 1 and 2 is not entirely driven by a non-linear response of the BPTX. The impact on the bunch-averaged σ_{vis} is negligible.

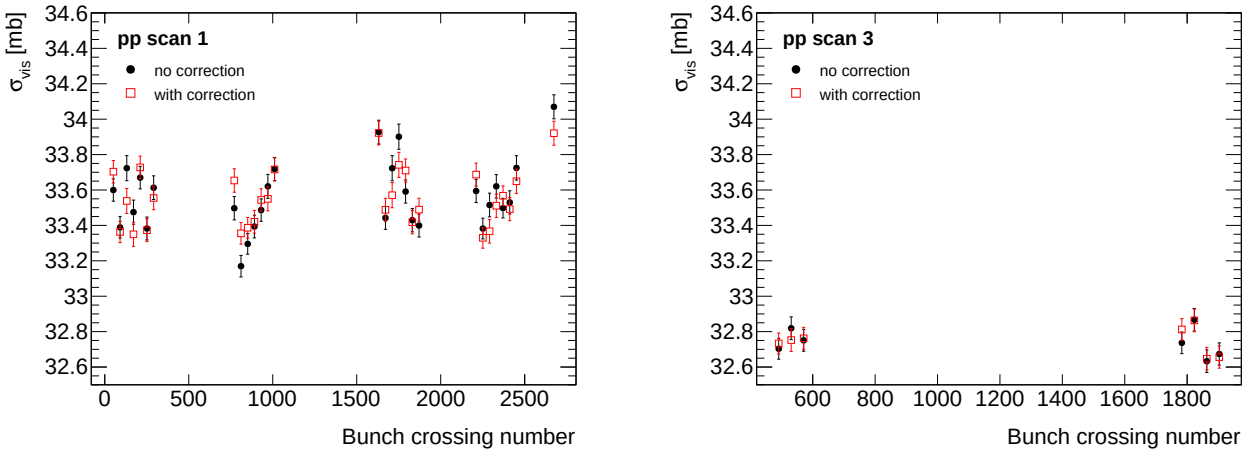

Figure 4.29: LucidOR σ_{vis} for each bunch crossing with and without offset correction in proton–proton scans 1 (left panel) and 3 (right panel). The error bars represent statistical uncertainty only.

The correlation between the measured σ_{vis} and $n_1 n_2$ is shown in figure 4.30. The level of the correlation is rather small and the offset correction has little impact, which again suggest that the bunch-by-bunch variations are not related to the $n_1 n_2$ measurement.

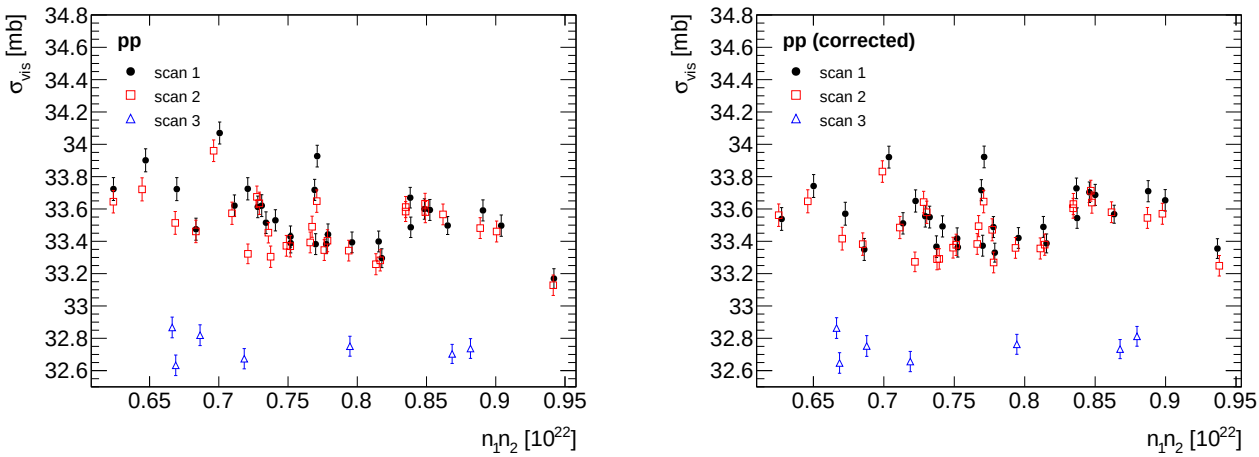

Figure 4.30: LucidOR σ_{vis} as a function of $n_1 n_2$ with (right panel) and without (left panel) offset correction in proton–proton scans. The error bars represent statistical uncertainty only.

Lead–lead

The fitted offsets in the lead–lead scans are in the range $-0.4 < b_1, b_2 < -0.3$ (10^9 protons). These correspond to about 3% of the typical bunch charge. Figure 4.31 shows the σ_{vis} for each bunch crossing with and without the offset correction in scans 1 and 3.

Even though the data do not show a clear systematic pattern, the reduced χ^2 of a constant fit decreased from 2.0–2.2 to 1.2–1.3 depending on the scan.

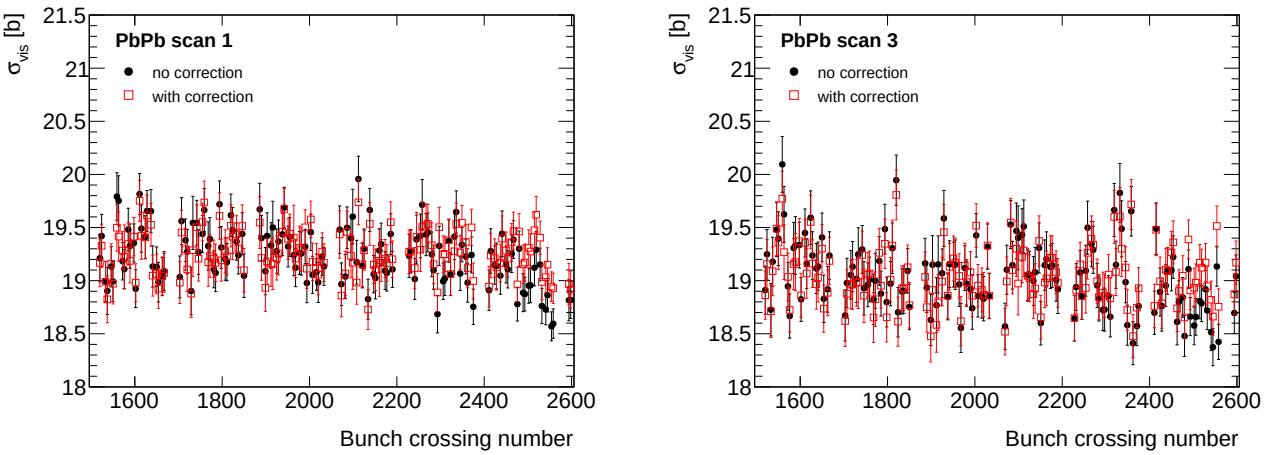

Figure 4.31: LucidOR σ_{vis} for each bunch crossing with and without offset fit correction in lead–lead scans 1 (left panel) and 3 (right panel). The error bars represent statistical uncertainty only. Only a subset of all bunch crossings are shown for clarity.

The offset correction has a large impact on the observed σ_{vis}–$n_1 n_2$ correlation, shown in figure 4.32. A strong negative correlation is clearly visible in all scans (even though the scan 3 is performed in a different fill). After the correction the slope of the correlation is consistent with zero within statistical uncertainties. The impact on the bunch-averaged σ_{vis} is about $-(0.1\text{–}0.2)\%$ depending on the scan.

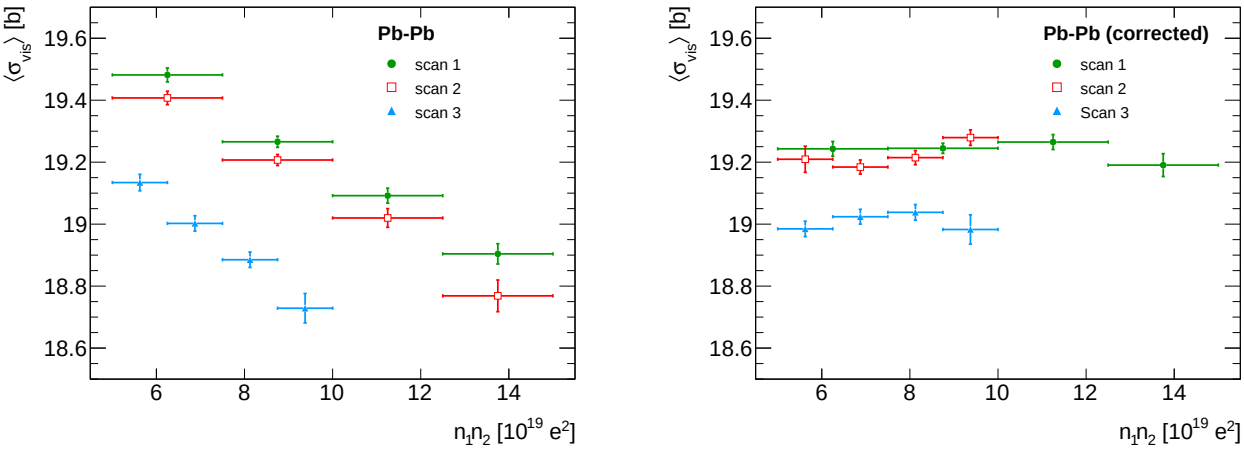

Figure 4.32: Bunch-averaged LucidOR σ_{vis} as a function of $n_1 n_2$ with (right panel) and without (left panel) offset correction in lead–lead scans. The error bars represent statistical uncertainty only.

Orbit drift correction

Instabilities in the orbit of the colliding beams lead to changes in the beam separation that bias the σ_{vis} measurement. This section summarizes the orbit drift correction in the 2015 proton–proton run that is detailed in Ref. [4]. Two independent beam-position monitors (BPMs) were fully operational during the 2015 proton–proton vdM scans: the arc-BPM and DOROS systems.

The arc-BPM system [82] is based on an array of electrodes that pick up the electric field of the beams. They are located in the LHC arcs ($|z| > 500$ m). The measurements at the electrodes locations are fitted to an orbit and extrapolated to the interaction region considering the magnetic lattice of the LHC. The arc-BPM system has been used throughout Run 1 for machine-protection purposes, and in previous vdM analysis [81, 83].

The DOROS system [84] is based on diode peak detectors coupled to electrode BPMs. The diode detectors transform the fast pulses induced in the electrodes to a slowly varying signal (O(Hz)) that can be digitized with high resolution. The DOROS system is embedded into the LHC collimators near the interaction point (at $|z| = 20$ m). It was commissioned during 2015.

Figure 4.33 shows the beam separation, obtained with DOROS data, as a function of time during scan 3. During the horizontal scan (first half of the period) the vertical separation remained close to zero and the horizontal separation increased linearly from about -600 μm to $+600$ μm, and vice-versa for the vertical scan.

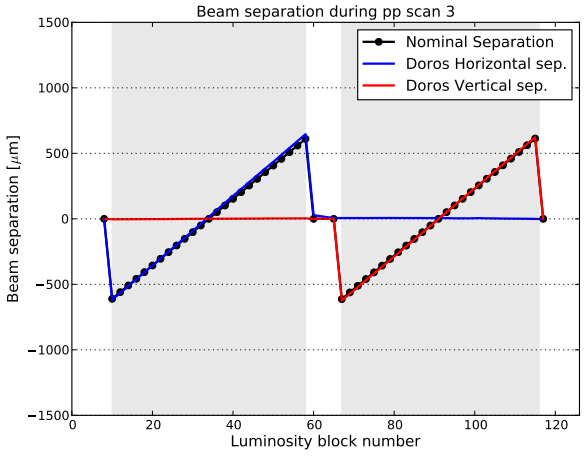

Figure 4.33: Beam separation as a function of luminosity block number, i.e. time, during the third proton–proton scan. The shaded bands represent the periods of the vdM scans.

The nominal beam separation is subtracted from the DOROS data to obtain the beam drift. Given that only relative beam separations are relevant for the vdM analysis, the drift is reported relative to the beam separation at the beginning of the scans.

Figure 4.34 shows the drift data as a function of time during scan 3. The horizontal drift increases roughly linearly with time during the horizontal scan and reaches about 27 μm, which is about 20% of the measured Σ_X. It is much less pronounced during the vertical scan, where a relative -7 μm drift is observed. The vertical drift is less than 3 μm over the whole scan. The agreement between the DOROS and arc-BPM data is excellent during the whole scan. The orbit drift during scans 1 and 2 are within ± 5 μm and show similar level of agreement as in scan 3.

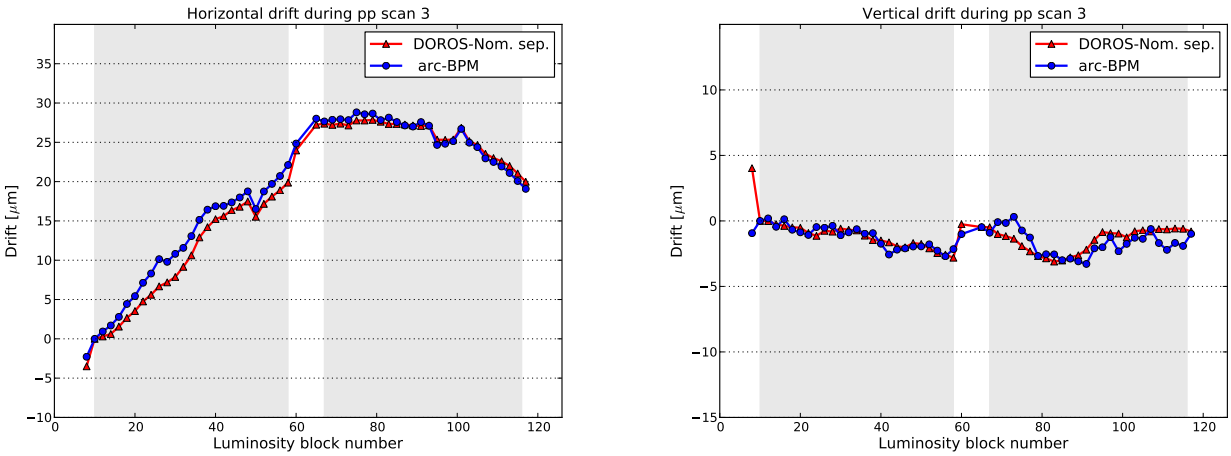

Figure 4.34: Horizontal (left panel) and vertical (right panel) orbit drift as a function of luminosity block number, i.e time, during the third proton–proton scan. The shaded bands represent the periods of the vdM scans.

The impact of orbit drift on the vdM calibration procedure is twofold: first, the drift in the direction that is being scanned distorts the shape of the scan curves and thus the extracted Σ; second, the drift in the orthogonal direction that is being scanned reduces the beam-overlap and thus the observed $\mu_{\mathrm{sp}}^{\mathrm{peak}}$. The former is taken into account by correcting the nominal separation, and the latter by considering an average non-scan plane separation, as explained in Ref. [81].

The corrections to the bunch-averaged σ_{vis} are -0.2%, $+0.1\%$ and $+2.1\%$ for scans 1, 2 and 3 respectively. The correction significantly improves the consistency of the results of scans. Table 4.13 summarizes the impact of the orbit drift correction on the

scan-averaged σ_{vis} on all datasets. Analysis of orbit drift data of the 2013 proton–lead and 2015 lead–lead runs are detailed in Refs. [3, 6].

Table 4.13: Impact of the orbit drift correction on the σ_{vis} determination.

	proton–lead	proton–proton	lead–lead
$\Delta\sigma_{\text{vis}}/\sigma_{\text{vis}}$	-0.3%	$+0.7\%$	—

Length scale correction

The vdM technique relies on absolute measurements of the beam separation at each scan point to measure Σ. The nominal beam displacement that is reported by the LHC systems (that ultimately relates to the current in the steering magnets) is calibrated to the actual beam displacement in dedicated "length scale scans". These scans are performed by displacing the colliding beams together in a stepwise manner over a range up to about $\pm 3\sigma_{\text{b}}$. Given that the beams remain in head-on collisions, their position can be measured with high accuracy by studying the spatial distribution of vertices reconstructed with the inner detector [49].

Length scale scans are performed close in time to each vdM scan. The analysis of the data relevant for this work is documented in Ref. [3] (proton–lead data) and Ref. [5] (proton–proton data). No dedicated analysis has been finalized in the lead–lead case.

No correction to the nominal separation is performed, but a multiplicative factor is applied to the measured Σ_{X} Σ_{Y} or equivalently to the measured σ_{vis}. Table 4.14 summarizes the length-scale correction for each analysed dataset.

Table 4.14: Impact of the length scale correction on the σ_{vis} determination.

	proton–lead	proton–proton	lead–lead
$\Delta\sigma_{\text{vis}}/\sigma_{\text{vis}}$	-2.1%	$+0.2\%$	—

Ghost charge and satellite bunches

The main sources of background for the $n_1 n_2$ measurement are "ghost charge" and "satellite charge". Ghost charge refers to protons (or ions) in nominally empty bunch

crossings. Satellite charge are protons (or ions) in radio-frequency buckets that are at least one period (2.5 ns) away from the filled bunches. These experience at most long-range encounters with the filled bunches in the other beam.

Both satellites and ghost charge can be measured with the LHC longitudinal density monitor (LDM) [85]. This instrument uses avalanche photo-diodes with 90 ps timing resolution to detect synchrotron radiation from the beams. These measurements are used to infer the ghost and satellite charge relative to the charge in the filled bunches.

For the LHC fills of the proton–lead scans, the ghost charge correction to $n_1 n_2$ ranges from -0.7 to -0.9% [3]. Data taken during the proton–proton scans show a negligible ghost charge [5]. Data taken during the lead–lead scans show a ghost charge of 3.5–4% in beam 1 and about 2.2–2.6% in beam 2 [6]. In both cases the ghost charge is observed to increase linearly with time. This can be attributed to intra-beam scattering that effective translates charge from the filled bunches to empty bunches.

The total satellite charge fraction in the lead–lead scans is about 0.7–0.9% in beam 1 and about 0.5–0.7% in beam 2; in the proton–lead scans it is about 0.1% in each beam; in the proton–proton scans it is negligible. The multiplicative factors that are applied to the measured $n_1 n_2$ at every scan point affect directly the σ_{vis} values. The corrections to σ_{vis} for each scan are shown in table 4.15.

Table 4.15: Impact of the ghost and satellite charge corrections on the σ_{vis} determination.

	proton–lead	proton–proton	lead–lead
Ghost charge	+0.8%	—	+5.8%
Satellite charge	+0.2%	—	+1.5%

4.7 Corrected results

The impact of each of the corrections applied to the vdM data is summarized in table 4.16. The total correction is -1.7% for the proton–lead data, $+2.8\%$ for the proton–proton data, and $+6.6\%$ for the lead–lead data. The dominant sources of correction are the length scale calibration, beam–beam deflection, and satellite and ghost charge, respectively.

Table 4.16: The relative effect of each correction, described in the text, on the LucidOR scan-averaged σ_{vis} in proton–lead, proton–proton, and lead–lead scans.

	proton–lead	proton–proton	lead–lead
Uncorrected σ_{vis}	1.623 b	33.26 mb	19.28 b
BPTX offset	−0.3%	< 0.1%	−0.2%
Beam–beam deflection	+0.2%	+2.0%	—
Orbit drift	−0.3%	+0.7%	—
Background subtraction	−0.2%	−0.1%	−0.5%
Length scale	−2.1%	+0.2%	—
Ghost charge	+0.8%	—	+5.8%
Satellite charge	+0.2%	—	+1.5%
Fully corrected σ_{vis}	1.595 b	34.20 mb	20.55 b

4.8 Systematic uncertainties of vdM analysis

This section describes the systematic uncertainties of the vdM analysis. The total uncertainty is divided into three categories that relate to the measurement of the bunch charge (section 4.8.1), beam conditions (section 4.8.2) and instrumental effects (section 4.8.3).

4.8.1 Bunch charge

The systematic uncertainties related to the $n_1 n_2$ measurement arise from the DCCT absolute scale, the relative per-bunch fractions, and ghost and satellite charges. These are described in the following sections and summarized in table 4.17.

DCCT scale variation

Dedicated studies of the DCCT data taken during the 2013 proton–lead run estimated its systematic uncertainty as $\pm 0.3\%$ [3]. The performance of the DCCT during the 2015 proton–proton run was at the level of the Run 1 or better[13]. An uncertainty of $\pm 0.2\%$ was estimated in Ref. [5]. No dedicated studies have been performed for the 2015 lead–lead run. But given that the total beam charge is similar to that in the 2015

[13]The DCCT readout system was upgraded before the Run 2.

proton–proton run and no instrumentation changes occurred, the same uncertainty is assigned for the lead–lead data [6].

Bunch charge fractions

The uncertainty on the bunch charge fractions is estimated by comparing results based on the FBCT and BPTX. The offset correction is applied to both. Figure 4.35 shows the $n_1 n_2$–σ_{vis} correlation in the first lead–lead scan with and without offset correction.

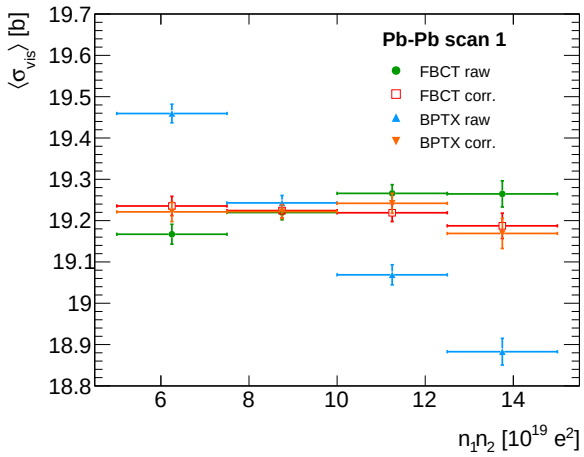

Figure 4.35: Correlation between the measured LucidOR σ_{vis} and $n_1 n_2$ with and without offset correction during the first lead–lead scan. The error bars represent statistical uncertainty only.

The corrected results are in excellent agreement. A similar level of agreement is observed in scans 1 and 2. Given that the offset correction is performed independently for both devices, this demonstrates the robustness of the correction. Studies of proton–proton and proton–lead data show a similar level of agreement [3, 5]. The bunch-averaged σ_{vis} obtained with BPTX and FBCT differ by at most $\pm 0.2\%$ in all datasets. These differences are taken as systematic uncertainty on the σ_{vis} determination.

Ghost and satellite charge

The methodology used to estimate the systematic uncertainty on the LHC LDM measurements is described in Ref. [85]. The systematic uncertainty on the ghost (satellite) charge measurement was estimated as $\pm 0.2\%$ ($< 0.1\%$) for the proton–lead data, and $\pm 0.5\%$

($\pm 0.7\%$) for the lead–lead data. Given the absolute correction for the proton–proton data is negligible, its uncertainty has no impact on the σ_{vis} determination.

Table 4.17: Systematic uncertainties on σ_{vis} determination related to the $n_1 n_2$ measurement.

	proton–lead	proton–proton	lead–lead
DCCT scale	0.3%	0.2%	0.2%
Bunch-by-bunch fractions	0.1%	0.2%	0.1%
Ghost charge	0.2%	$< 0.1\%$	0.5%
Satellite charge	$< 0.1\%$	$< 0.1\%$	0.7%
Subtotal, bunch-charge product	**0.4%**	**0.3%**	**0.9%**

4.8.2 Beam-conditions

This section describes the sources of systematic uncertainties related to beam conditions during the the vdM scans. These are described separately in the following sections, and the result is summarized in table 4.19.

Orbit drift

The systematic uncertainty on the orbit drift correction to the proton–proton data is estimated as the difference of the scan-averaged σ_{vis} corrected results based on the arc-BPM and DOROS data. This corresponds to a systematic uncertainty of $\pm 0.1\%$ on σ_{vis}. Studies of proton–lead [3] and lead–lead [6] data estimated the corresponding systematic uncertainties as $\pm 0.2\%$.

Beam jitter

The beam separation at each step of the vdM scans is affected by random deviations of the beam positions. As described in Ref. [49], the magnitude of this "jitter" can be estimated with orbit drift data and is typically less than 1 μm (RMS). The resulting systematic uncertainty on σ_{vis} is obtained with pseudo-experiments by varying the nominal beam displacements at each scan point by a random jitter.

The resulting systematic uncertainty on σ_{vis} has been estimated as $\pm 0.3\%$ for the proton–lead data [3]. No dedicated studies have been performed for the proton–proton or lead–lead data, so a conservative estimate of $\pm 0.6\%$ is assigned for these datasets.

Emittance growth

The Σ_{X}, Σ_{Y} and horizontal and vertical $\mu_{\text{sp}}^{\text{peak}}$ are measured with data taken at different times. The time lapse between the peaks of the horizontal and vertical scans is typically around 30 minutes. If the emittance growth in the horizontal and vertical scans is exactly the same, the relative decrease in the average $\mu_{\text{sp}}^{\text{peak}}$ would be the same as the relative increase of Σ_{X} Σ_{Y}; thus yielding a null effect on σ_{vis}.

Following Ref. [49], the potential non-cancellation of emittance growth effects are estimated by interpolating linearly the values of the bunch-averaged Σ_{X}, Σ_{Y} and $\mu_{\text{sp}}^{\text{peak}}$ to a reference time and recalculating the σ_{vis} with these interpolated values. The rate of change is obtained from the changes of these quantities from one scan to the next.

The reference time used in the interpolation is arbitrary and the estimated effect should be independent of the time chosen. Three references times are used: at the peak of the horizontal scan, at the peak of the vertical scan, and the average of these times. When the time of the peak of the horizontal scan is used, neither the horizontal $\mu_{\text{sp}}^{\text{peak}}$ nor the Σ_{X} need to be interpolated. Similarly, when the time of the peak of the vertical scan is used, neither the vertical $\mu_{\text{sp}}^{\text{peak}}$ nor the Σ_{Y} need to be interpolated. When the average time is used all quantities are interpolated.

For the proton–lead scans, the largest difference between the interpolated and nominal σ_{vis} is found to be -0.3%. It is also found that the estimated effect is dependent on the reference time chosen. Thus, a systematic uncertainty of $\pm 0.3\%$ is assigned on σ_{vis}. Studies of the proton–proton [5] and lead–lead [6] data estimated the corresponding uncertainties as $\pm 0.1\%$ and $\pm 0.2\%$ respectively.

Scan-to-scan reproducibility

The scan-to-scan variation of the σ_{vis} after all corrections have been applied is taken as a systematic uncertainty on σ_{vis}. The maximum relative difference between the scan-averaged LucidOR σ_{vis} in any pair of scans is about 0.8% for the proton–lead scans, 0.4% for the proton–proton scans and 0.3% for the lead–lead scans. Similar differences are

found in the BcmHOR and BcmVOR results, indicating that the inconsistency among scans is not driven by an instrumental effect.

Bunch-to-bunch σ_{vis} consistency

The σ_{vis} measurements for each bunch crossing in a given scan are not always statistically compatible. The non-statistical component of this bunch variation, $\sigma_{\text{b2b,irr}}$, is defined as:

$$\sigma_{\text{b2b,irr}} = \sqrt{\sigma^2_{\text{b2b,total}} - \langle\sigma_{\text{stat}}\rangle^2}. \tag{4.19}$$

Here $\sigma^2_{\text{b2b,total}}$ is the variance over all bunch crossings and $\langle\sigma_{\text{stat}}\rangle$ is the bunch-averaged statistical uncertainty on σ_{vis}. The largest value of $\sigma_{\text{b2b,irr}}$ for any scan is taken as a systematic uncertainty on σ_{vis}; these are shown in table 4.18.

Table 4.18: Bunch-averaged statistical uncertainty on σ_{vis} and irreducible bunch variation.

	proton–lead	proton–proton	lead–lead
$\langle\sigma_{\text{stat}}\rangle$	0.3%	0.2%	1.0%
$\sigma_{\text{b2b,irr}}$	0.5%	0.6%	0.7%

Fit model

As an alternative to the default fit a numerical integration (trapezoidal rule) is performed to the scan curves to obtain Σ_{X} and Σ_{Y}. The maximum difference between the bunch-averaged σ_{vis} values obtained with the default fit method and the numerical integration is taken as a systematic uncertainty on σ_{vis}. The algorithm used in this estimation is LucidOR due to its larger statistical precision. The systematic uncertainty corresponds to $\pm\,0.5\%$ for the proton–lead scans; $\pm\,0.5\%$ for the proton–proton scans and $\pm\,0.8\%$ for the lead–lead scans. This also covers possible fit biases attributed to systematic patterns in the fit pulls presented in section 4.6.2.

Non-factorization effects

The analysis presented in section 4.6 assumes that the particle densities in each bunch can be factorized into independent horizontal and vertical components. If the factorization

assumption is violated, the Σ extracted in one plane is not independent of the beam separation in the other plane.

The beam-separation dependence of the position and shape of the three-dimensional luminous centroid ("beam-spot"), can be used to constrain possible biases to the luminosity calibration from such effects [49, 86].

In previous proton–proton vdM analyses, the possible effects of non-linear correlations of the x and y beam profiles have been successfully evaluated using a combination of beam-spot analyses and separate vdM analyses using non-factorizable fit functions. Typically, those two different methods have provided consistent results indicating good control over non-factorization effects [49].

The non-factorizable fit function "coupled GPn" is given by:

$$f(\Delta_x, \Delta_y) = e^{-\Delta_x^2/2\sigma_x^2 - \Delta_y^2/2\sigma_y^2} \times \left(1 + \sum_i^n \left[C_{ix}\left(\frac{\Delta_x}{\sigma_x}\right)^i + C_{iy}\left(\frac{\Delta_y}{\sigma_y}\right)^i\right]\right). \tag{4.20}$$

It differs from the product of uncoupled GPn functions because it does not have cross-terms involving both x and y. Then, for example, the $n = 2$ terms in equation 4.20, multiply the Gaussian with an elliptic profile. The odd powers allow for asymmetric modulation of the Gaussian profile.

The coupled-profile fits were performed using the procedures described in Ref. [87]. In particular, the coupled fits require a simultaneous fit to the horizontal and vertical vdM profiles. To evaluate the magnitude of non-factorization effects, the results of the coupled GPn were compared to the uncoupled GPn results used for the primary analysis described in section 4.6.2. The biases on the scan-averaged σ_{vis} are estimated to be -1.0% for the proton–lead data, -0.5% for the proton–proton data and $+0.2\%$ for the lead–lead data.

Estimates based on the beam-spot method suggest biases of $+(1\text{–}2)\%$ depending on the scan for the proton–lead data, and $+0.5\%$ for the proton–proton data. Given that the beam-spot and coupled-fit methods differ, no correction is applied to the nominal results and a conservative uncertainty of $\pm 2.0\%$ is assigned on σ_{vis} for the proton–lead data and $\pm 1.0\%$ for the proton–proton data. No beam-spot analysis has been finalized for the lead–lead data, so a conservative systematic uncertainty of $\pm 1.0\%$ is assigned to the lead–lead results.

Beam-beam effects

The uncertainty on the beam–beam deflection calculation is estimated by varying the β^* values that are used as input by an uncertainty that has been estimated to be $\pm 20\%$ [50]. The corresponding systematic uncertainties on σ_{vis} are $\pm 0.1\%$ for the proton–lead data, $\pm 0.3\%$ for the proton–proton data. Given that the beam–beam deflection is negligible for the lead–lead data, its uncertainty does not impact the σ_{vis} determination.

Table 4.19: Systematic uncertainties on σ_{vis} determination from beam-conditions.

	proton–lead	proton–proton	lead–lead
Orbit drift	0.2%	0.1%	0.2%
Beam-position jitter	0.3%	0.6%	0.6%
Emittance growth	0.3%	0.1%	0.2%
Scan-to-scan reproducibility	0.8%	0.4%	0.3%
Bunch-to-bunch σ_{vis} consistency	0.5%	0.6%	0.7%
Fit model	0.5%	0.5%	0.8%
Non-factorization effects	2.0%	1.0%	1.0%
Beam-beam effects	0.1%	0.3%	$< 0.1\%$
Subtotal, beam conditions	**2.4%**	**1.5%**	**1.6%**

4.8.3 Instrumental effects

This section describes systematic uncertainties related to instrumental effects. These are described separately in the following sections, and the result is summarized in table 4.20.

Reference specific luminosity

The so called specific luminosity, $\mathcal{L}_{\text{spec}} \equiv (\Sigma_{\text{X}}\Sigma_{\text{Y}})^{-1}$, only depends on the beam properties and thus should be independent of the luminometer used to measure it. The difference between the scan-averaged $\mathcal{L}_{\text{spec}}$ results obtained with the LUCID and BCM algorithms is used to estimate a systematic uncertainty on σ_{vis}. This corresponds to $\pm 0.2\%$ for the proton–lead and proton–proton scans and $\pm 0.6\%$ for the lead–lead scans.

Background subtraction

The background subtraction impacts the LucidOR σ_{vis} determination at a few permil level. A conservative systematic uncertainty of half the correction is assigned on the σ_{vis} determination. This corresponds to an uncertainty of $\pm 0.1\%$ for the proton–lead data, $< 0.1\%$ for the proton–proton data and $\pm 0.2\%$ for the lead–lead data.

Length scale calibration

The major source of systematic uncertainty on the length scale calibration arises from orbit drift during the length-scale scans. It has been estimated as $\pm 0.4\%$ for the proton–lead data [3] and $\pm 0.1\%$ for the proton–proton data [5]. Given that no dedicated length-scale calibration analysis has been finalized for the lead–lead data, a conservative uncertainty of $\pm 5.0\%$ is assigned. This is about twice the largest correction applied in any previous analysis.

Inner detector length scale

The length scale correction applied to the σ_{vis} measurements relies on the inner detector accuracy. The inner detector length scale uncertainty has been estimated to be $\pm 0.3\%$ for Run 1 data [49]. No dedicated studies have been performed with Run 2 data. A conservative estimate of $\pm 0.6\%$ is assigned for Run 2 data. This is motivated by the changes in the detector hardware made between Run 1 and Run 2.

Table 4.20: Systematic uncertainties on σ_{vis} determination from instrumental effects.

	proton–lead	proton–proton	lead–lead
Reference specific luminosity	0.2%	0.2%	0.6%
Background subtraction	0.1%	$< 0.1\%$	0.2%
Length scale calibration	0.4%	0.1%	5.0%
Inner detector length scale	0.3%	0.6%	0.6%
Subtotal, instrumental effects	**0.5%**	**0.7%**	**5.1%**

4.9 Luminosity monitoring

This section describes luminosity measurements performed during the 2015 lead–lead run. Similar studies for the 2013 proton–lead and 2015 proton–proton runs are documented in Refs. [3, 5].

The σ_{vis} values obtained with the vdM scans are used to calibrate the measurements of LucidOR, LucidBiOR, BcmHOR and BcmVOR. Figure 4.36 shows the calibrated luminosity as a function of time during a representative lead–lead run. The luminosity starts at about 2.6×10^{27} cm^{-2}s^{-1} and decreases by a factor of three over the duration of the run (about 6 hours). This decrease is mainly due to intra-beam scattering that reduces the beam intensity.

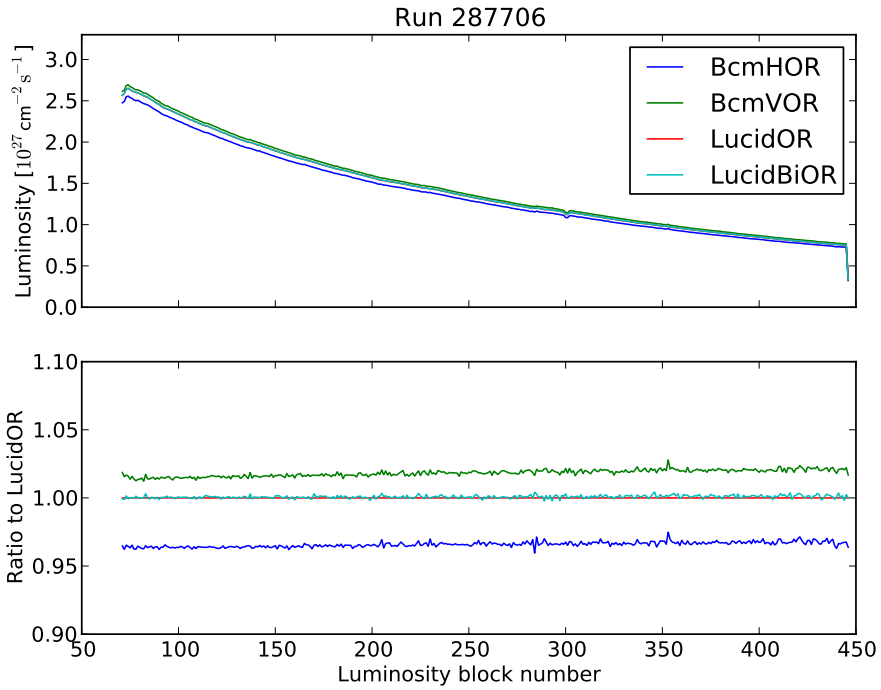

Figure 4.36: Top panel: luminosity as a function of luminosity block number (i.e. time) during a representative lead–lead run. Bottom panel: ratio to LucidOR measurements. Each luminosity block lasts about 60 seconds.

The LucidOR and LucidBiOR results are in excellent agreement. This indicates that the background subtraction of the bismuth signal in the LucidBiOR algorithm does not bias the measurement. The results of BcmHOR and BcmVOR differ by up to $\pm 4\%$ relative to LucidOR. The fact that the results differ by a constant factor suggests that the relative response of the detectors is similar.

Figure 4.37 shows the measured integrated luminosity in each of the 33 lead–lead runs in 2015. The different values obtained among runs is driven by differences in the beam emittance and total duration of the run.

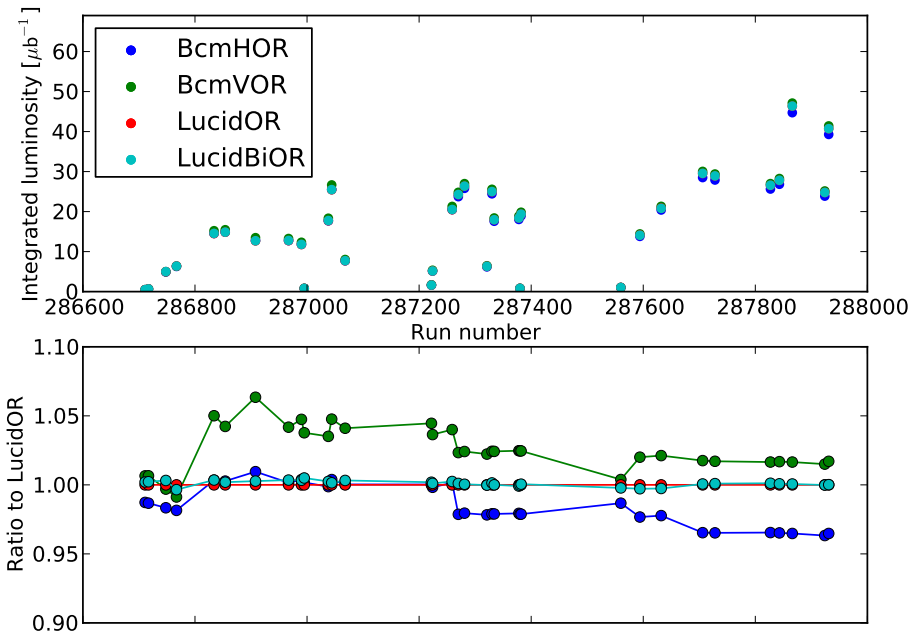

Figure 4.37: Top panel: integrated luminosity per run during the 2015 lead–lead runs. Bottom panel: ratio to the LucidOR measurements. The lines join the data points to guide the eye.

The LucidOR and LucidBiOR are in excellent agreement in all runs. The consistency between the BcmHOR and LucidOR algorithms is better than 2% for the first half of the runs and about 2.5–4% level in the latter half. The agreement between BcmVOR and LucidOR is better than 1% in the first two-runs, but then the BcmVOR efficiency increases relative to LucidOR by about 4–6% during the first half of the runs. A stepwise decrease of BcmVOR efficiency relative to LucidOR at the middle of the run period is observed at about 2.5%. This is consistent to what it is observed for BcmHOR.

The long-term stability of the response of the different algorithms can be studied with measurements of cumulative integrated luminosity. This may reveal changes in the response of the luminometers with total rather than instantaneous luminosity. Figure 4.38 shows the cumulative luminosity as a function of run.

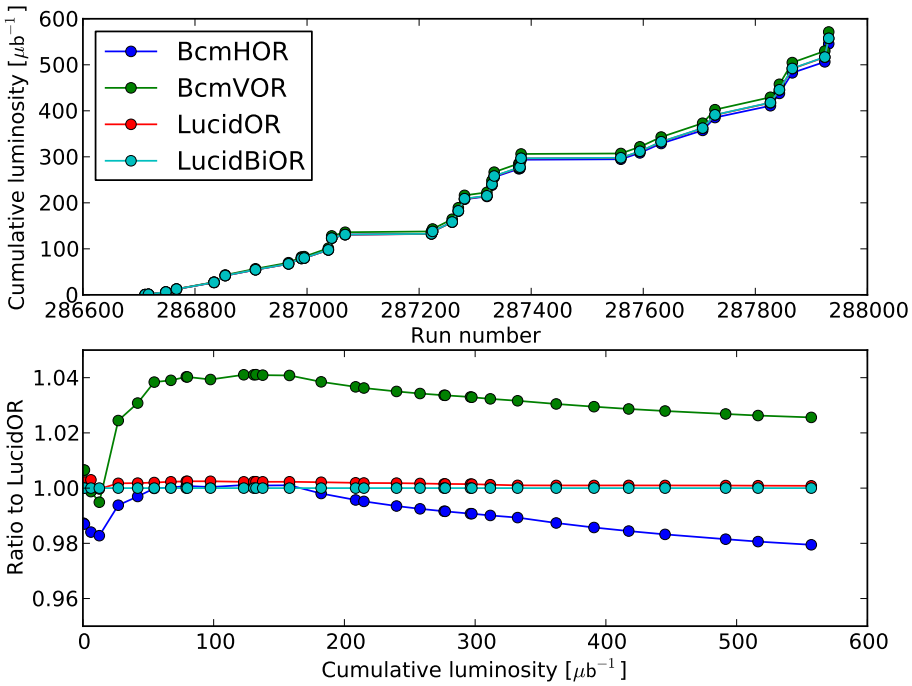

Figure 4.38: Top panel: cumulative luminosity as a function of run number during the 2015 lead–lead runs. Bottom panel: ratio to the results of LucidOR shown as a function of cumulative luminosity (by LucidOR). The lines join the data points to guide the eye.

The ratio plot shows that the response of the BCM algorithms increases relative to LucidOR by 2–4% with the first $\approx 50\ \mu b^{-1}$. This effect is consistent with the so called "charge-pumping" that increases the BCM efficiency due to radiation damage[14]. The first observation of the charge-pumping effect was in the 2013 proton–lead data [3]. Then the response of the BCM algorithms stays constant up to $\approx 150\ \mu b^{-1}$ but then drifts by up to -2% relative to LucidOR.

The total integrated luminosity measured by BcmVOR (BcmHOR) is 3% higher (2% lower) than LucidOR. The LucidOR measurement is taken as the nominal result and a systematic uncertainty of $\pm 3\%$ is assigned to cover the discrepancy with the BCM results.

[14]Caused by saturation of the charge-trapping mechanisms in the diamond lattice.

4.10 Summary and conclusions

The luminosity scale for the 2013 proton–lead, and the 2015 proton–proton and lead–lead runs has been calibrated with the vdM technique. The luminosity measurement during physics running was performed with the LUCID and BCM detectors. Table 4.21 summarizes all systematic uncertainties on the total integrated luminosity for each dataset. The accuracy achieved in all cases meets the goals discussed in section 4.5.

Table 4.21: Summary of the systematic uncertainties on integrated luminosity. The total uncertainty considers more significant digits than shown in the table. The results in parenthesis apply for the low-luminosity proton–proton data.

	proton–lead	proton–proton	lead–lead
Subtotal, bunch charge product	0.4%	0.3%	0.9%
Subtotal, beam conditions	2.4%	1.5%	1.6%
Subtotal, instrumental effects	0.5%	0.7%	5.1%
Total syst. uncertainty on σ_{vis}	2.5%	1.7%	5.4%
Luminosity monitoring	1.0%	1.3 (1.0)%	3.0%
Total, luminosity uncertainty	**2.7%**	**2.1 (1.9)%**	**6.1%**

The proton–proton calibration is one of the most accurate ATLAS results to date, and comes close to the most accurate result at the LHC (by LHCb, 1.2% [68]). This result has been exploited in practically all ATLAS publications that use 2015 proton–proton data; either directly in the case of cross-section measurements, or indirectly via background normalizations or pile-up corrections. The accuracy achieved is most relevant for the measurements of electroweak bosons [56] and the inelastic proton–proton cross section [7]. The latter is described in the next chapter.

The proton–lead and lead–lead calibrations are the first of their kind in ATLAS. The accuracy of the proton–lead calibration is similar to the best of its kind (by LHCb, 2.4% [68]) and has been used in some of the first measurements of quarkonia [88] and electroweak bosons [72] in proton–nucleus collisions at the TeV scale. The preliminary lead–lead calibration is one of the most precise of its kind (after ALICE, 5.8% [89]) and has been used in the first ATLAS measurement of photon–photon scattering [73], that opens up a rich physics program based on the study of so called "ultra-peripheral" collisions in ATLAS.

Chapter 5

Measurement of the inelastic proton–proton cross-section at 13 TeV

5.1 Introduction

Cross-sections for hadron–hadron scattering are amongst the most basic observables in particle physics experiments and yet they provide a direct probe of the dynamics of the strong interactions, which is described by the theory of quantum chromo dynamics (QCD). Inclusive cross-sections such as the total elastic and inelastic proton–proton (or proton–antiproton) cross-sections probe the low-scale dynamics of the theory that is connected to the confinement of quarks and gluons.

Because the strength of the interactions is larger at low energy scales these cross-sections cannot be predicted from QCD via perturbative calculations. Nevertheless, phenomenological approaches have been able to describe the data with reasonable success even before the development of QCD. During the 1950's, data from several fixed-target experiments showed that hadronic cross-sections (pp, $p\bar{p}$, πp, Kp, etc.) presented only a mild energy dependence at the highest energies available at the time. In this context, the mainstream theory predicted that all hadronic cross-sections would become constant at sufficiently large energies [90].

This prediction motivated the experimental community to explore the asymptotic behaviour; however, given that the center-of-mass energy achievable in fixed-target experiments rises only as the square root of the beam energy, progress was limited until the advent of the collider era.

The intersecting storage rings (ISR) at CERN was the first hadron collider and began operation in 1971. It delivered proton–proton collisions with a center-of-mass energy up to 62 GeV, almost five times higher than had previously been explored. One of the high priorities of the ISR experiments was the search for the predicted asymptotic behaviour of the proton–proton cross-section. As explained in chapter 4, this motivated an extensive research program that led to the development of the van der Meer technique.

The experiments measuring total cross-sections at the ISR were of two types: one used hodoscopes to trigger on inelastic interactions and the other used tracking detectors located near the beam pipe to measure small-angle elastic scattering.

The hodoscope based method is conceptually simple: it aims to measure the total rate of inelastic interactions, which together with an independent measurement of the ISR luminosity enabled the cross-section measurement. The main experimental challenge was the luminosity calibration with the newly invented vdM method.

The second approach aimed to measure elastic scattering down to very low angles (small momentum transfer, t) to estimate the differential elastic cross-section at $t \to 0$. This can be related to the total proton–proton cross-section via the optical theorem. The main experimental challenge was the development of tracking detectors located very close to the beam pipe; these are today known as "Roman pots".

Arguably one of the most important results from ISR experiments was the discovery that the total proton–proton cross-section (elastic plus inelastic) increases with center-of-mass energy; data showed an increase of 10% from 30 to 62 GeV [91, 92]. This result was not what most of the physics community expected, although a few had predicted it.

In 1952, Heisenberg predicted the rise of the total hadronic cross-sections and even postulated the energy dependence was a log-square, $\ln^2 s$, based on the kinematics of a model of meson production [93]. In 1961, Froissart and Martin postulated that $\ln^2 s$ was an upper limit of the s-dependence of the total proton–proton cross-section, based on a general argument that depends on unitarity (conservation of probability) [94]. In 1970, Cheng and Wu predicted a monotonic rise of the total proton–proton cross-section with energy based on a massive version of quantum electrodynamics [95].

Moreover, studies of data from cosmic-ray experiments suggested that hadronic cross-sections were much higher than at laboratory energies [96]. This result influenced the development of the theory of strong interactions, and gave rise to the modern version of "Pomeron physics" [90].

Experiments that succeeded the ISR at the CERN SPS and Fermilab Tevatron confirmed the rise of the proton–proton and proton–antiproton cross-sections up to the TeV scale [97]. In 2011, the LHC delivered proton collisions at 7 TeV and several experiments measured the elastic, inelastic and total proton–proton cross-sections, as described in the next section.

Figure 5.1 shows the collection of all measurements of the proton–proton and proton–antiproton total cross-section to date. This includes measurements from a large fraction of all accelerator-based experiments since the beginning of particle physics, as well as measurements of extensive cosmic-ray air showers at even higher energies. The data remain consistent with Heisenberg's log-square dependence[1].

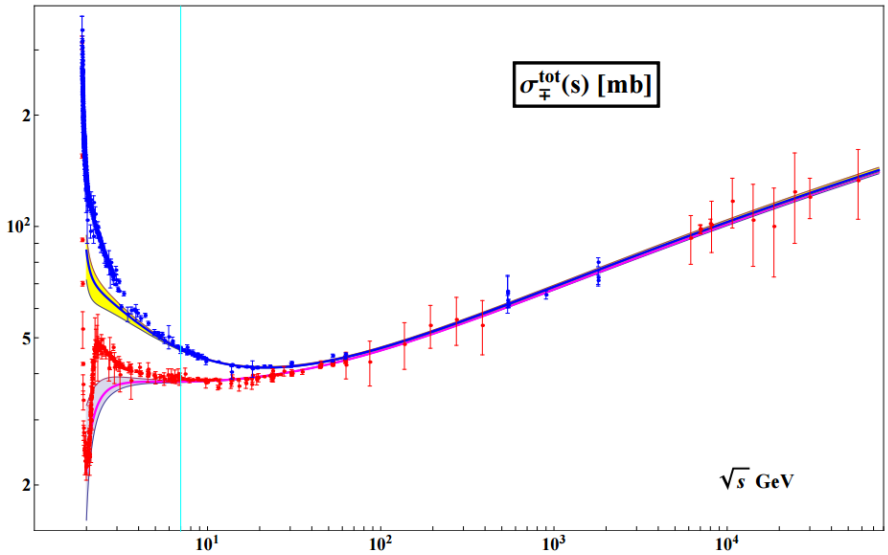

Figure 5.1: Total proton–proton (in red) and proton–antiproton (in blue) cross-section as a function of center-of-mass energy measured by several fixed-target, collider, and cosmic-ray experiments. The lines represent a fit to the data with a function quadratic in $\ln(s/s_0)$, where s_0 is a constant represented by the vertical light-blue line. Source: Ref. [98] and references therein.

Measurements at the LHC

The measurements of elastic and inelastic cross-sections at the LHC, summarized in table 5.1, have relied on techniques similar to those developed for the ISR experiments. The ATLAS, ALICE, CMS, and LHCb experiments have measured the inelastic cross-section with detectors that play a similar role to the ISR hodoscopes (scintillators in

[1] Other parametrizations such as power law are also consistent with the data.

ATLAS [99] and ALICE [100], a tracker in LHCb [101], and calorimeters in CMS [102]). Dedicated experiments known as TOTEM and ALFA are based on Roman pot detectors. These experiments measure elastic scattering and use the optical theorem to extract the total proton–proton cross-section [103, 104].

Both methods rely on a luminosity calibration[2] obtained primarily using the vdM method. The uncertainty on the luminosity calibration varies between 1.2% and 3.5% and is the dominant systematic uncertainty for the Roman pot experiments.

The CMS experiment also used an alternative method that exploits the fact that the number of proton–proton interactions in each bunch crossing follows a Poisson distribution with mean μ; therefore, the inelastic cross-section can be extracted using a Poisson fit to the distribution of reconstructed vertices [102].

Table 5.1: Summary of measurements of elastic, inelastic and total proton–proton cross-section at the LHC. These are described in more detail in the main text.

Experiment	Energy [TeV]	Experimental technique	Measure	Requires $\mathcal{L}$
ATLAS [99]	7	Hodoscopes (scintillators)	σ_{inel}	Yes
ALFA [103]	7	Roman pots	σ_{el}, σ_{total}	Yes
ALICE [100]	2.76, 7	Hodoscopes (scintillators)	σ_{inel}	Yes
CMS [102]	7	Hodoscopes (calorimeters)	σ_{inel}	Yes
CMS [102]	7	Pile-up counting	σ_{inel}	Yes
LHCb [101]	7	Hodoscopes (tracker)	σ_{inel}	Yes
TOTEM [104]	7, 8	Roman pots	σ_{el}, σ_{total}	Yes
TOTEM [105, 106]	7, 8	Roman pots and Hodoscopes (scintillators)	σ_{el}, σ_{inel}, σ_{total}	No

Motivation for a measurement at 13 TeV

In 2015, the LHC delivered proton collisions at a center-of-mass energy of 13 TeV. Measurements of elastic and inelastic cross-sections at this new energy will probe non-perturbative QCD, test the center-of-mass energy dependence and the Froissart bound. Moreover, the combined data of all LHC experiments will test whether the asymptotic

[2] TOTEM has also measured independently the inelastic and elastic cross-sections which enables a luminosity-independent measurement [105, 106].

"black-disk" limit of hadronic interactions has been reached or not; this limit predicts a ratio of the inelastic to total cross-section of $1/2$ as $s \to \infty$. Some studies claimed evidence for this behaviour with 7 TeV data from the LHC and 57 TeV data from the Auger experiment [107, 108], but others have contested that claim [109, 110].

From a detector performance perspective, the measurement of the inelastic cross-section enables a measurement of the pile-up parameter μ (see equation 4.11). This is an input for simulations that aim to accurately describe the detector performance. Moreover, it is an input for the corrections to the calorimeter response that ensure precise measurements of jet transverse momenta and missing-transverse energy [20, 111].

The impact of a measurement of the inelastic cross-section also permeates to astroparticle physics. The LHC is the first accelerator to probe the energy regime above which the cosmic ray spectrum presents a discontinuity in the slope that remains to be explained. This is known as the "knee". One of the primary ways to address the origin of the knee is to measure the composition of cosmic rays.

Detectors in satellites or balloons can identify cosmic rays before they hit the atmosphere. Their limitation is the low flux for energies above $\approx 10^{15}$ eV. Large-area detectors at ground level provide an alternative as these measure the shower of particles produced after cosmic rays hit the atmosphere. Protons and heavier nuclei produce different types of showers; this is exploited to infer the cosmic ray type. But this method suffers from a drawback: it relies on empirical models of hadronic interactions. One of the key inputs for these models is the inelastic proton–proton cross-section [112] as it influences the mean-free path of cosmic rays in the atmosphere[3].

5.2 Types of inelastic interactions

Inelastic proton–proton interactions can be classified according to their products. In diffractive scattering there is only a small energy transfer[4] between the protons, but one or both disintegrates. The prominent signature of diffractive scattering is the presence of voids in the angular distribution of the produced particles, known as "rapidity gaps".

[3] The inelastic proton–proton cross-section is extrapolated to the inelastic cross-sections of heavier nuclei and air by using models of nuclear interactions [113].

[4] Theoretically this is described as pomeron exchange.

Single-diffractive scattering occurs when only one of the protons disintegrates, and double-diffractive scattering when both do so. Diffractive interactions are about 25% of the total inelastic cross-section at LHC energies. The remainder corresponds to non-diffractive scattering characterized by a large energy transfer and high particle multiplicities. Figure 5.2 shows an illustration of single-diffractive, double-diffractive, and non-diffractive interactions.

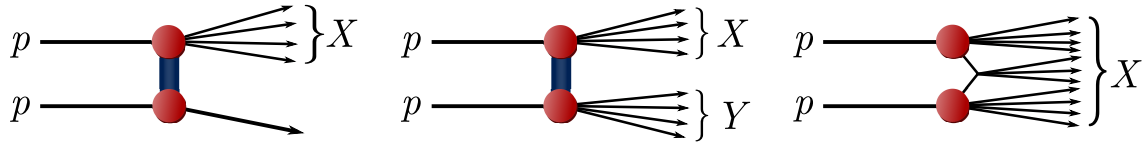

Figure 5.2: Sketches of proton–proton inelastic scattering types: single-diffractive scattering (left), double-diffractive scattering (middle) and non-diffractive scattering (right). The blue line represents pomeron exchange. Source: Ref. [7].

In diffractive interactions, the invariant mass of each dissociated system is related to the maximum pseudorapidity of the particles that compose them. That is, experimentally, the selection of events based on a detector with limited acceptance leads to a lower bound on the measurable invariant masses. This is discussed in more detail in section 5.8.

5.3 Overview of the analysis

The work presented below uses 13 TeV data and follows a similar methodology as the ATLAS analysis of 7 TeV data [99]. The analysis is based on highly-efficient plastic scintillators to trigger and select events containing inelastic interactions. The majority of proton–proton inelastic scattering events are non-diffractive and yield high particle multiplicities that are detected with optimal efficiency. However, a fraction of diffractive scattering events produce particles only at forward angles that are not covered by the scintillators (inevitable because of the beam-pipe) and thus are not detected.

The measurement is performed with data collected in a low-luminosity run to minimize the impact of pile-up. The method consists of counting the selected events and dividing by the luminosity. A number of small corrections are applied to take into account trigger and selection efficiencies and backgrounds. The efficiency of the scintillators is large enough that uncertainties on its value do not impact the measurement significantly. Backgrounds from induced radiation and beam-gas interactions are estimated with a

data-driven method using dedicated triggers; these represent a sub-percent correction. The final result is obtained by applying a data-driven correction for undetected events that is about 15%.

5.4 Experimental set up

Minimum bias trigger scintillator

The minimum bias trigger scintillator (MBTS) is a two-arm hodoscope made of plastic scintillators. Its primary goal is, as its name suggests, to serve as a trigger for minimum-bias interactions. It is located on the front face of each end-cap calorimeter at $|z|= 3.6$ m. It covers $2.4^\circ < \theta < 14.4^\circ$ ($2.07 < |\eta| < 3.86$) and is divided into two regions: the outer ring covers $3.7^\circ < \theta < 14.4^\circ$ ($2.07 < |\eta| < 2.76$) and is divided azimuthally into 4 equal segments; the inner ring covers $2.4^\circ < \theta < 3.7^\circ$ ($2.76 < |\eta| < 3.86$) and is divided azimuthally into 8 equal segments. Figure 5.3 shows an illustration of the MBTS.

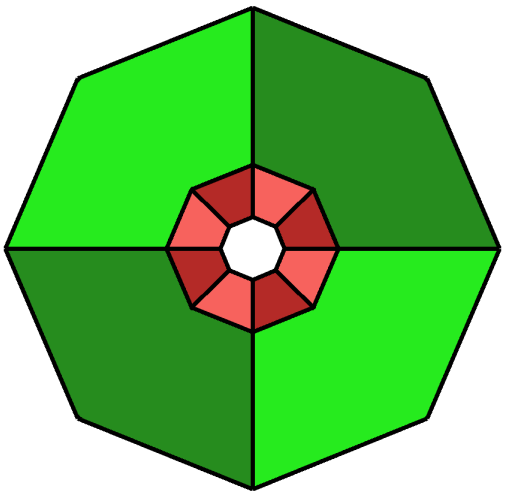

Figure 5.3: Illustration of the geometry of the minimum-bias trigger scintillator during Run-2. The inner modules cover $14.9 < r < 44.5$ cm and the outer modules $44.5 < r < 98.5$ cm. The colour illustrates the different modules. Source: Ref. [7].

The light from the scintillators is readout with wavelength-shifting fibres to photo-multipliers located about 15 meters away (to avoid the high magnetic fields that would deteriorate their performance). The MBTS provides single-particle detection capability and a time resolution of about 10 ns.

The MBTS is primarily used during low-luminosity running such as at the commissioning phase of each running period. At high luminosity, it saturates due to pile-up and suffers from severe radiation damage. For this reason, it was completely replaced after Run-1. The Run-2 geometry (described in the previous paragraphs) follows closely the Run-1 design but incorporated small changes in the geometric acceptance, segmentation, and routing of the wavelength-shifting fibres.

Trigger strategy

Table 5.2 summarizes the triggers used. The main trigger is named `MBTS_1` and required at least one MBTS module to have a signal above a threshold that ranged from 0.5 pC to 0.7 pC, depending on the module[5], in coincidence with a filled bunch crossing. No requirements were applied at the trigger software level.

Table 5.2: Summary of triggers used in this analysis.

Trigger name	Used to measure
`MBTS_1`	Signal events
`MBTS_1_UNPAIRED`	Background
`MBTS_1_EMPTY`	Background
`RD0_FILLED`	`MBTS_1` efficiency, background
`LHCF_1`	`MBTS_1` efficiency
`LUCID_1`	`MBTS_1` efficiency

Two other triggers similar to `MBTS_1` but in coincidence with "unpaired" or "empty" bunch crossings were also used. Unpaired bunch crossings refer to 25 ns slots when only one beam had a filled bunch; empty bunch crossings are those that had no filled bunch. The corresponding triggers are named `MBTS_1_UNPAIRED` and `MBTS_1_EMPTY`. The data collected with these triggers are used in the background estimation presented in section 5.9.

A random trigger in coincidence with filled bunch crossings, named `RD0_FILLED`, is used in the measurement of the `MBTS_1` efficiency (section 5.10) and background (sections 5.9). Additionally, triggers based on signals from either LUCID or the LHCf detector were used to measure the `MBTS_1` efficiency. These are called `LUCID_1` and

[5]This is slightly below the most probable value expected from a minimum-ionizing particle.

`LHCF_1`. The LUCID detector was described in chapter 4; LHCf is a tungsten-scintillator calorimeter located at $|z| = 140$ m.

Datasets

The datasets used were collected during a series of low-luminosity runs in July 2015. These were some of the first runs of the LHC Run-2. The peak luminosity was $1.5\text{–}3.5 \times 10^{28}$ cm^{-2}s^{-1}, and the pile-up was $\mu \approx (2.6\text{–}6.1) \times 10^{-3}$, depending on the run. Technically, the low per-bunch luminosity was achieved by using low beam intensities ($\approx 10^{10}$ protons), large β^* (19 m), and by offsetting the beams. The LHC bunch pattern in the ATLAS interaction region contained 37 paired bunches and 4 unpaired bunches.

The motivation for this special low-luminosity run is to suppress pile-up. This is essential because this analysis is by definition inclusive to all inelastic interactions and thus only minimal selection criteria are applied (described in section 5.6). This makes the analysis particularly sensitive to pile-up because it would complicate the event counting and lead to a degradation of the MBTS response.

The runs and their details are listed in table 5.3. The first dataset, with the lowest μ, is used for the main result. It contains about 8.7 million events recorded with `MBTS_1` trigger. Other runs with larger μ are used for cross-checks and efficiency measurements.

The data quality selection criteria applied required the MBTS to be fully operational, but did not use information from any other sub-detectors. From the dataset passing the data quality requirements, only a subset in which no prescales were applied to the `MBTS_1` is considered in the analysis. The dataset used corresponds to an integrated luminosity of 60.1 ± 1.1 μb^{-1}, or 4.3 million events.

Table 5.3: Summary of the runs used. Shown are the run number, the center-of-mass energy, the peak and integrated luminosities (before any selection), the average number of interactions per crossing, the number of recorded events, and the main use of each run.

Run	$\sqrt{s}$ [TeV]	Peak $\mathcal{L}$ [cm^{-2}s^{-1}]	$\int \mathcal{L}\, \mathrm{d}t$ [μb^{-1}]	μ	Events	Use
267358	13	1.5×10^{28}	100	2.6×10^{-3}	8,665,704	Main measurement
267359	13	3.5×10^{28}	290	6.1×10^{-3}	12,295,654	Cross-checks
267599	13	2.2×10^{30}	7222	0.50	105,611,891	Trigger efficiency

5.5 Simulation datasets

Several Monte Carlo (MC) simulation datasets are used to obtain the selection efficiency (section 5.8), compare with the measured MBTS hit distributions (section 5.13), and extrapolate the measured cross-section (section 5.13) to account for undetected events.

The MC generators used are Pythia8 [114, 115], Epos LHC [116], and QGSJet-II [117]. Pythia8 is the default generator used in ATLAS and other LHC experiments to model "minimum bias" interactions (i.e. inclusive inelastic scattering). Two different "tunes", i.e. sets of parameters and parton density functions, of Pythia8 are used: the A2 [118] and Monash [119] tunes. Epos LHC and QGSJet-II are extensively used in the modelling of cosmic-ray showers and differ from Pythia8 in their hadronization modelling. Epos LHC has been constrained with minimum-bias data from LHC experiments.

One of the main differences between these models lies in the modelling of diffractive interactions. In Pythia8 this is, by default, based on a model by Schuler and Sjöstrand (SS) [120, 121]. Alternative MC samples are generated with the Donnachie and Landshoff (DL) [122] model, and the Rockefeller (MBR) [123] model. The DL model has two main parameters, ε and α'. The values considered ($\varepsilon = 0.06,\ 0.085,\ 0.10$; $\alpha' = 0.25$) are motivated by previous diffractive cross-section measurements [124, 125]. Epos LHC and QGSJet-II use a "cut pomeron" [116] and "single-pole pomeron" [117] model respectively.

Table 5.4 summarizes the MC generators used, their tunes and models of diffraction. In all simulated samples, the detector response is modelled using Geant4 [126]. The MC simulated events are processed with the same reconstruction and analysis algorithms as are used for data.

Table 5.4: Summary of Monte Carlo generators used, specific tunes (if applicable) and model of diffractive interactions.

MC Generator name	Tunes	Model of diffractive interactions
Pythia8	A2 and Monash	Schuler and Sjöstrand [120, 121]
		Donnachie and Landshoff [122]
		Rockefeller [123]
Epos LHC	—	Cut pomeron [116]
QGSJet-II	—	Simple-pole pomeron [117]

5.6 Event selection

Hit threshold

Figure 5.4 shows the charge distributions in events selected by `MBTS_1_EMPTY` and in simulation. The data are well described by the sum of a narrow and a very wide Gaussian function. The noise model in the simulation was by default a single-Gaussian (shown as "MC Noise") but it was changed to a double-Gaussian to match the data. For the "MinBias MC" (Pythia8 default) the tail towards positive charges is much larger due to contributions from real particles. Based on this distribution the hit threshold is set to 0.15 pC. Increasing it further does not reduce the noise substantially since the second Gaussian is rather broad.

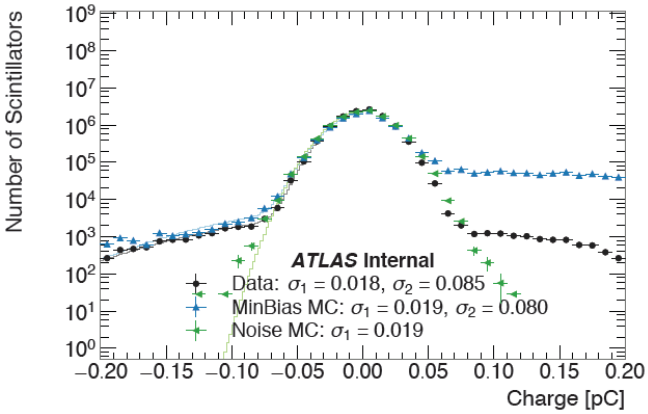

Figure 5.4: MBTS charge distribution in randomly-triggered events and MC simulations. The results of Gaussian fits to the negative tails are shown as lines. Source: Ref. [127].

Selection

The subset of the triggered events that have at least two MBTS hits is considered for the analysis, where a hit is a signal above 0.15 pC. The requirement on the number of hits is motivated by the background studies described in section 5.9. A total of 4,159,074 events pass this selection; these are called the "inclusive sample".

An additional "single-sided" selection is defined requiring at least two hits on one side of the detector and no hits on the other. A total of 442,192 events pass this selection. This sample is used for the background estimation (section 5.9) and to constrain the simulations used in the analysis (section 5.11).

5.7 Fiducial region definition

MC simulation shows that the selection requirement of at least two hits effectively limits the acceptance of diffractive scattering events to those in which at least one of the protons dissociated into a hadronic system with invariant mass $M_X > 13$ GeV or equivalently $\xi = M_X^2/s > 10^{-6}$. These values are discussed in more detail in section 5.8.

Thus, the measurement is performed in a restricted "fiducial region" corresponding to the subset of inelastic events[6] for which $\xi > 10^{-6}$. The measured fiducial inelastic proton–proton cross-section, $\sigma_{\text{inel}}^{\text{fid}}$, is given by:

$$\sigma_{\text{inel}}^{\text{fid}}\left(\xi > 10^{-6}\right) = \frac{N - N_{\text{BG}}}{\epsilon_{\text{trig}} \times \int \mathcal{L}} \times \frac{1 - f_{\xi<10^{-6}}}{\epsilon_{\text{sel}}}, \tag{5.1}$$

where N is the number of events in the inclusive sample; N_{BG} is the number of background events; ϵ_{trig} and ϵ_{sel} are the trigger and selection efficiencies; $1 - f_{\xi<10^{-6}}$ is a factor to account for migration of events with $\xi < 10^{-6}$ into the fiducial region due to detector resolution; and $\int \mathcal{L}$ is the integrated luminosity of the data sample.

The correction for migration of events and the event selection efficiency are combined in a single factor C_{MC} defined as:

$$C_{\text{MC}} \equiv \frac{1 - f_{\xi<10^{-6}}}{\epsilon_{\text{sel}}}. \tag{5.2}$$

The C_{MC} factor is close to unity by construction, i.e. by the judicious choice of the fiducial region. The values of C_{MC} for different simulations are presented in section 5.11.

5.8 Selection efficiency

The selection efficiency is determined from MC simulation as a function of ξ, which is defined in a manner that is applicable to both diffractive and non-diffractive events: in each event, the largest rapidity gap between any two final-state hadrons is used to define the boundary between two collections of hadrons (as illustrated in figure 5.5). These collections define the dissociation systems, named X and Y. The system with the largest mass is by definition X, and ξ is defined to be M_X^2/s.

[6] As explained in section 5.8, the ξ can be also defined in non-diffractive interactions.

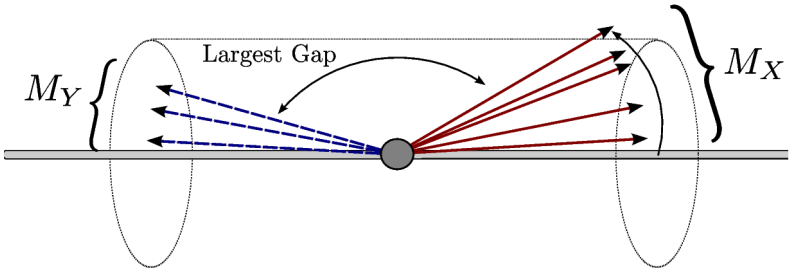

Figure 5.5: Illustration of M_X, and M_Y in a double-diffractive event. Source: Ref. [7].

Figure 5.6 shows the ξ distribution of inclusive inelastic interactions for different MC simulations and the fraction of events that pass the event selection, i.e. the selection efficiency. The distribution is dominated by the non-diffractive component which peaks at large values of ξ. The selection efficiency turns on at about $\xi = 10^{-7}$ and reaches a plateau of unity at about $\xi = 10^{-5}$ for most simulations.

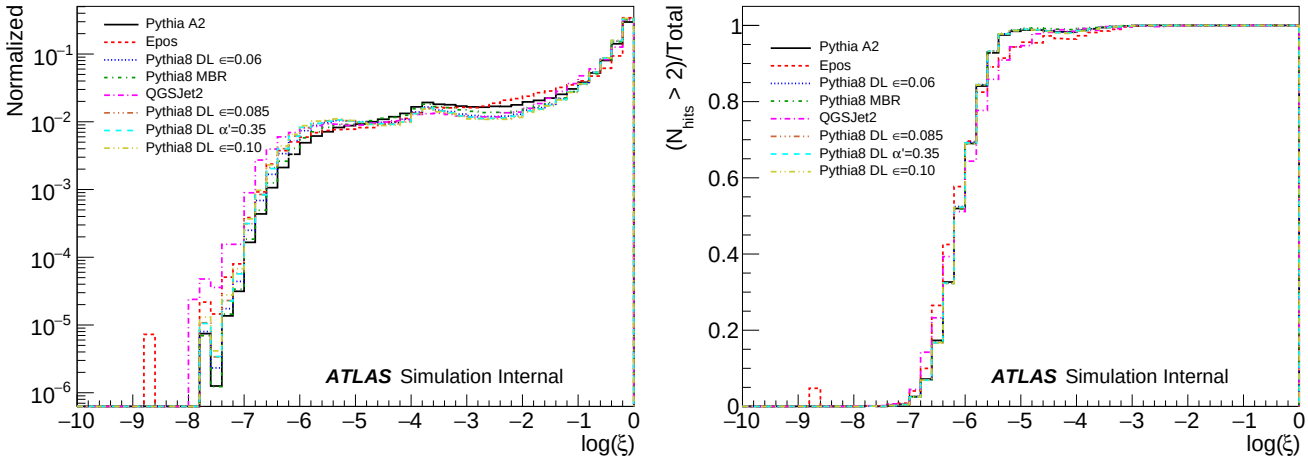

Figure 5.6: Left panel: distribution of $\log(\xi)$, without detector effects, for different MC simulations. Right panel: selection efficiency as a function of $\log(\xi)$. Source: Ref. [127].

The requirement of at least two hits corresponds approximately to requiring $\xi > 10^{-6}$ (where the selection efficiency exceeds $\approx 50\%$). This defines the fiducial volume of the measurement. The corresponding values of C_{MC} are shown in section 5.11, after data-driven constraints are applied to the simulations.

The estimation of the selection efficiency relies on an adequate simulation of the response of the MBTS. Data-driven estimates of the single-particle efficiency show that it is above 98.5% for all modules. The MC simulations were adjusted to match the efficiency measured in data, as described in Ref. [127].

5.9 Background estimation

The main sources of background are beam-gas interactions and afterglow. These were described in chapter 4. Backgrounds from cosmic rays and instrumental noise are negligible. The afterglow and beam-gas background is estimated with events selected by `MBTS_1_EMPTY` and `MBTS_1_UNPAIRED` triggers. In run 267358 there were 4 unpaired bunch crossings and 3023 empty bunch crossings. The `MBTS_1_UNPAIRED` trigger was unprescaled, whereas the `MBTS_1_EMPTY` trigger was prescaled by factors of 100, 200, and 2000 at different periods. Event weights are applied to correct for these prescale factors.

Events per bunch crossing

Figure 5.7 shows the number of events with two or more hits as a function of bunch crossing number. The bunch-to-bunch variation in the filled bunch crossings stems from the different per-bunch luminosity (emittance). The empty bunch crossings show larger activity and a decaying tail after a filled bunch crossing; this is consistent with afterglow.

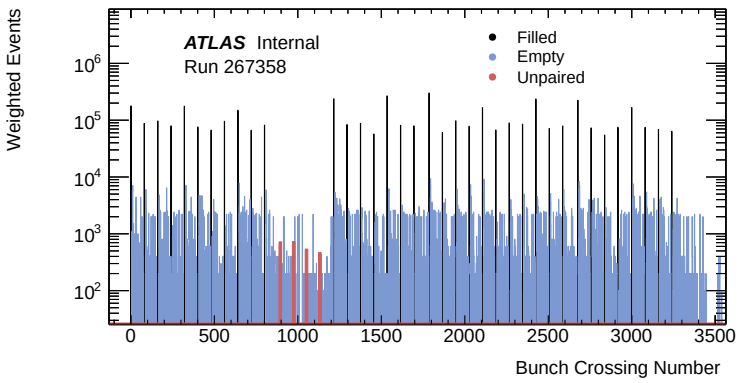

Figure 5.7: Number of events with two or more MBTS hits as a function of bunch crossing number.

Figure 5.8 shows a clearer comparison of the empty and paired bunch crossings using a linear vertical scale. This observed pattern of peaks in the empty bunch crossings is roughly similar to that of the filled bunch crossings. This suggest that afterglow is roughly proportional to the bunch crossing luminosity.

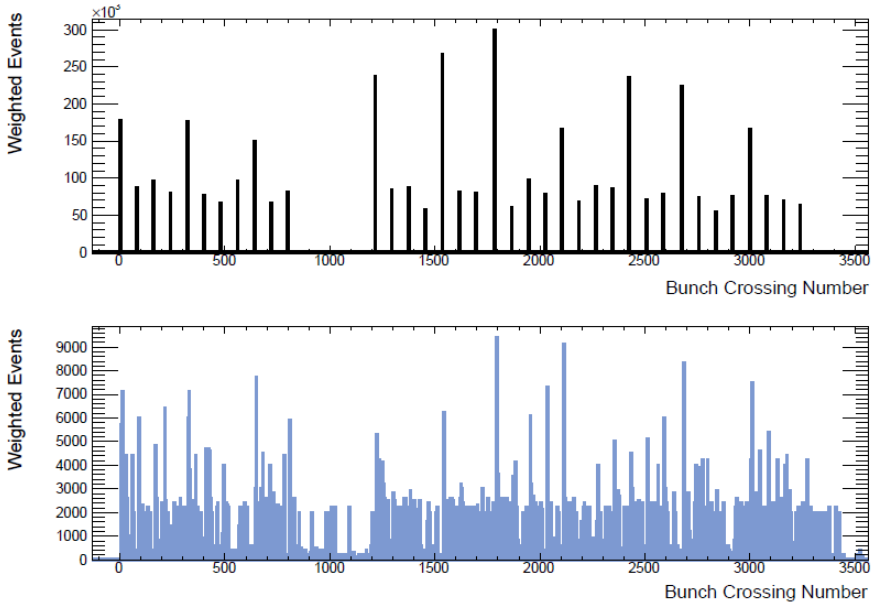

Figure 5.8: Comparison between number of events with two or more MBTS hits in filled bunch crossings (top panel) and empty bunch crossings (bottom panel).

Isolated empty bunches

Figure 5.9 is the same as figure 5.7 but for a restricted range of bunch crossings in a region containing no filled bunch crossings, the 4 unpaired bunch crossings, and 359 empty bunch crossings (some of them with no counts). The empty bunch crossings that contain more counts than the unpaired bunch crossings have a large statistical uncertainty. These correspond to a few events that were weighted by a factor of about 2000.

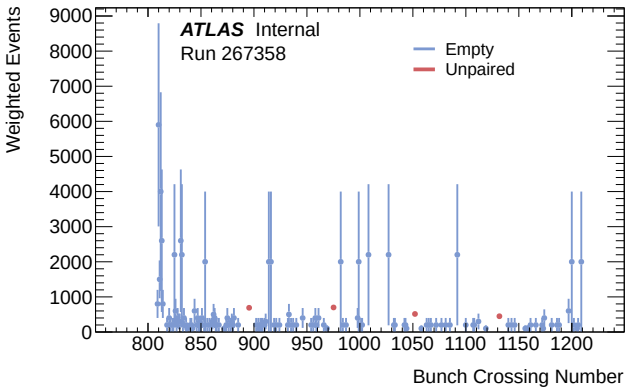

Figure 5.9: Number of events with two or more MBTS hits as a function of bunch crossing number for empty and unpaired crossings. The error bars represent statistical uncertainty only.

The empty bunch crossings in this range, called "isolated empty", are the least affected by the afterglow. This sample is used to estimate the afterglow in the unpaired bunch crossings. This is estimated as the weighted average over all the 359 bunch crossings in this sample. However, figure 5.9 also shows that more events are concentrated in the early bunch crossings, suggesting an afterglow tail. The limitations of this approach are considered in the assessment of systematic uncertainties.

Figure 5.10 shows the inclusive and single-sided hit distributions in unpaired bunch crossings and isolated-empty bunch crossings. The latter distributions are scaled by 4/359 to take into account the different number of bunch crossings in these categories.

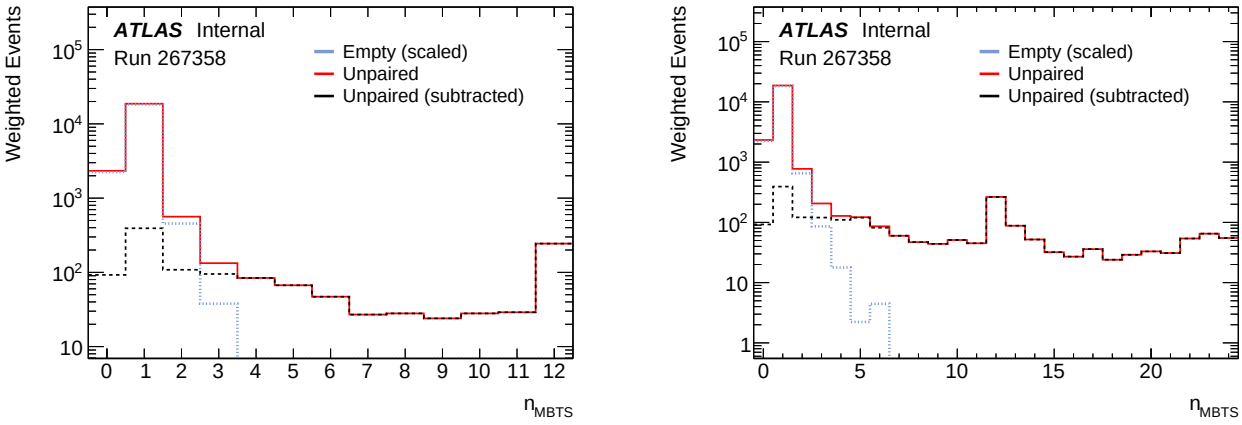

Figure 5.10: Single-sided (left panel) and inclusive (right panel) hit distributions in unpaired and empty bunch crossings.

The distribution from unpaired bunch crossings has a tail that reaches even the highest multiplicities. This is consistent with beam-gas interactions. On the other hand, the distribution from empty bunch crossings falls more steeply. This is consistent with afterglow. The afterglow estimate reproduces well the first few bins of the distribution from unpaired bunch crossings. The distribution that is left after the afterglow subtraction, which is dominated by beam-gas background, is much flatter in the low-multiplicity region.

Bunch currents in unpaired bunches

Beam-gas interactions are proportional to the bunch charge. Thus, a reliable estimation of this background with events selected by `MBTS_1_UNPAIRED` can only be made if the bunch charge in unpaired bunch crossings is similar to that in filled bunch crossings.

Figure 5.11 shows the average bunch charge[7], taken over the duration of the run, as a function of bunch crossing number. There are bunch-to-bunch variations on the bunch charge of about $\pm 25\%$ around a mean value of about 8×10^9 protons. The charge in the unpaired bunches is similar to the average bunch charge in the filled bunches.

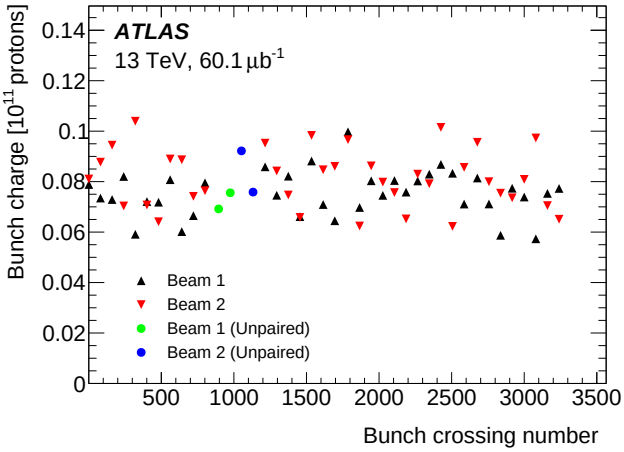

Figure 5.11: Average bunch charge over the duration of run 267358 as a function of bunch crossing number. The standard deviations of the measurements are smaller than the marker sizes.

Background estimate

The estimate of beam-gas background per bunch is taken as the number of events collected with `MBTS_1_UNPAIRED` minus the estimate of afterglow in isolated bunches, averaged over the 4 unpaired bunch crossings sampled:

$$N_{\text{perbunch}}^{\text{beam-gas}} = \frac{N_{\text{MBTS\_1\_UNPAIRED}} - N_{\text{ISO}}^{\text{afterglow}}}{4}. \tag{5.3}$$

The estimate of afterglow and noise per-bunch crossing is taken as the number of weighted events collected with `MBTS_1_EMPTY`, averaged over the 3023 empty bunch crossings sampled:

$$N_{\text{perBC}}^{\text{afterglow}} = \frac{N_{\text{MBTS\_1\_EMPTY}}}{3023}. \tag{5.4}$$

[7]As explained in chapter 4, the bunch charge is measured with the LHC FBCT and ATLAS BPTX. No significant differences between BPTX and FBCT measurements are observed.

Thus, the estimated background from beam-gas and afterglow for the 37 bunch crossings sampled by `MBTS_1` is given by:

$$N_{\mathrm{BKG}} = 37 \times \left(2 \times N_{\mathrm{perbunch}}^{\mathrm{beam-gas}} + N_{\mathrm{perBC}}^{\mathrm{afterglow}}\right). \tag{5.5}$$

The factor of 37 accounts for the number of filled bunch crossings, and the factor of 2 accounts for the fact that in a collision both bunches can produce beam-gas interactions.

Figure 5.12 shows the inclusive hit distribution and the corresponding background estimate, as well as the relative background fraction. The background estimate describes the increase of events in the low-multiplicity region well, suggesting a large background contamination in this region.

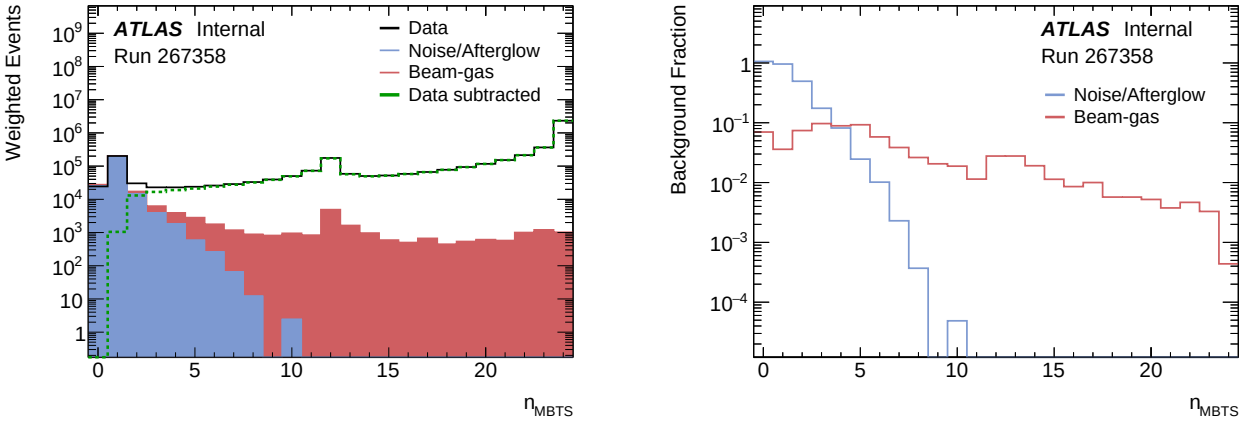

Figure 5.12: Left panel: MBTS hit distribution and estimated background in the inclusive sample. Right panel: background fraction as a function of hit multiplicity.

The afterglow estimate is about 100% for events with zero or one hit[8], about 40% for events with two hits, and decreases steeply for events with higher hit multiplicity. The beam-gas estimate is about 10% for events with less than five hits and slowly decreases with multiplicity. The peak of the beam-gas distribution at 12 hits is due to asymmetric events produced within the detector volume. The corresponding peak in the data is due to diffractive interactions, as will be shown in section 5.13.

Figure 5.13 shows the MBTS hit distribution measured in single-sided events and the corresponding background estimate. The hit distribution and background estimate show qualitatively similar features to those presented in figure 5.12.

[8]The background fraction for 0 hits is larger than 100%, which suggest a slight overestimation.

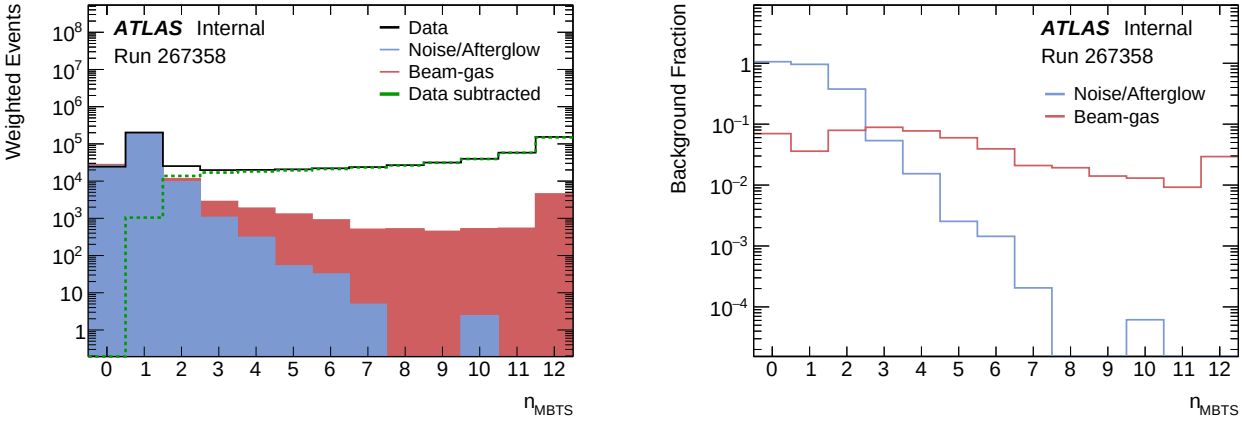

Figure 5.13: Left panel: MBTS hit distribution and estimated background in the single-sided sample. Right panel: background fraction as a function of hit multiplicity.

An event is considered to be a part of the signal if it contains at least two hits. This is a compromise between the impact of background contamination and the selection efficiency. So, the total background estimate is obtained by integrating the background distributions in the bins with at least two hits.

The estimated background in the signal region is: 0.7% from beam-gas interactions and 0.5% from afterglow and noise. In the single-sided sample these estimates increase to 3.3% and 2.5% respectively. This is expected as most beam-gas interactions pass the single-sided selection due to the asymmetry of the collisions but only about a tenth of the signal events do. The uncertainty on these estimates is described in section 5.12.

The background estimate is tested with data from run 267359. The luminosity in this run was about 2.5 times higher than in run 267358. The difference in μ is driven by different beam separations (the bunch-currents were similar). The background estimate for this run is 0.3% from beam-gas and 0.5% afterglow for the inclusive sample and 1.2% and 2.5% for the single-sided sample.

The estimated background is consistent with expectations: the rate of beam-gas background in both runs is similar (because it is driven by the bunch currents), but the rate of signal increases with μ. Thus, the relative background due to beam-gas interactions should be about 2.5 times lower in run 267359. On the other hand, the background from afterglow is expected to be proportional to μ. Thus the relative fraction should be the same in both runs.

In-time afterglow

The background estimate considered in the previous sections does not consider in-time afterglow, i.e. the afterglow due to the collision that triggered the event. This would cause late hits. Figure 5.14 shows the MBTS time distribution for events with different hit multiplicity in data and simulation. In all cases, the background estimate can account for most of the negative tail of the distributions, but does not account for the pronounced positive tail that is well described by the MC simulation.

Given that the MC simulation provides a reasonable[9] (ad hoc) model of the in-time afterglow effect, no explicit subtraction is needed as this effect is incorporated in the selection efficiency. Moreover, any potential bin migration in the hit distributions due to this effect is also simulated.

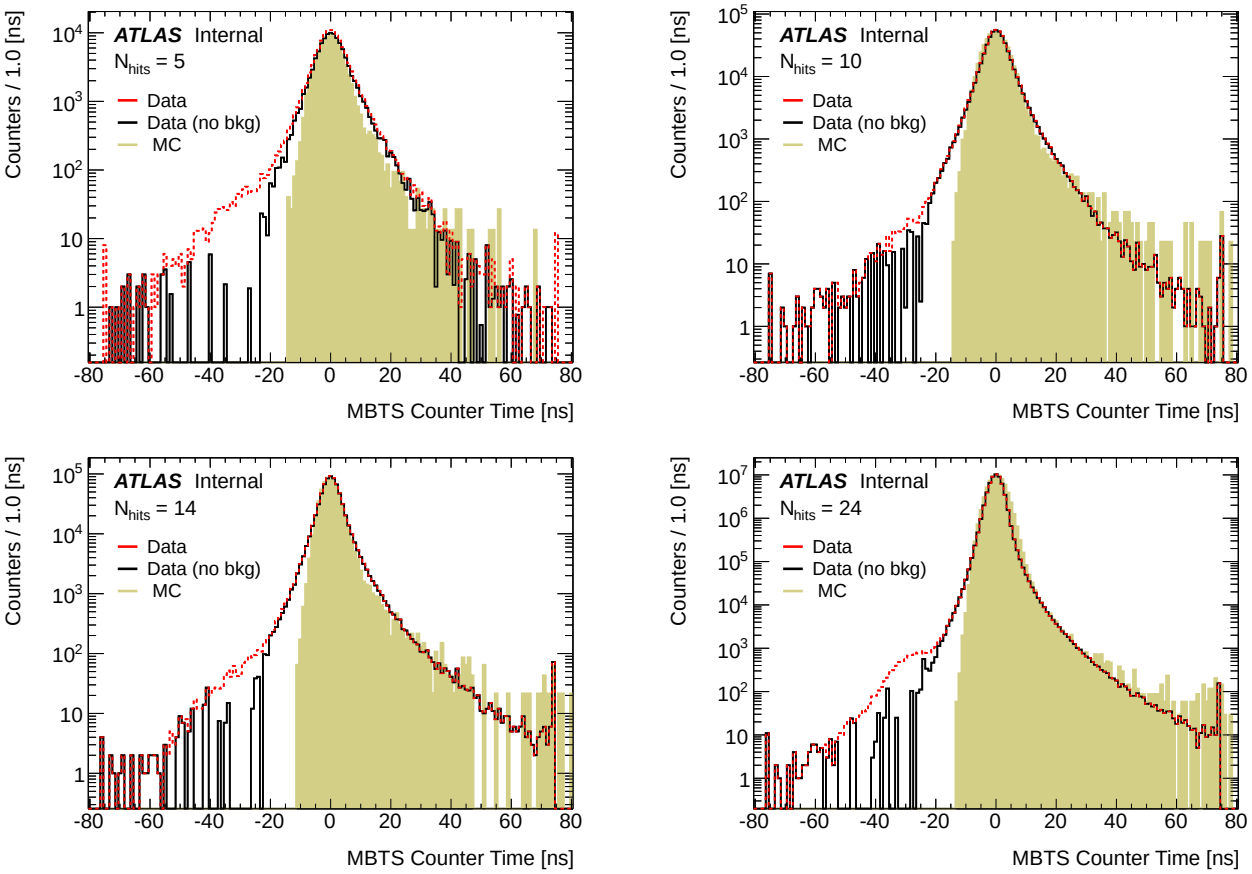

Figure 5.14: Distribution of MBTS time for modules with charge above threshold (0.15 pC) in events with a given number of hits. The solid black (dashed red) line represents the data after (before) background subtraction; the filled yellow histogram represents MC simulation.

[9]The negative tails of the data are not well modelled, but this has no effect on the results.

5.10 Trigger efficiency

The efficiency of the MBTS_1 trigger is measured in events selected by an independent trigger. It is defined as:

$$\epsilon_{\text{trig}} = \frac{\text{Number of events passing MBTS\_1 and independent trigger}}{\text{Number of events passing independent trigger}}. \tag{5.6}$$

The independent triggers used are: RD0_FILLED, LUCID_1, and LHCF_1. The samples of events selected by these triggers are referred to as random, LUCID and LHCf samples. The LUCID and LHCf triggers require activity at higher pseudorapidity than the MBTS acceptance ($5.6 < |\eta| < 6.0$ and $|\eta| > 8.4$ respectively), thus do not bias the efficiency measurement.

The MBTS hit distribution in the random sample is shown in figure 5.15 with and without the MBTS_1 trigger requirement. The random sample is strongly peaked towards zero hits. This peaking feature and the studies of the backgrounds (section 5.9) suggest that the random sample is completely dominated by afterglow. To avoid large systematic uncertainties associated with noise and background subtraction, this sample is discarded.

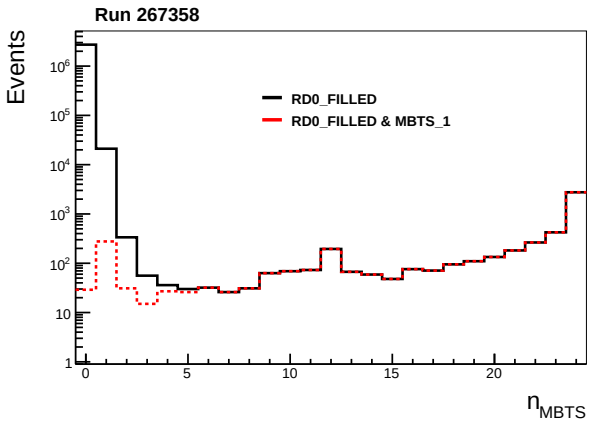

Figure 5.15: MBTS hit distribution in events selected by a random trigger (in black), and in events selected by both the random trigger and the MBTS_1 trigger (in red).

The MBTS hit distribution in the LUCID and LHCf samples is shown in figure 5.16 with and without the MBTS_1 trigger requirement. Neither the LHCf-sample nor the LUCID-sample present the background level observed in the random-sample[10]. This is

[10]The peak at 0 hits is expected as LUCID or LHCf can trigger on events outside the MBTS acceptance.

attributed to the fact that the LUCID_1 and LHCf_1 triggers select mainly proton–proton collisions, whereas the random trigger only selects an average number of inelastic events equal to μ, which in the run used in this study is very low.

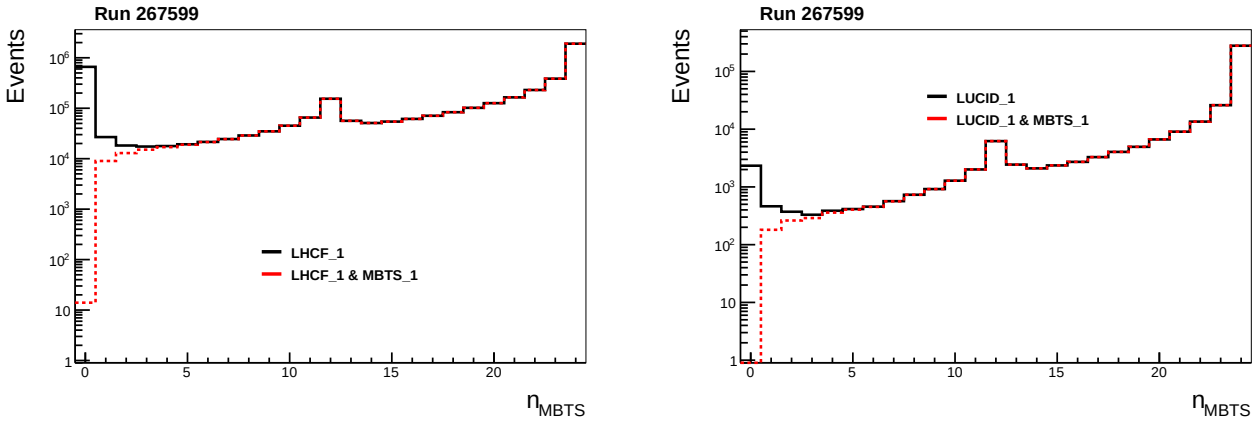

Figure 5.16: MBTS hit distribution in events selected by LHCf (left) and LUCID (right) triggers. The distribution in events that also pass the MBTS_1 trigger selection is shown in red.

Figure 5.17 shows the MBTS_1 trigger efficiency, measured in the LHCf and LUCID samples, as a function of hit multiplicity. The measurements relevant for the signal region are consistent within statistical uncertainties. The inefficiency observed for events with a low number of hits above the offline threshold (0.15 pC) is expected because the trigger hardware threshold for this run was set to 100–150 mV (0.5–0.7 pC).

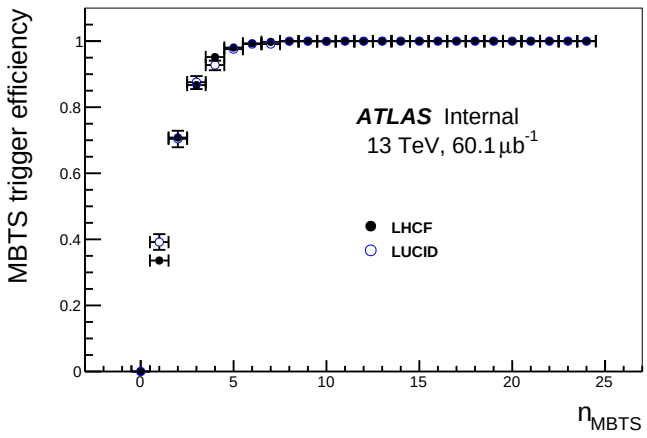

Figure 5.17: Trigger efficiency as a function of number of MBTS hits, measured in events triggered by LHCf and LUCID.

The MBTS_1 trigger is fully efficient for events with 6 or more hits. The vast majority of the events in the signal region fall into that category. This is quantified by using the measured trigger efficiency to weight individual events. The relative difference between the number of events in the signal region and the corresponding number of weighted events is 0.29% and 0.27% for the LUCID and LHCf-based results respectively, corresponding to a 99.71% and 99.73% efficiency. The systematic uncertainty on these estimates is described in section 5.12.

5.11 Tuning the diffraction fraction in simulation

The MC generators described in section 5.5 differ not only in their modelling of diffractive interactions but also in the breakdown of the inelastic cross-section into single-diffractive, double-diffractive and non-diffractive cross-sections (σ_{SD}, σ_{DD}, and σ_{ND} respectively). Given that diffractive interactions tend to produce large rapidity gaps, a significant fraction of those will produce hits in just one arm of the MBTS. The probability for a non-diffractive event to produce that topology is negligible ($< 0.01\%$ according to simulations). Therefore, the diffractive fraction,

$$f_{\mathrm{D}} = (\sigma_{\mathrm{SD}} + \sigma_{\mathrm{DD}})/\sigma_{\mathrm{inel}}, \tag{5.7}$$

can be constrained in a data-driven way by measuring the following ratio:

$$R_{\mathrm{SS}} = \frac{\text{Corrected number of events in single-sided sample}}{\text{Corrected number of events in inclusive sample}}. \tag{5.8}$$

The corrections involved in the ratio are related to trigger efficiency and backgrounds. The background contamination in the single-sided events is 5.8% and is significantly larger than in the inclusive selection that is 1.2%. The trigger efficiency is also different, mainly because the single-sided selection has a larger fraction of events in the low-multiplicity region where the trigger is not fully efficient. This amounts to an efficiency of 99.7% in the inclusive sample and 97.5% in the single-sided sample.

The result is $R_{\mathrm{SS}} = 10.4 \pm 0.4\%$. The statistical uncertainty on R_{SS} is negligible, and the dominant source of systematic uncertainty arises from the background subtraction. Figure 5.18 shows the MC predictions and their f_{D} dependence along with the measured R_{SS}. The dependence of R_{SS} and f_{D} is different for each model but in all cases it increases monotonically, as expected.

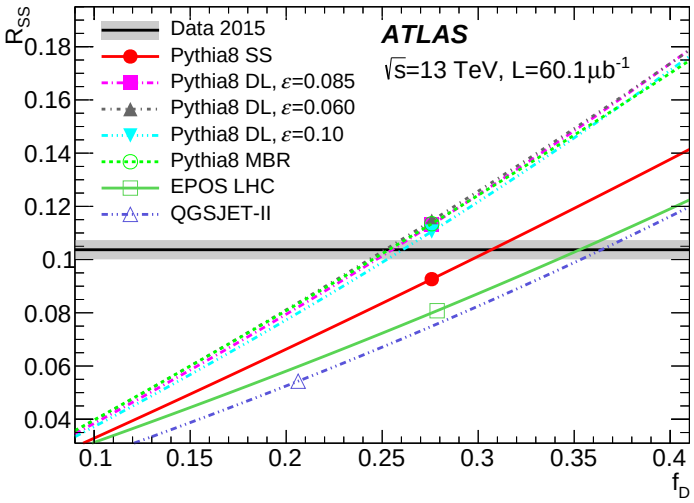

Figure 5.18: R_{SS} as a function of f_{D} for different simulations. The nominal value of f_{D} is shown with a marker along each curve. The black line and grey band represent the measured R_{SS} and its uncertainty. Source: Ref. [7]

For each MC simulation, f_{D} is varied to match the measured R_{SS}. The fitted value of f_{D} in the PYTHIA8 samples is between 25 % and 31 %, depending on the model (compared to the default value of 28 %). For the QGSJET-II (EPOS LHC) model the fitted f_{D} is 36 % (35 %), and differs significantly from the default value of 21 % (28 %). This method cannot constrain the relative fraction of single-diffractive and double-diffractive interactions; the sensitivity to this is checked by taking the diffractive events to be entirely single-diffractive or to be evenly divided between single and double-diffractive. No significant effect is observed. The adjusted f_{D} is then used to determine C_{MC}, shown in table 5.5, and to obtain the MBTS hit distributions in simulation (shown in section 5.13).

Table 5.5: Fitted f_{D} values and corresponding C_{MC} value for different simulations. The uncertainty on C_{MC} is derived from the uncertainty on f_{D}.

Generator	Model	C_{MC}	f_{D}
PYTHIA8	SS	0.9970 ± 0.0003	0.31 ± 0.01
PYTHIA8	DL, $\epsilon = 0.085$	0.9940 ± 0.0003	0.26 ± 0.01
PYTHIA8	DL, $\epsilon = 0.060$	0.9931 ± 0.0003	0.25 ± 0.01
PYTHIA8	DL, $\epsilon = 0.10$	0.9923 ± 0.0003	0.26 ± 0.01
PYTHIA8	MBR	0.9943 ± 0.0003	0.25 ± 0.01
EPOS LHC	–	0.9925 ± 0.0009	0.35 ± 0.02
QGSJET-II	–	0.9904 ± 0.0007	0.36 ± 0.02

5.12 Systematic uncertainties

This section describes the systematic uncertainties on the measurement of the fiducial cross-section. In decreasing order of importance, these are: luminosity calibration, background subtraction, MC simulation, and trigger efficiency. An additional systematic uncertainty related to the extrapolation of the fiducial cross-section to the full phase-space is discussed in section 5.13.

Luminosity

The luminosity uncertainty is $\pm 1.9\%$ as described in detail in chapter 4.

Background

The background estimate presented in section 5.9 contains a series of first-order approximations, such as: neglecting the difference of the per-bunch charge in the estimate of beam-gas background; taking an average over all empty bunch crossings as an estimate of afterglow; and neglecting time variations of the rate of beam-gas or afterglow over the duration of the run. Further studies are limited by the available data sample. A conservative uncertainty of $\pm 50\%$ of the net correction is assigned. This corresponds to $\pm 0.6\%$ systematic uncertainty on $\sigma_{\text{inel}}^{\text{fid}}$.

MC correction factor

The sources of systematic uncertainty on the C_{MC} determination include the uncertainty on the MBTS module efficiency, the uncertainty on the fitted f_{D}, and the difference among the MC simulations. As detailed in Ref. [7], the MBTS module efficiency uncertainty is estimated to be $\pm 1\%$; varying them by this amount in simulation has a $\pm 0.2\%$ impact on $\sigma_{\text{inel}}^{\text{fid}}$. The uncertainty on f_{D} arises from the uncertainty on the measured R_{SS}; this has a $\pm 0.1\%$ impact on $\sigma_{\text{inel}}^{\text{fid}}$. The average of all C_{MC} results (shown in table 5.5) is taken as the nominal value (0.993); the full envelope of the results ($\pm 0.4\%$) is taken as systematic uncertainty. Thus, the total systematic uncertainty on $\sigma_{\text{inel}}^{\text{fid}}$ that arises from C_{MC} is estimated to be $\pm 0.5\%$.

Trigger efficiency

Most events in the inclusive sample have large hit multiplicities. The overall correction is rather small (+0.3%). Hence, the small differences between the LHCF-based and LUCID-based results in the low-multiplicity region (see figure 5.17) have a negligible impact on the $\sigma_{\text{inel}}^{\text{fid}}$ measurement. However, given that no dedicated triggers for background studies in the LUCID or LHCf acceptance are available, a conservative uncertainty of $\pm 0.3\%$ is estimated to cover possible biases due to background. This variation also covers changes in the trigger efficiency caused by varying the hit threshold from the nominal 0.15 pC to higher values of 0.3 pC or 0.5 pC. This corresponds to a $\pm 0.3\%$ systematic uncertainty on $\sigma_{\text{inel}}^{\text{fid}}$.

5.13 Results

Hit distributions

The measured hit distributions in the inclusive and single-sided samples are shown in figure 5.19 along with expectations from different MC simulations. About 60% of the events in the inclusive sample fired all 24 MBTS modules; these are mainly non-diffractive events. A tail extends to lower multiplicities. A prominent feature at 12 hits is expected from diffractive interactions that fired all MBTS modules in one side. This feature is reasonably well described by the simulations. Individually, none of the simulations describe the data in the low multiplicity region. However, if considered as an ensemble they bracket the data.

The hit distribution in single-sided events is also dominated by high-multiplicity events, about 40% fired all the 12 MBTS modules. The tail at low multiplicities is dominated by low-mass diffractive scattering; this is well described by the Pythia8 DL and MBR models. The Epos LHC, QGSJet-II, and Pythia8 SS models do not describe the distribution well.

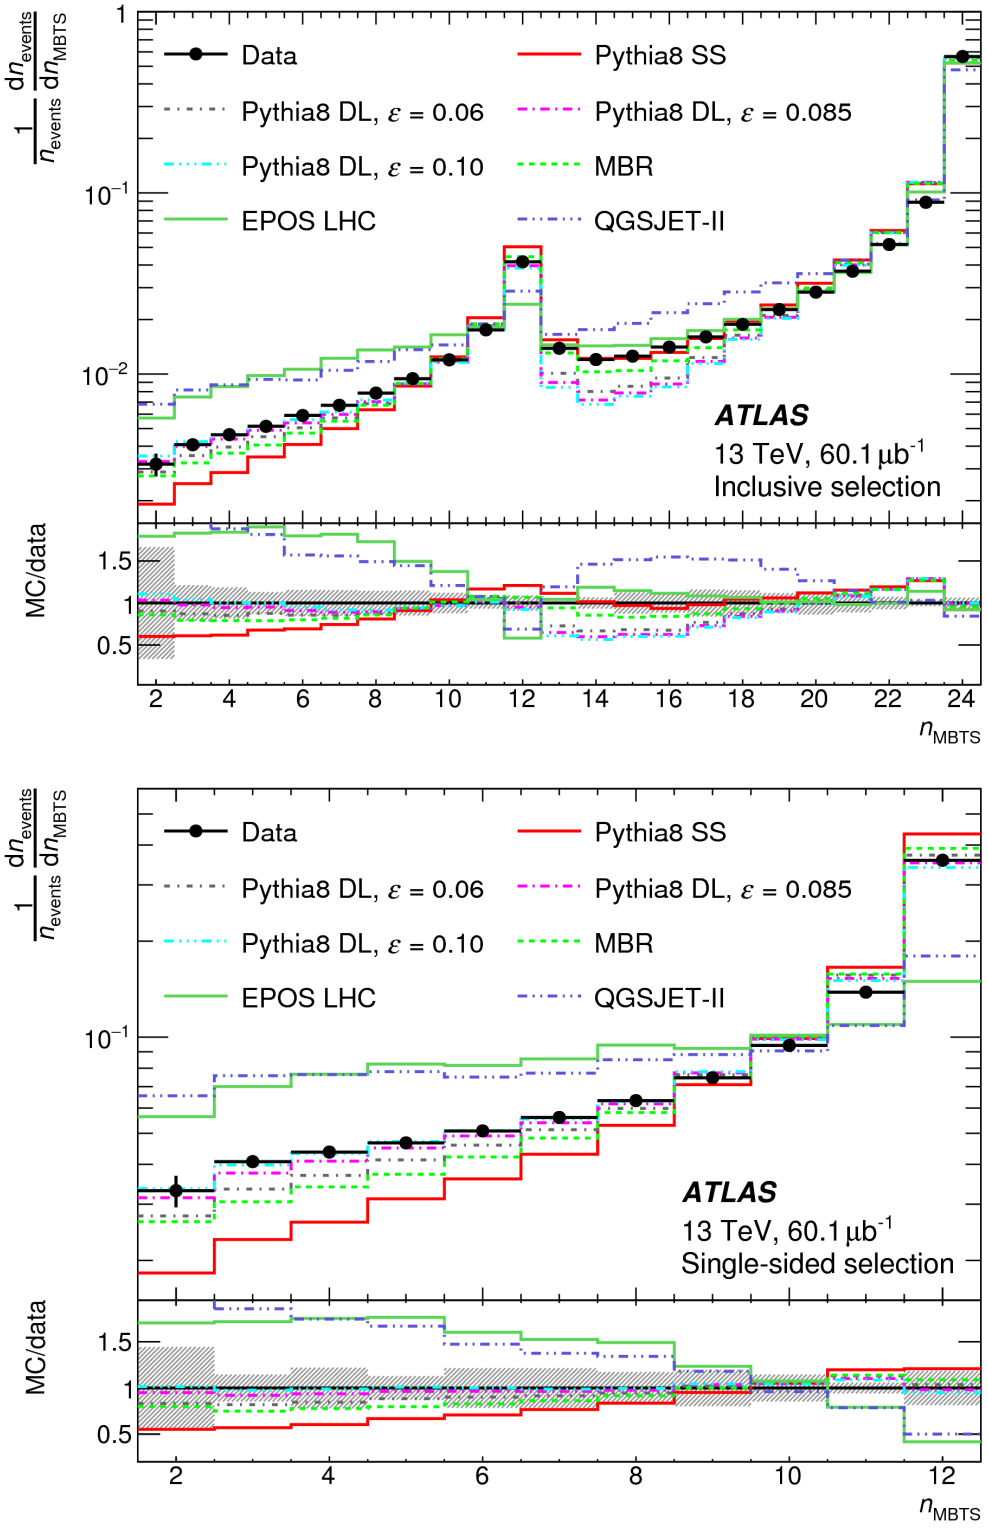

Figure 5.19: MBTS hit distribution in data, after background subtraction, compared to several MC samples for the inclusive selection (top panel) and the single-sided selection (bottom panel). The ratio of the various MC models to the data is also shown. Error bars represent statistical uncertainties only and are smaller than the marker size in most bins. The systematic uncertainty on the data is shown as a shaded band in the ratio plot. Source: Ref. [7].

Fiducial cross-section

The values of all quantities needed for the fiducial cross-section calculation are shown in table 5.6 along with their relative systematic uncertainty and the corresponding impact on $\sigma_{\text{inel}}^{\text{fid}}$.

Table 5.6: Values used for the fiducial cross-section calculation, their relative systematic uncertainties and corresponding impact on the result.

Factor	Value	Rel. uncertainty	$\Delta\sigma_{\text{inel}}^{\text{fid}} / \sigma_{\text{inel}}^{\text{fid}}$
Number of selected events (N)	4159074	–	–
Number of background events (N_{BG})	51187	$\pm 50\%$	$\pm 0.6\%$
Integrated luminosity [μb^{-1}] ($\mathcal{L}$)	60.1	$\pm 1.9\%$	$\pm 1.9\%$
Trigger efficiency (ϵ_{trig})	99.7%	$\pm 0.3\%$	$\pm 0.3\%$
MC correction factor (C_{MC})	99.3%	$\pm 0.5\%$	$\pm 0.5\%$

All uncertainties are assumed to be Gaussian and uncorrelated. The fiducial cross-section, calculated according to equation 5.1, is:

$$\sigma_{\text{inel}}^{\text{fid}} = 68.1 \pm 0.6 \text{ (exp.)} \pm 1.3 \text{ (lumi.) mb}, \tag{5.9}$$

where the first uncertainty ("experimental") includes all sources of systematic uncertainty except the luminosity uncertainty. The luminosity uncertainty is the leading source of uncertainty of the measurement. The statistical uncertainty is negligible.

Figure 5.20 shows the comparison of the measured fiducial cross-section with predictions from MC simulations. The Pythia8 DL model predicts values of 71.0 mb, 69.1 mb and 68.1 mb for $\varepsilon = 0.06$, 0.085 and 0.10 respectively; all of these are compatible with the measurement. The Pythia8 MBR model predicts 70.1 mb, also in agreement with the measurement. The Epos LHC (71.2 mb) and QGSJet-II (72.7 mb) predictions exceed the data by 2–3 σ. The Pythia8 SS model predicts 74.4 mb, and thus exceeds the measured value by $\approx 4\ \sigma$.

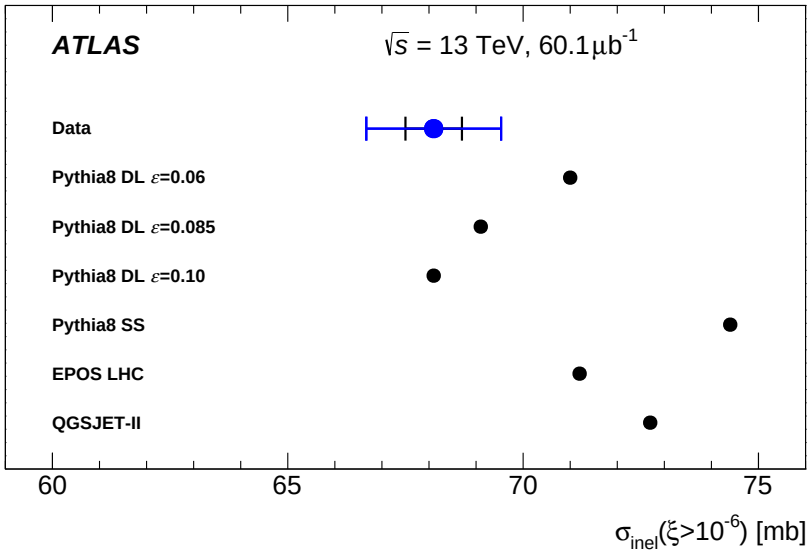

Figure 5.20: The measured fiducial inelastic proton–proton cross-section compared with predictions from several MC simulations. The innermost error bar on the measurement shows the experimental uncertainty not including the luminosity uncertainty, and the outermost error bars show the total uncertainty. Source: Ref. [7].

Extrapolation to full cross-section

The extrapolation of the measured fiducial cross-section to the total inelastic cross-section is based on previous ATLAS measurements at 7 TeV to minimize the model dependence. The results based on the ALFA sub-detector include the total cross-section [103], while the result based on the MBTS sub-detector corresponds to a fiducial cross-section in the range $\xi > 5 \times 10^{-6}$ [99]. The difference between these two measurements yields the cross-section outside the fiducial region (at 7 TeV):

$$\begin{aligned} \sigma_7(\xi < 5 \times 10^{-6}) &= \sigma_7(\xi > m_p^2/s) - \sigma_7(\xi > 5 \times 10^{-6}) \\ &= (71.34 \pm 0.91)\ \text{mb} - (60.33 \pm 2.10)\ \text{mb} \\ &= 11.0 \pm 2.3\ \text{mb}. \end{aligned}$$

The uncertainties of the ALFA and MBTS measurement are treated as uncorrelated. Both uncertainties are dominated by the luminosity calibration, but in 2010 (relevant for the MBTS data), it was dominated by the beam current measurement, whereas in 2011 and afterwards (relevant for the ALFA data) the uncertainty on the beam current measurements is negligible.

The cross-section outside the fiducial region at 7 TeV can be used to extrapolate the fiducial measurement at 13 TeV to the full cross-section. The fiducial region at 13 TeV is different than at 7 TeV because the same acceptance covers lower ξ values due to the increase in center-of-mass energy. MC simulation is used to estimate a scale factor, shown in table 5.7, to account for this difference.

Table 5.7: Ratio of fiducial cross-sections at 7 and 13 TeV for different MC simulation models. The statistical uncertainty on the ratio is negligible.

Model	$\sigma_{13}^{\mathrm{MC}}(\xi < 10^{-6})/\sigma_{7}^{\mathrm{MC}}(\xi < 5\times10^{-6})$
Pythia8 SS	0.782
Pythia8 DL $\epsilon = 0.085$	1.015
Pythia8 DL $\epsilon = 0.06$	1.04
Pythia8 DL $\epsilon = 0.10$	0.999
Pythia8 MBR	1.097
Epos LHC	1.051
QGSJet-II	1.093

The value for the nominal Pythia8 DL model with $\epsilon = 0.085$ is used; a systematic uncertainty that covers differences among all the models except Pythia8 SS is assigned. Pythia8 SS is neglected because it fails to describe the fiducial cross-section data (at more than 4 σ as shown in figure 5.20). Thus, the correction factor is taken to be:

$$\frac{\sigma_{13}^{\mathrm{MC}}(\xi < 10^{-6})}{\sigma_{7}^{\mathrm{MC}}(\xi < 5\times10^{-6})} = 1.015 \pm 0.081. \tag{5.10}$$

The total inelastic proton–proton cross-section at 13 TeV is estimated as:

$$\sigma_{\mathrm{inel}} = \sigma_{\mathrm{inel}}^{\mathrm{fid}} + \sigma_{7}(\xi < 5\times10^{-6}) \times \frac{\sigma_{13}^{\mathrm{MC}}(\xi < 10^{-6})}{\sigma_{7}^{\mathrm{MC}}(\xi < 5\times10^{-6})}. \tag{5.11}$$

The overall additive correction to the fiducial cross-section at 13 TeV is 11.2 ± 2.5 mb, where the uncertainty includes the data uncertainty on the 7 TeV measurements[11]. The

[11] The TOTEM experiment measured 72.9 ± 1.5 mb for the total inelastic cross-section at 7 TeV. This is slightly higher than the 71.34 ± 0.91 mb measured using ALFA. If the uncertainty-weighted average of the two measurements were used, 71.76 mb, for the extrapolation, then the total cross-section measurement would be 0.43 mb higher. This change is well within the uncertainties.

final result is:

$$\begin{aligned}\sigma_{\text{inel}} &= 79.3 \pm 0.6\ (\text{exp.}) \pm 1.3\ (\text{lum.}) \pm 2.5\ (\text{extrap.})\ \text{mb} \\ &= 79.3 \pm 2.9\ \text{mb}.\end{aligned}$$

Here the last uncertainty is related to the extrapolation procedure. This result agrees well with predictions from the PYTHIA8 DL models of 78.4 mb; PYTHIA8 MBR of 76.6 mb; EPOS of 79.0 mb and QGSJET-II of 80.2 mb. The measurement is compared with various theoretical predictions in figure 5.21. These calculations [128–132] are based on QCD-inspired models that attempt to encapsulate the non-perturbative aspects of the theory. The predictions agree within the uncertainty on the data.

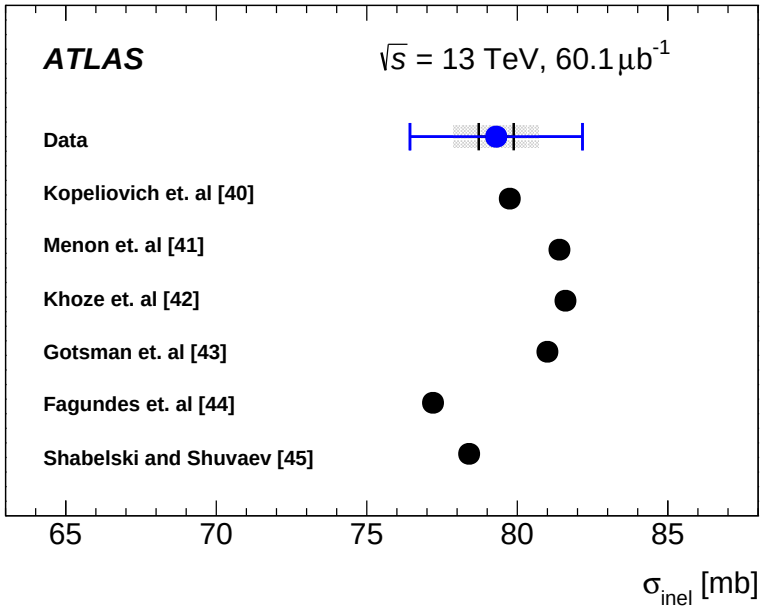

Figure 5.21: The measured inelastic cross-section at 13 TeV compared with various theoretical predictions. The innermost error bar shows the experimental uncertainty only, the shaded region shows the experimental and extrapolation uncertainties together, and the outermost error bar shows the total uncertainty. Source: Ref. [7].

Figure 5.22 shows the measured inelastic proton–proton cross-section and compares it with previous data from fixed-target, collider and cosmic ray experiments. Compared to the measurement with the ALFA detector at 7 TeV the cross-section at 13 TeV is increased by $11 \pm 4\%$. The correlation between the measurements (arising from the use of ALFA 7 TeV data in the extrapolation) is taken into account in the uncertainty.

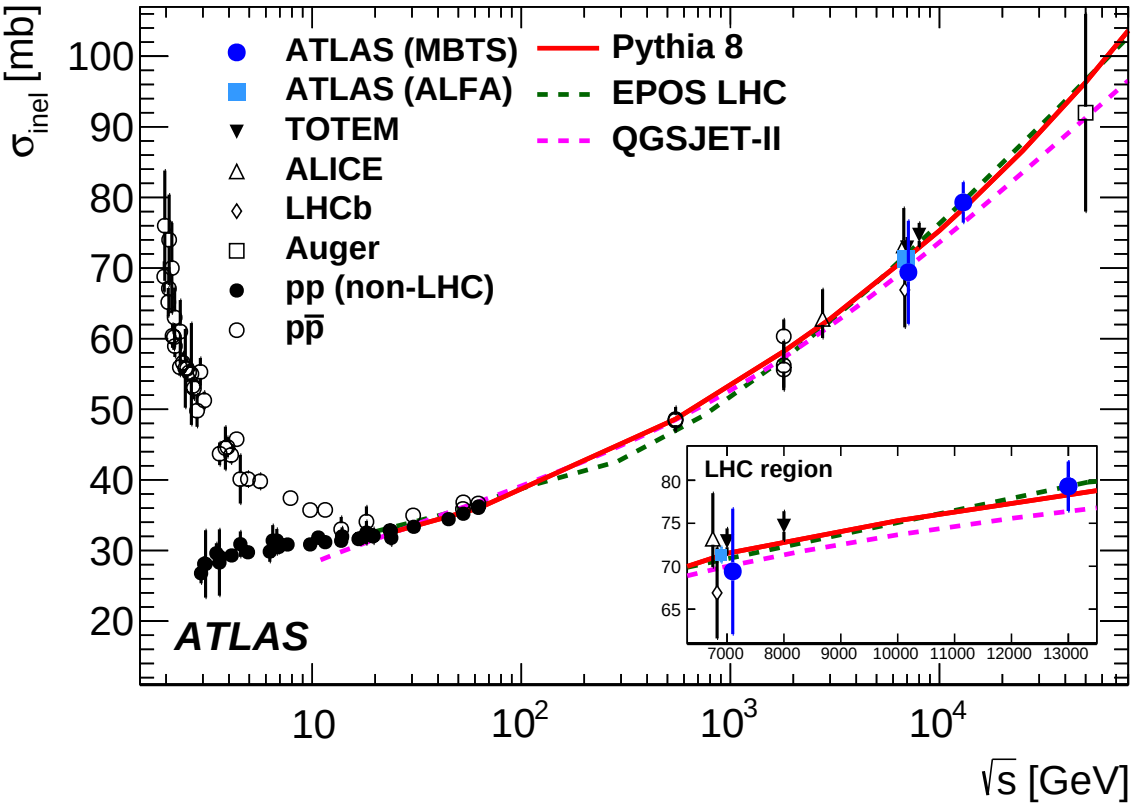

Figure 5.22: Inelastic proton–proton and proton–antiproton cross-section as a function of center-of-mass energy. Measurements of several fixed-target, collider and cosmic ray experiments are shown. Some of the LHC data points have been slightly shifted in the horizontal position for clarity. The data are compared to the predictions from Pythia8, Epos LHC, and QGSJet-II. The inset shows a zoom to the LHC measurements. Source: Ref. [7].

5.14 Conclusions and future prospects

This work presents the first measurement of the proton–proton inelastic cross-section at 13 TeV. It is consistent with an increase with respect to lower energies, continuing a trend that has been observed since 1971. The work presented in the previous chapter enabled a luminosity measurement with a precision of 1.9%, which is one the most precise luminosity calibrations performed by ATLAS to date. The precision achieved in this cross-section measurement (total uncertainty of 3.7%) represents a significant improvement with respect to a similar ATLAS measurement at 7 TeV [99] (total uncertainty of 11%).

Future prospects for similar measurements include an indirect measurement based on elastic scattering via the optical theorem with the ALFA sub-detector. As ALFA is specifically designed to measure the elastic cross-section, it is expected to surpass

the precision of the measurement presented in this work. Nevertheless, these two measurements complement each other as they involve completely different techniques.

Another possibility would be to measure the proton–proton and proton–lead inelastic cross-sections at 5 TeV. ATLAS already collected 5 TeV proton–lead data in 2013 and 5 TeV low-luminosity proton–proton data in 2015. A direct comparison would enable an estimate of nuclear effects and put constraints on the models used to describe the nuclear physics involved in interactions of cosmic rays and the atmosphere.

The luminosity calibration of the proton–lead data set, presented in chapter 4, has an accuracy of 2.7%; this suggests that a cross-section measurement would be possible at the 4–5% level. Similar precision could be achieved in the 5 TeV proton–proton case (as the work presented in this thesis demonstrates). These measurements have the potential to improve the modelling of cosmic-ray and atmosphere interactions, which is key to understanding the origin of the knee in the cosmic-ray energy spectrum.

Bibliography

[1] B. Hommels et al., *Detailed studies of full-size ATLAS12 sensors*, Nucl. Instr. Meth. Phys. A. doi:10.1016/j.nima.2016.03.042.

[2] K. Hara et al., *Charge collection and field profile studies of heavily irradiated strip sensors for the ATLAS inner tracker upgrade*, Nucl. Instr. Meth. Phys. A. doi:10.1016/j.nima.2016.04.035.

[3] M. Arratia et al., *Luminosity determination in p+Pb collisions at* $\sqrt{s_{NN}}$ *= 5.02 TeV in 2013*, ATL-COM-DAPR-2014-001, https://cds.cern.ch/record/1645872/.

[4] M. Arratia, *Orbit drift, beam-beam, and non-factorization corrections for the analysis of vdM scan data for calibration of pp collisions at 13 TeV in 2015* , ATL-COM-DAPR-2016-003, https://cds.cern.ch/record/2143254/.

[5] ATLAS Luminosity Task Force, *Final 2015 luminosity results*, https://indico.cern.ch/event/525538/.

[6] M. Arratia, *Luminosity determination for the 2015 PbPb run*, ATL-COM-DAPR-2016-004, https://cds.cern.ch/record/2152730/.

[7] ATLAS Collaboration, *Measurement of the Inelastic Proton-Proton Cross Section at* $\sqrt{s}$ *= 13 TeV with the ATLAS Detector at the LHC*, `arXiv:1606.02625 [hep-ex]`.

[8] L. Evans and P. Bryant, *LHC Machine*, JINST **3** (2008) S08001.

[9] J. J. Goodson, *Search for Supersymmetry in States with Large Missing Transverse Momentum and Three Leptons including a Z-Boson*, CERN-THESIS-2012-053, http://cds.cern.ch/record/1449722.

[10] CERN, *"LHC dipole magnet"*, http://www.lhc-closer.es/dipole.

[11] H. Wiedemann, *Particle Accelerator Physics*. Springer, third ed., 2007.

[12] M. Kuhn, *Emittance Preservation at the LHC*, CERN-THESIS-2013-031, http://cds.cern.ch/record/1544573/.

[13] *MAD-X programme*, `http://frs.web.cern.ch/frs/Xdoc/mad-X.html`.

[14] J. M. Jowett, *The LHC as a Nucleus-Nucleus Collider*, J. Phys. **G35** (2008) 104028.

[15] *CERN Press Release, November 2015*, https://press.cern/press-releases/2015/11/lhc-collides-ions-new-record-energy.

[16] C. A. Salgado et al., *Proton-Nucleus Collisions at the LHC: Scientific Opportunities and Requirements*, J. Phys. **G39** (2012) 015010.

[17] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08003.

[18] ATLAS Collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*, ATLAS-PERF-2015-01-002, https://cds.cern.ch/record/2158820.

[19] ATLAS Collaboration, *Electron performance measurements with the ATLAS detector using the 2010 LHC proton-proton collision data*, Eur. Phys. J. **C72** (2012) 1909.

[20] ATLAS Collaboration, *Jet energy measurement and its systematic uncertainty in proton-proton collisions at* $\sqrt{s} = 7$ *TeV with the ATLAS detector*, Eur. Phys. J. **C75** (2015) 17.

[21] ATLAS Collaboration, *ATLAS magnet system: Technical Design Report*, CERN-LHCC-97-018, https://cds.cern.ch/record/338080/.

[22] ATLAS Collaboration, *Muon reconstruction performance of the ATLAS detector in proton-proton collision data at* $\sqrt{s}$ *=13 TeV*, Eur. Phys. J. **C76** (2016).

[23] S. White, *The ATLAS zero degree calorimeter*, Nucl. Instrum. Meth. **A617** (2010) 126–128.

[24] G. Apollinari, O. Bruning, and L. Rossi, *High Luminosity LHC Project Description*, CERN-ACC-2014-0321, https://cds.cern.ch/record/1974419.

[25] *Physics at a High-Luminosity LHC with ATLAS*, in *Proceedings, Community Summer Study 2013: Snowmass on the Mississippi (CSS2013): Minneapolis, MN, USA, July 29-August 6, 2013*. `arXiv:1307.7292 [hep-ex]`.

[26] ATLAS Collaboration, *Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment*, CERN-LHCC-2012-022, https://cds.cern.ch/record/1502664.

[27] L. Rossi, *Pixel Detectors*. Springer, 2006.

[28] H. Spieler, *Semiconductor Detector Systems*. Oxford University Press, 2007.

[29] G. Lutz, *Semiconductor Radiation Detectors*. Springer, 2007.

[30] F. Hartmann, *Evolution of Silicon Sensor Technology in Particle Physics*. Springer, 2009.

[31] O. Pooth, *The CMS Silicon Strip Tracker*. Springer, 2010.

[32] M. Krammer, *Silicon detectors in high energy physics experiments*, Scholarpedia, 10(10):32486. doi:10.4249/scholarpedia.32486.

[33] Y. Unno et al., *Development of n-on-p silicon sensors for very high radiation environments*, Nucl.Instrum.Meth. **A636** (2011) S24–S30.

[34] Y. Unno et al., *Development of n^+-in-p large-area silicon microstrip sensors for very high radiation environments: ATLAS12 design and initial results*, Nucl. Instrum. Meth. **A765** (2014) 80–90.

[35] ATLAS Collaboration, *Supply of Silicon Microstrip Sensor of ATLAS12 specification*, https://indico.cern.ch/event/331594/.

[36] G. Lindstrom et al., *Radiation hardness of silicon detectors: A challenge from high-energy physics*, Nucl. Instrum. Meth. **A426** (1999) 1–15.

[37] P. S. Miyagawa and I. Dawson, *Radiation background studies for the Phase II inner tracker upgrade.*, ATL-UPGRADE-INT-2012-002, https://cds.cern.ch/record/1502669.

[38] J. Bohm et al., *Evaluation of the bulk and strip characteristics of large area n-in-p silicon sensors intended for a very high radiation environment*, Nucl. Instrum. Meth. **A636** (2011) S104–S110.

[39] M. Marcela et al., *Evaluation of Bulk and Surface Radiation Damage of Silicon*

Sensors for the ATLAS Upgrade, PoS **Vertex2014** (2015) 050.

[40] M. Hernandez, *Alibava : A portable readout system for silicon microstrip sensors*, doi:10.1109/NSSMIC.2008.4775030. (2007) 5 p.

[41] CMS Tracker Group Collaboration, W. Adam et al., *Impact of low-dose electron irradiation on n^+p silicon strip sensors*, Nucl. Instrum. Meth. **A803** (2015) 100–112.

[42] H. Bichsel, *Straggling in Thin Silicon Detectors*, Rev. Mod. Phys. **60** (1988) 663–699.

[43] G. Lindstrom et al., *Radiation hard silicon detectors developments by the RD48 (ROSE) Collaboration*, Nucl. Instrum. Meth. **A466** (2001) 308–326.

[44] M. Mikestikova et al., *Study of surface properties of ATLAS12 strip sensors and their radiation resistance*, Nucl. Instr. Meth. Phys. A. doi:10.1016/j.nima.2016.03.056.

[45] M. Mori et al., *Evaluation of the performance of irradiated silicon strip sensors for the forward detector of the ATLAS Inner Tracker Upgrade*, Nucl. Instr. Meth. Phys. A. doi:10.1016/j.nima.2016.04.044.

[46] K. Lalwani, G. Jain, R. Dalal, K. Ranjan, and A. Bhardwaj, *Study the radiation damage effects in Si microstrip detectors for future HEP experiments*, Nucl. Instrum. Meth. **B379** (2016) 262–264.

[47] CMS Tracker Collaboration, *Charge Collection Efficiency Simulations of Irradiated Silicon Strip Detectors*, JINST **9** no. 12, (2014) C12010.

[48] CMS Tracker Collaboration, *A Method to Simulate the Observed Surface Properties of Proton Irradiated Silicon Strip Sensors*, JINST **10** no. 04, (2015) C04025.

[49] ATLAS Collaboration, *Improved luminosity determination in pp collisions at $\sqrt{s}$ = 7 TeV using the ATLAS detector at the LHC*, Eur. Phys. J. **C73** (2013) 2518.

[50] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s}$ = 8 TeV using the ATLAS detector at the LHC*, `arXiv:1608.03953 [hep-ex]`.

[51] P. Grafstrom and W. Kozanecki, *Luminosity determination at proton colliders*, Prog. Part. Nucl. Phys. **81** (2015) 97–148.

[52] O. S. Bruning et al., *LHC Design Report Vol.1: The LHC Main Ring*, CERN-2004-003-V1, https://cds.cern.ch/record/782076/.

[53] *CERN Press Release, April 2011*, http://press.cern/press-releases/2011/04/lhc-sets-world-record-beam-intensity.

[54] B. Richter, *High Energy Colliding Beams; What Is Their Future?*, Rev. Accel. Sci. Tech. **7** (2014) 1–8.

[55] I. Hinchliffe et al., *Luminosity goals for a 100 TeV pp collider*, Int. J. Mod. Phys. **A30** (2015) 1544002.

[56] ATLAS Collaboration, *Measurement of $W^{\pm}$ and Z-boson production cross sections in pp collisions at $\sqrt{s}$ =13 TeV with the ATLAS detector*, Phys. Lett. **B759** (2016) 601–621.

[57] ATLAS and CMS Collaborations, *Combination of ATLAS and CMS top-quark pair cross-section measurements using proton-proton collisions at $\sqrt{s}$ = 8 TeV*, ATLAS-CONF-2014-054, https://cds.cern.ch/record/1951322/.

[58] ATLAS Collaboration, *Measurement of the $W \to \ell\nu$ and $Z/\gamma^* \to \ell\ell$ production cross sections in proton-proton collisions at $\sqrt{s}$ = 7 TeV with the ATLAS detector*, JHEP **12** (2010) 060.

[59] CMS Collaboration, *Measurements of Inclusive W and Z Cross Sections in pp Collisions at $\sqrt{s}$ = 7 TeV*, JHEP **01** (2011) 080.

[60] R. Jones, *First Results from the LHC Beam Instrumentation Systems*, CERN-LHC-Performance-Note-006, https://cds.cern.ch/record/1156348.

[61] V. Zhilich, *The luminosity measurement at the Belle detector*, Nucl. Instrum. Meth. **A494** (2002) 63–67.

[62] ALEPH Collaboration, *Measurement of the absolute luminosity with the ALEPH detector*, Z. Phys. **C53** (1992) 375–390.

[63] H1 Collaboration, *Determination of the Integrated Luminosity at HERA using Elastic QED Compton Events*, Eur. Phys. J. **C72** (2012) 2163, [Erratum: Eur. Phys. J.C74,2733(2012)].

[64] G. Watt, *Parton distribution function dependence of benchmark Standard Model total cross sections at the 7 TeV LHC*, JHEP **09** (2011) 069.

[65] S. van der Meer, *Calibration of the Effective Beam Height in the ISR*, CERN-ISR-PO-68-31, https://cds.cern.ch/record/296752/.

[66] M. Ferro-Luzzi, *Proposal for an absolute luminosity determination in colliding beam experiments using vertex detection of beam-gas interactions*, Nucl. Instrum. Meth. **A553** (2005) 388–399.

[67] C. Barschel, *Precision Luminosity Measurement at LHCb with Beam-gas Imaging*, CERN-THESIS-2013-301, https://cds.cern.ch/record/1693671.

[68] LHCb Collaboration, *Precision luminosity measurements at LHCb*, JINST **9** (2014) P12005.

[69] O. Viazlo et al., *ATLAS LUCID detector upgrade for LHC Run 2*, ATL-FWD-PROC-2015-004, https://cds.cern.ch/record/2062038/.

[70] ATLAS Collaboration, *Luminosity Determination in pp Collisions at $\sqrt{s} = 7$ TeV Using the ATLAS Detector at the LHC*, Eur. Phys. J. **C71** (2011) 1630.

[71] *ATLAS luminosity results*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults.

[72] ATLAS Collaboration, *Z boson production in p+Pb collisions at $\sqrt{s_{NN}}$ =5.02 TeV measured with the ATLAS detector*, Phys. Rev. **C92** (2015) 044915.

[73] ATLAS Collaboration, *Measurement of high-mass dimuon pairs from ultraperipheral lead-lead collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector at the LHC*, ATLAS-CONF-2016-025, http://cds.cern.ch/record/2157689.

[74] C. Barschel et al., *Results of the LHC DCCT Calibration Studies*, CERN-ATS-2012-026, https://cds.cern.ch/record/1425904/.

[75] D. Belohrad et al., *The LHC Fast BCT system: A comparison of Design Parameters with Initial Performance*, CERN-BE-2010-010, https://cds.cern.ch/record/1267400/.

[76] C. Ohm and T. Pauly, *The ATLAS beam pick-up based timing system*, Nucl. Instrum. Methods Phys. Res., A **623** (2009) 558–560. 3 p.

[77] G. Anders et al., *Study of the Relative Bunch Populations for Luminosity Calibration*, CERN-ATS-2012-028, https://cds.cern.ch/record/1427726/.

[78] CMS Collaboration, *Measurement of the Inelastic Cross Section in Proton-Lead*

Collisions at $\sqrt{s_{NN}} = 5.02$ TeV, Phys. Lett. **B759** (2016) 641–662.

[79] ALICE Collaboration, *Measurement of the Cross Section for Electromagnetic Dissociation with Neutron Emission in Pb-Pb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Rev. Lett. **109** (2012) 252302.

[80] W. Kozanecki et al., *Observation of Beam-beam Deflections with LHC Orbit Data*, CERN-ACC-2013-0006, https://cds.cern.ch/record/1581723/.

[81] G. Anders, *Absolute luminosity determination for the ATLAS experiment*, CERN-THESIS-2013-111, https://cds.cern.ch/record/1595219.

[82] J. P. Koutchouk, *Measurement of the beam position in the LHC main rings*, LHC-BPM-ES-0004, https://cds.cern.ch/record/1068133/.

[83] H. Pirumov, *Impact of beam orbit drifts on 2012 calibration*, ATL-COM-DAPR-2015-001, https://cds.cern.ch/record/1981519/.

[84] J. Olexa et al., *Prototype system for phase advance measurements of LHC small beam oscillations*, CERN-ATS-2013-038, https://cds.cern.ch/record/1546401/.

[85] A. Boccardi et al., *LHC Luminosity calibration using the Longitudinal Density Monitor*, CERN-ATS-2013-034, https://cds.cern.ch/record/1556087/.

[86] S. Webb, *Determining single beam density profiles from beam spot and luminosity data in vdM scans*, ATL-DAPR-2014-004, https://cds.cern.ch/record/1662745/.

[87] B. Cole et al., *Luminosity calibration for February 2013 2.76 TeV pp measurements*, ATL-COM-DAPR-2014-009, https://cds.cern.ch/record/1747966.

[88] ATLAS Collaboration, *Measurement of differential J/ψ production cross sections and forward-backward ratios in p+Pb collisions with the ATLAS detector*, Phys. Rev. **C92** (2015) 034904.

[89] ALICE Collaboration, *Measurement of visible cross sections in proton-lead collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV in van der Meer scans with the ALICE detector*, JINST **9** (2014) P11003.

[90] U. Amaldi, *An ISR Discovery: The Rise of the Proton Proton Cross-Section*, Adv. Ser. Dir. High Energy Phys. **23** (2015) 257–286.

[91] S. R. Amendolia et al., *Measurement of the total proton proton cross-section at the ISR*, Phys. Lett. **B44** (1973) 119–124.

[92] U. Amaldi et al., *The Energy dependence of the proton proton total cross-section for center-of-mass energies between 23 and 53 GeV*, Phys. Lett. **B44** (1973) 112–118.

[93] W. Heisenberg, *Production of Mesons as a Shock Wave Problem*, Z. Phys. **133** (1952) 65–79.

[94] M. Froissart, *Asymptotic behaviour and subtractions in the Mandelstam representation*, Phys. Rev. **123** (1961) 1053.

[95] H. Cheng and T. T. Wu, *Limit of Cross-Sections at Infinite Energy*, Phys. Rev. Lett. **24** (1970) 1456–1460.

[96] G. B. Yodh et al., *Evidence for rapidly rising p-p total cross-section from cosmic ray data*, Phys. Rev. Lett. **28** (1972) 1005–1008.

[97] K. Nakamura et al. (Particle Data Group), *Accelerator Physics of Colliders*, J. Phys. G **36** (2010) 075021.

[98] Particle Data Group Collaboration, K. A. Olive et al., *Review of Particle Physics*, Chin. Phys. **C38** (2014) 090001.

[99] ATLAS Collaboration, *Measurement of the Inelastic Proton-Proton Cross-Section at $\sqrt{s} = 7$ TeV with the ATLAS Detector*, Nature Commun. **2** (2011) 463.

[100] ALICE Collaboration, *Measurement of inelastic, single- and double-diffraction cross sections in proton–proton collisions at the LHC with ALICE*, Eur. Phys. J. **C73** (2013) 2456.

[101] LHCb Collaboration, *Measurement of the inelastic pp cross-section at a centre-of-mass energy of $\sqrt{s} = 7$ TeV*, JHEP **02** (2015) 129.

[102] CMS Collaboration, *Measurement of the inelastic proton-proton cross section at $\sqrt{s} = 7$ TeV*, Phys. Lett. B **722** (2013) 5.

[103] ATLAS Collaboration, *Measurement of the total cross section from elastic scattering in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, Nucl. Phys. B **889** (2014) 486.

[104] TOTEM Collaboration, *Evidence for non-exponential elastic proton–proton differential cross-section at low $|t|$ and $\sqrt{s} = 8$ TeV*, Nucl. Phys. **B899** (2015) 527–546.

[105] TOTEM Collaboration, *Luminosity-independent measurements of total, elastic and inelastic cross-sections at* $\sqrt{s} = 7$ *TeV*, Europhys. Lett. **101** (2013) 21004.

[106] TOTEM Collaboration, *Luminosity-Independent Measurement of the Proton-Proton Total Cross Section at* $\sqrt{s} = 8$ *TeV*, Phys. Rev. Lett. **111** (2013) 012001.

[107] M. Block and F. Halzen, *Experimental Confirmation that the Proton is Asymptotically a Black Disk*, Phys. Rev. Lett. **107** (2011) 212002.

[108] M. Block et al., *Comprehensive fits to high energy data for* σ*,* ρ*, and* B *and the asymptotic black-disk limit*, Phys. Rev. **D92** (2015) 114021.

[109] D. Fagundes et al., *Exploring central opacity and asymptotic scenarios in elastic hadron scattering*, Nucl. Phys. **A946** (2016) 194–226.

[110] V. V. Anisovich et al., *Hadron collisions at ultrahigh energies: black disk or resonant disk modes?*, Phys. Rev. **D90** (2014) 074005.

[111] CMS Collaboration, *Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV*, `arXiv:1607.03663 [hep-ex]`.

[112] D. d'Enterria et al., *Constraints from the first LHC data on hadronic event generators for ultra-high energy cosmic-ray physics*, Astropart. Phys. **35** (2011) 98–113.

[113] M. Block, *Ultra-high Energy Predictions of Proton-Air Cross Sections from Accelerator Data: an Update*, Phys. Rev. **D84** (2011) 091501.

[114] T. Sjöstrand et al., *PYTHIA 6.4 physics and manual*, JHEP **05** (2006) 026.

[115] T. Sjöstrand et al., *A Brief Introduction to PYTHIA 8.1*, Comput. Phys. Comm. **178** (2008) 852.

[116] T. Pierog et al., *EPOS LHC : test of collective hadronization with data measured at the CERN Large Hadron Collider*, Phys. Rev. C **92** (2015) 034906.

[117] S. Ostapchenko, *Monte Carlo treatment of hadronic interactions in enhanced Pomeron scheme: I. QGSJET-II model*, Phys. Rev. D **83** (2011) 014018.

[118] ATLAS Collaboration, *Summary of ATLAS Pythia 8 tunes*, ATL-PHYS-PUB-2012-003, https://cds.cern.ch/record/147410.

[119] P. Skands et al., *Tuning PYTHIA 8.1: the Monash 2013 Tune*, Eur. Phys. J. C **74** (2014) 3024.

[120] A. Schuler and T. Sjöstrand, *Hadronic diffractive cross-sections and the rise of the total cross-section*, Phys. Rev. **D49** (1994) 2257.

[121] S. Navin, *Diffraction in Pythia*, `arXiv:1005.3894 [hep-ph]`.

[122] A. Donnachie and P. Landshoff, *Elastic scattering and diffraction dissociation*, Nucl. Phys. B **244** (1984) 322.

[123] R. Ciesielski and K. Goulianos, *MBR Monte Carlo Simulation in PYTHIA8*, `arXiv:1205.1446`.

[124] ATLAS Collaboration, *Rapidity gap cross sections measured with the ATLAS detector in pp collisions at* $\sqrt{s} = 7$ *TeV*, Eur. Phys. J. **C72** (2012) 1926.

[125] CMS Collaboration, *Measurement of diffraction dissociation cross sections in pp collisions at* $\sqrt{s} = 7$ *TeV*, Phys. Rev. **D92** (2015) 012003.

[126] S. Agostinelli et al., *Geant4 - A Simulation Toolkit*, Nucl. Instrum. Meth. A **506** (2003) 250.

[127] M. Arratia et al., *Measurement of the Inelastic Proton-Proton Cross-Section at* $\sqrt{s} = 13$ *TeV*, ATL-COM-PHYS-2015-423, https://cds.cern.ch/record/2016959/.

[128] B. Z. Kopeliovich et al., *Two-scale hadronic structure and elastic pp scattering: predicted and measured*, Phys. Rev. D **86** (2012) 051502.

[129] M. Menon and P. Silva, *An updated analysis on the rise of the hadronic total cross-section at the LHC energy region*, Int. J. Mod. Phys. A **28** (2013) 1350099.

[130] V. Khoze et al., *High Energy Elastic and Diffractive Cross Sections*, Eur. Phys. J. C **74** (2014) 2756.

[131] E. Gotsman et al., *A model for strong interactions at high energy based on the CGC/saturation approach*, Eur. Phys. J. C **75** (2015) 18.

[132] D. Fagundes et al., *Soft edge of hadron scattering and mini-jet models for the total and inelastic pp cross-sections at LHC and beyond*, Phys. Rev. D **91** (2015) 114011.