Construction, Commissioning and Performance Measurements of the Inner Tracking System Upgrade (ITS2) of ALICE at the LHC

Thesis submitted in accordance with the requirements of the University of Liverpool for the degree of Doctor in Philosophy by

James Philip Iddon[1,2]

April 2022

Supervisors:

Marielle Chartier[1]

Felix Reidt[2]

Roy Lemmon[3]

[1]University of Liverpool

[2]CERN

[3]Daresbury Laboratory

For my father,

Phil Iddon,

who believed my brother and I could do anything and was proud of us to try.

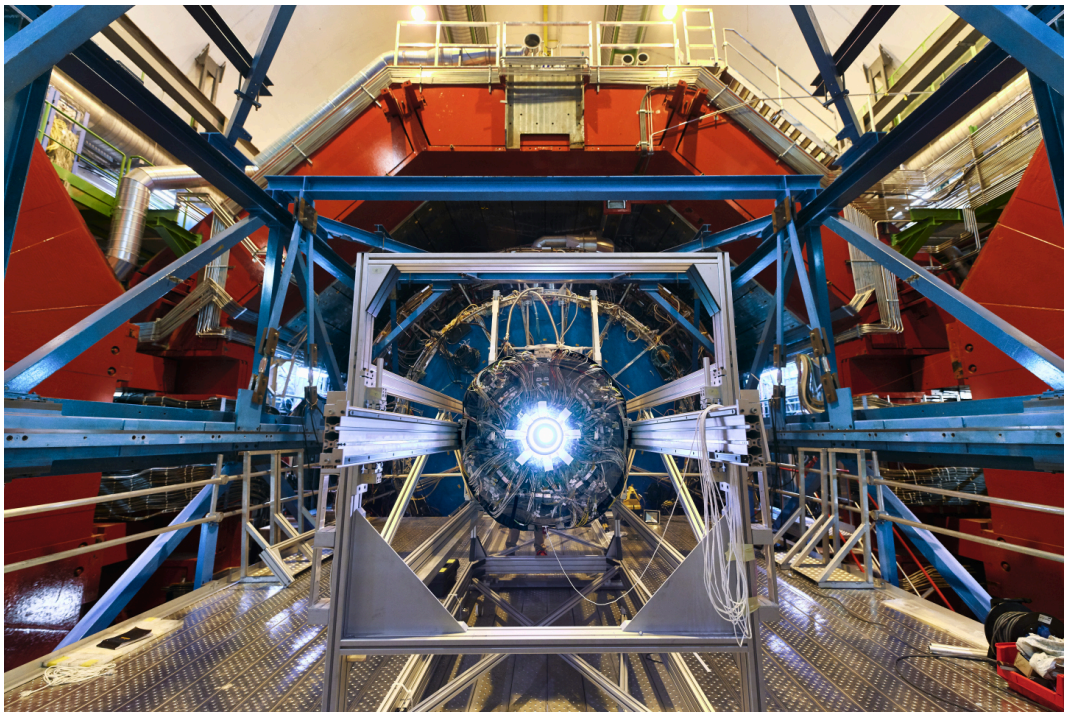

The LHC Runs 1 and 2 ALICE Inner Tracking System being removed from the ALICE detector in February 2019 during Long Shutdown 2. © CERN. Taken from [1].

Abstract

A Large Ion Collider Experiment (ALICE), situated on the Large Hadron Collider (LHC), is optimised to study relativistic heavy-ion collisions in order to quantitatively characterise the Quark Gluon Plasma. The inner-most detector system, the Inner Tracking System (ITS), is responsible for primary and secondary vertex reconstruction in the vicinity of the interaction point, as well as stand-alone tracking of particles unable to reach the rest of the ALICE sub-detectors, due to momentum cut-off and acceptance limitations.

The ITS was upgraded over the LHC Long Shutdown 2 (late 2018 - early 2022) with the aims of improving the impact parameter resolution, tracking efficiency, transverse momentum resolution, and readout rate. This was achieved by replacing the ITS used in LHC Runs 1 and 2 by seven concentric layers of silicon pixel detector, utilising a Monolithic Active Pixel Sensors (MAPS) design known as ALPIDE, together with the reduction of the radius of the inner-most layer by 16 mm, the reduction of the material budget and the decrease of the pixel pitch.

Over the last few years, the upgraded ITS, ITS2, was successfully constructed at several international sites before being assembled into the concentric barrel structure of the final tracker at CERN. The detector was then integrated into the final services including the readout and power systems before undergoing a comprehensive commissioning campaign.

This thesis concerns the construction, characterisation, and commissioning of the outer-most four layers of the ITS2, known as the Outer Barrel. Namely, the construction and characterisation of the detector modules and staves, the verification of the final system, and the acquisition of the first real particle data set and subsequent performance measurements.

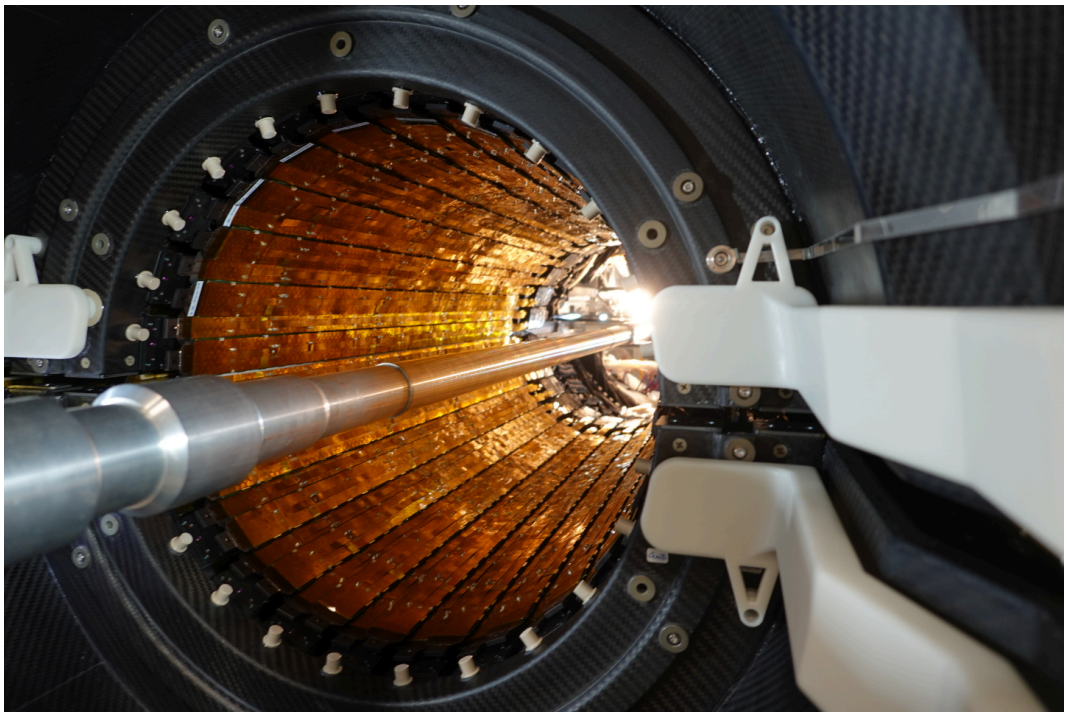

The Outer Barrel of the upgraded ALICE Inner Tracking System fully installed in the ALICE detector in March 2021. © CERN. Taken from [2].

Acknowledgements

I would like to thank Marielle Chartier, my primary supervisor, for giving me the opportunity to work within the ALICE ITS collaboration, at the University of Liverpool, Daresbury Laboratory and CERN. You have been my supervisor from my Master's project until the submission of this thesis and I am grateful for every second of it.

My time at Daresbury Laboratory was supervised by Roy Lemmon, whom I would like to thank for allowing me to integrate into the Daresbury group and inspiring me to start training for climbing. The lunchtime discussions with Roy, Marielle and John Dainton were always very interesting.

I would like to express my gratitude towards Felix Reidt, my CERN supervisor, for welcoming me into the CERN ITS group, entrusting me with interesting, important work and providing me with answers to any technical questions I could think of. The depth of your knowledge of the ITS never failed to impress me.

Matteo Lupi, whom I regard as an unofficial supervisor during my time at CERN, deserves a special mention for always extending my 'quick two minute' questions into fruitful discussions. Despite having unrestricted access to the coffee machine on your desk, I expect that I consumed more knowledge of the readout system than coffee in your office.

I want to thank Arild Velure for welcoming me into our shared office and always providing me with insightful answers to my questions.

I would like to thank Marcello Borri for the lifts to Daresbury Laboratory where the discussions about work were advantageous and the discussions about food, although sometimes confrontational, were always a lot of fun.

I want to express my gratitude towards Matthew Buckland, Giacomo Contin, Andrew Hill, Terry Lee, Roy Lemmon, Jian Liu, Gary Markey and Marcin Poblocki, my colleagues at Daresbury Laboratory and the LSDC, for ensuring the stave construction was as enjoyable as it was productive.

My time at CERN was enormously rewarding thanks to all of the ALICE ITS upgrade group. I would like to express my gratitude towards Gianluca Aglieri Rinella, Dimitra Andreou, Lee Barnby, Iouri Belikov, Stefania Beole, César Ceballos Sanchez, Domenico Colella, Alberto Collu, Giacomo Contin, Lorenzo De Cilladi, Antonello Di Mauro, Piero Giubilato, Leo Greiner, Fabrizio Grosa, Ryan Hannigan, Hartmut Hillemanns, Antoine Junique, Markus Keil, Jochen Klein, Vít

Kuĉera, Tatiana Lazareva, Jian Liu, Matteo Lupi, Magnus Mager, Paulo Martinengo, Luciano Musa, Ivan Ravasenga, Felix Reidt, Joachim Schambach, Serhiy Senyukov, Miko Šuljić, and Arild Velure. It was a privilege to have been able to work with you all.

I would also like to thank those who proof read all or parts of my thesis: Matthew Buckland, Marielle Chartier, John Dainton, Sheila Iddon, Roy Lemmon, Matteo Lupi, Jaime Norman, Felix Reidt, Júlia Tena Vidal and Benedikt Vöelkel.

The submission of this thesis represents four and a half years of hard work, which would not have been possible if it wasn't for the adventures and therapy offered by my friends. I would like to thank those I met in Geneva: Kévin André, Ed Andrews, Peter Bell, Bruno Borbely, Ben Bouvet, Savannah Clawson, James Cruikshank, Laura Fennelly, Joao Florencio, Hannes Gamper, Marta Gomis, Victor Guillen Humbria, Karolina Kulesz, Audrey Kvam, Alex Laut, Matteo Lupi, Phil Marshall, André Rattinger, Adam Rennie, Sølve Slettebak, Francesca Stocker, and Benedikt Vöelkel. I will never forget the adventures and parties we had together and I am grateful for how fun you made my time at CERN. I would also like to thank those who have known me for a long time and have always been there for me: Rory Boyle, Greg Heyes, Dom Riley, James Szylobryt, Matt Turner, Scott Wallis, and Will Wenborn. You always make me laugh and cheer me up when I need it. També m'agradaria donar les gràcies a Julio i Dolors, pel menjar increïble i la càlida benvinguda a casa seva durant les millors vacances possibles quan encara necessitava treballar. Additionally, I need to thank my family: my mum Sheila Iddon, and my brother Chris Iddon, for their unwavering support and for making sure I was well looked after behind the scenes.

Last, but certainly not least, I would like to thank Júlia Tena Vidal for always making me happy. I am fortunate to have been to many wonderful places with you, and I have always preferred to admire your reaction to a beautiful view than the beautiful view itself.

James Philip Iddon
22nd April, 2022

Contents

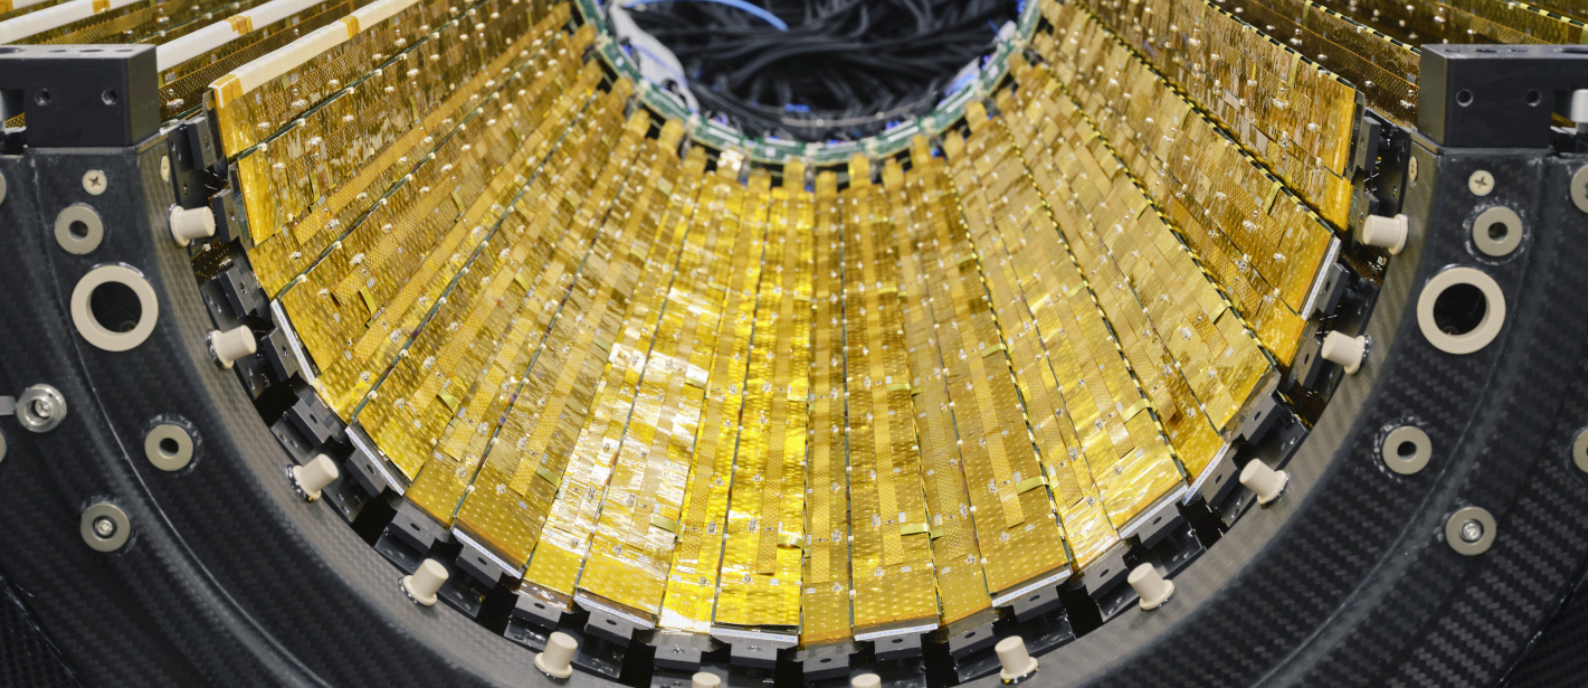

Chapter 1:
Introduction

Particle physics is concerned with understanding the fundamental building blocks of matter and the interactions between them. The first notion that matter is made up of tiny, indivisible units first emerged in ancient Greece, from Democritus and his teacher, Leucippus. The purely philosophical argument hypothesised that the division of matter into smaller units to infinity is impossible, therefore there must exist some fundamental unit from which all matter is comprised. Since the times of ancient Greece, the study of the most fundamental constituents of matter has been significantly developed, no longer existing as a purely philosophical argument but maturing into an experimental science, at which CERN in Geneva is at the forefront.

At CERN, the Large Hadron Collider (LHC) accelerates protons and nuclei to close to the speed of light before colliding them. The energy densities created during the collisions are similar to those seen moments after The Big Bang, creating a unique window through which humans can peer at the early universe. Four detectors sit on the Large Hadron Collider (LHC) beam-line, each optimised to study different processes. One detector, A Large Ion Collider Experiment (ALICE), is optimised to detect the results of heavy-ion collisions. The temperature of these collisions can reach 100,000 times the temperature of the centre of the sun [4]. The energy density

Header photograph shows the fully constructed top half of the Outer Barrel in the commissioning laboratory at CERN. © CERN. Taken from [3].

is so great that conditions resemble the universe mere microseconds after The Big Bang [5], permitting the recreation of the most primordial state of matter ever seen in a Laboratory on Earth.

The primordial state of matter studied by ALICE consists of quarks and gluons. Gluons are the gauge bosons responsible for strong force interactions. When quarks and gluons exist within a region where the energy density is lower than a critical value, they can only be observed in bound states, called hadrons. Removal of a quark from a hadron is not possible because the strength of the strong force between bound quarks increases as the distance between the quarks increases. This is known as confinement. However, when the energy density of a hadronic system surpasses the critical value, quarks and gluons are no longer bound within hadrons and are instead free to roam within the confines of the medium. This state of matter is known as the Quark Gluon Plasma (QGP).

The mechanism of confinement is not fully understood. Recreating the Quark-Gluon Plasma (QGP) at the LHC and studying it with ALICE could shed light on this mechanism, and improve the understanding of why quarks are never observed alone with the energy density seen in the modern Universe. In addition, it is not fully understood where hadrons get their mass. Adding the bare mass of the constituent quarks and gluons of a nucleon, for example, results in $\approx 1\%$ of the total nucleon mass. The remaining 99% is unaccounted for, but it is theorised to be generated from the restoration of chiral symmetry, which occurs close in time to the hadronisation of the QGP. By studying the QGP, ALICE aims to understand how hadronic matter is organised and where it gets its mass.

Since commissioning ALICE has confirmed the existence of the QGP, found by earlier QGP experiments at CERN SPS [6] and BNL RHIC [7], and observed the QGP at unmatched temperatures, densities, and volumes. To continue taking boundary-pushing measurements, ALICE was recently upgraded to better pursue its evolving physics programme, during the second LHC long shutdown (LS2) [8], from 2018 - 2022.

To detect QGP observables at the high particle multiplicities produced by Pb-Pb collisions at the LHC, several detector components work in unison, of which the Inner Tracking System (ITS) is a key element. This is the component of the ALICE detector that sits closest to the beam pipe. The previous ITS consisted of two layers of three different types of silicon tracking detectors: pixel, drift, and strip detectors [9], all of which rely on the generation of a signal when charged particles ionise in

the depletion region of a reverse-biased silicon p-n junction.

The pointing resolution to the secondary decay vertex of particles with a low transverse momentum (p_T), which are more susceptible to the effects of multiple scattering in the ITS layers, is limited by the thickness of the detector. Heavy flavour hadrons preserve their mass, flavour and colour throughout their interaction with the QGP, which means, when they are thought of as Brownian particles within the QGP, their interaction history can be reconstructed using kinematic distributions [9]. As heavy-flavour hadrons traverse the QGP, they are shifted to lower p_T and at sufficiently low p_T will thermalise in the QGP. Charm and beauty hadrons are difficult to detect due to their short lifetime (the $c\tau$ of Λ_c is two times smaller than that of D^0) meaning precise tracking and impact parameter resolution is needed [10].

The main goals of the ITS Upgrade, known as ITS2, are to improve the ability of the ITS to measure heavy-flavour hadrons, thermal photons and low-mass dileptons by extending the measurements to a lower p_T [9], as well as increasing the read-out capabilities, and improving the tracking efficiency such that it is comparable to what could be achieved previously with both the ITS and Time Projection Chamber (TPC) but to much lower values of p_T [9]. The readout rate increase is necessary to incorporate the full interaction at the higher LHC luminosity[1] expected after LS2 of $\mathscr{L} = 6 \times 10^{27}\,\mathrm{m}^{-2}\,\mathrm{s}^{-1}$, corresponding to an interaction rate for Pb-Pb of 50 kHz [9].

The upgrade goals have been realised by replacing the original ITS layers with 7 layers of Monolithic Active Pixel Sensor (MAPS), together with reducing the first ITS layer radius by 16 mm, reducing the material budget and decreasing the pixel size.

Monolithic Active Pixel Sensors (MAPS) are Complementary Metal Oxide Semiconductor (CMOS) pixel sensors where the active volume, signal amplification and read-out circuitry are all integrated onto one chip. MAPS were first used in a collider experiment in the STAR detector at RHIC [11]. The MAPS chip designed for the ALICE ITS is called ALice PIxel DEtector (ALPIDE). A key feature of ALPIDE is a deep p-well shielding an n-well, meaning p-channel Metal Oxide Semiconductor (PMOS) and n-channel Metal Oxide Semiconductor (NMOS) transistors can be incorporated into the chip to form an in-pixel discriminator, amplifier, signal shaper and multiple event buffers. If a particle generates enough charge within the active

[1] Luminosity is a measure of collision frequency per unit area and is defined in Sec. 4.1.

volume of the ALPIDE pixel to surpass a preset threshold, the discriminator outputs a signal. A three signal buffer is implemented in-chip to ensure it is unlikely that a signal is detected and not read out.

Using ALPIDE chips will reduce the material budget to 0.35% of the radiation length, X_0, for the innermost three layers of ITS2, known as Inner Barrel (IB), and to 1.1% X_0 for the outermost four layers of ITS2, known as the Outer Barrel (OB), compared to 1.14% X_0 in the previous ITS [9] extending the physics reach to a lower p_T and making it the thinnest silicon tracking detector at the LHC. The readout rate will be increased to be twice the upgrade requirement (100 kHz for Pb-Pb collisions). Together with the reduction of the radius of the first layer, and the decrease in pixel pitch, the use of MAPS will improve the resolution of the track impact parameter by a factor of three for p_T below 1 GeV/c and a factor of six above 10 GeV/c.

The new ALICE ITS2 was successfully constructed before undergoing an extensive commissioning period. It is currently installed within the ALICE apparatus ready for the start of the LHC Run 3. The ITS2 has approximately 10 m^2 of active area, segmented into nearly 13 billion pixels, making it the largest and most granular silicon pixel tracker ever built. It is the first silicon tracker to use solely MAPS technology.

1.1 Thesis outline

This thesis is concerned with the construction, commissioning and first performance measurements of ITS2. The thesis is organised as follows:

Chapter 2: Silicon tracking technology
A brief overview of semiconductor physics is presented before outlining the use of silicon in tracking detectors and the necessary electronics. An overview of silicon tracking technology is given, as well as an explanation of some concepts necessary to understand the quality of a silicon particle tracker.

Chapter 3: Quark Gluon Plasma
The ALICE physics programme is mainly concerned with the characterisation of the QGP. An overview of the QGP is presented in this chapter, as well as some of its

experimentally observable properties. An inner tracking system ideal for measuring the QGP observables is outlined.

Chapter 4: ALICE upgrade over LS2
An overview of the ALICE detector used in LHC Runs 1 and 2 is presented, before describing the ALICE upgrade programme over LS2. The ALICE ITS upgrade is focussed on and described in depth.

Chapter 5: HIC and stave construction
The ITS2 is made up of a number of staves, where staves consist of a number of Hybrid Integrated Circuits (HICs). The construction of these objects is outlined in this section.

Chapter 6: Stave characterisation
The metrological and electrical characterisation of Outer Barrel staves is presented here.

Chapter 7: ITS readout system
An overview of the ITS2 readout system is presented with a focus on aspects necessary to the ITS2 commissioning campaign.

Chapter 8: Outer Barrel verification
Staves were installed in the barrel-geometry and final services of the ITS2 at CERN. Each stave was tested with the final readout system. The testing protocol and results for the Outer Barrel are outlined here.

Chapter 9: Measurement of cosmic muons
A cosmic muon data set was obtained with the Outer Barrel at the end of 2020. The data collection campaign is discussed and the first measurement of detection efficiency of the Outer Barrel of the ITS2 as a whole is made.

1.2 Author's contribution

I was heavily involved in the construction of HICs and staves at the Liverpool Semiconductor Detector Centre (LSDC) and Daresbury Laboratory respectively. This

work is described in Chap. 5. I was also responsible for the testing of constructed staves at Daresbury Laboratory. The electrical and metrological testing procedure is outlined in Chap. 6, where some results I gathered for all staves in the Outer Barrel are presented. I then moved to Geneva as a CERN Doctoral Student within the ALICE ITS Upgrade group, where I developed the testing software used to characterise the staves with the final readout system. This work is described in Chap. 8. I developed the software which I used to gather a cosmic muon data set with the Outer Barrel at the end of 2020. From this data, I made the first measurement of the detection efficiency of the ITS2 Outer Barrel as a whole. The data acquisition and analysis are presented in Chap. 9.

All the work presented in this thesis is the result of my own work, except where explicit reference is made to the work of others.

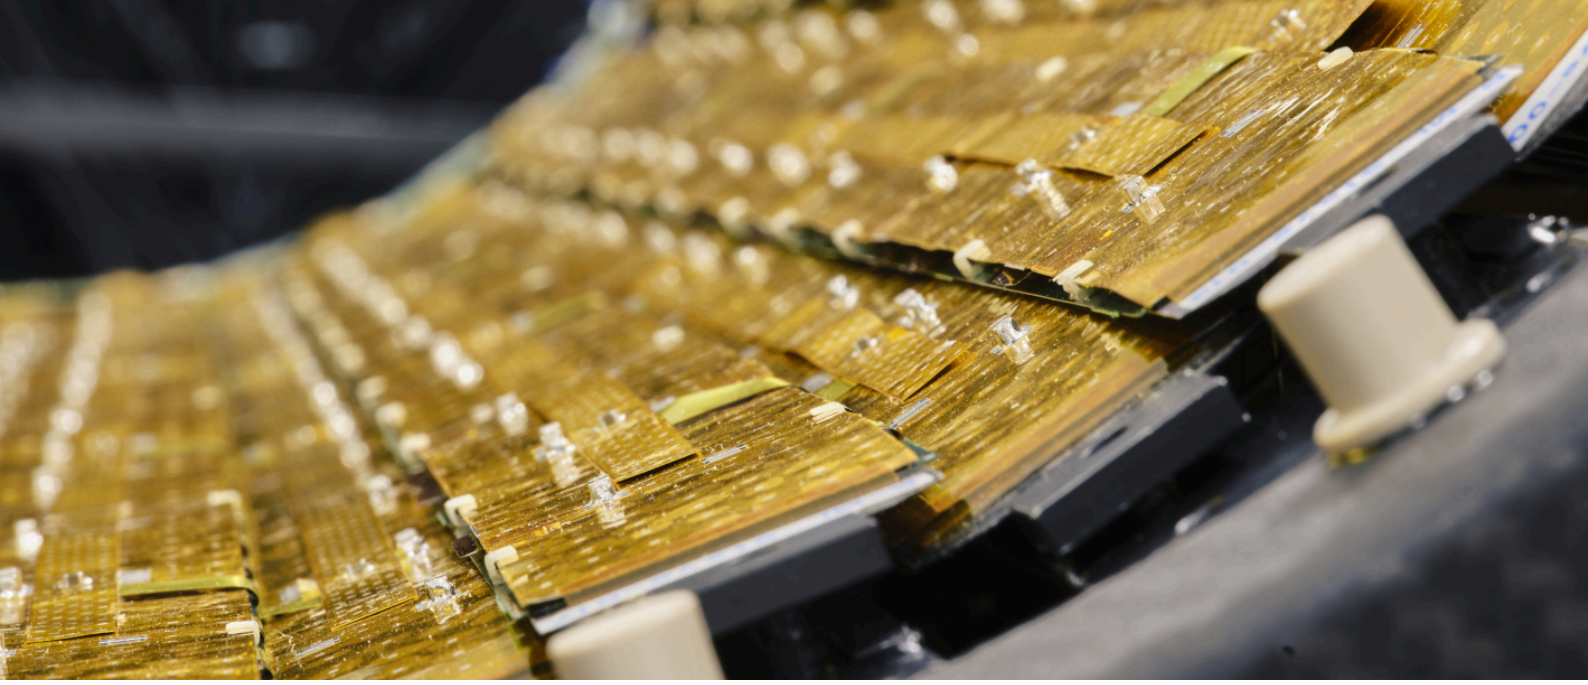

Chapter 2:
Silicon tracking technology

Charged particles can ionise the atoms of the material through which they traverse. The resulting free charge carriers can be collected and measured. This is the basic principle of charged particle detection used by silicon tracking detectors. This chapter will describe the semiconductor physics and semiconductor structures necessary to understand the working principles of silicon tracking detectors, before outlining different types of silicon trackers, finally focussing on the silicon pixel detector, which is fundamental to the design of the ITS2.

2.1 Interaction of particles with matter

The ALICE ITS is designed to track charged particles and photons by exploiting their interaction with silicon. The interaction of charged particles and photons with matter is outlined here.

2.1.1 Charged particles

Ionisation is the dominant energy loss mechanism for charged particles traversing matter. Charged particles traversing through a medium lose energy via elastic col-

Header photograph shows a close up of a fraction of layer three staves installed in the ITS2 barrel. © CERN. Taken from [3].

lisions with electrons within that medium. The average rate of energy loss of a moderately relativistic heavy charged particle to the medium through which it travels is given by the Bethe-Bloch formula [12]:

$$\frac{dE}{dx} = 2\pi N_0 r_e^2 m_e c^2 \rho \frac{Z}{A}\frac{z^2}{\beta^2}\left[ln\left(\frac{2m_e\gamma^2 v^2 W_{max}}{I^2}\right) - 2\beta^2 - \delta - 2\frac{C}{Z}\right], \qquad (2.1)$$

where:

$2\pi N_0 r_e^2 m_e c^2 = 0.1535\,\mathrm{MeV c^2/g}$,

x is the path length in $\mathrm{g/cm^2}$,

$r_e = \frac{e^2}{4\pi m_e c^2} = 2.817 \times 10^{-13}$ cm which is the classical electron radius,

m_e is the mass of an electron in $\mathrm{eV/c^2}$,

I is the average ionisation potential of electrons in eV,

Z and A are the atomic number and atomic weight of the medium respectively,

ρ is the density of the medium in $\mathrm{g/cm^2}$,

z is the charge of the ionising particle in units of electron charge,

δ is the density correction,

C is the shell correction and

W_{max} is the maximum energy transfer in a single collision in eV.

For cases where the mass of the ionising particle $M >> m_e$, $W_{max} \approx 2m_e c^2\beta^2\gamma^2$. The Bethe-Bloch formula describes the region $0.1 < \beta\gamma < 10^4$ to within an accuracy of a few percent.

Energy loss spectrum

The Bethe-Bloch formula describes the average energy loss. The energy loss via ionisation of a particle in matter is statistical in nature since each elastic collision is an independent event. Therefore, the energy loss is better described by a probability density function, the straggling function [13]. Figure 2.1 shows the straggling function for 500 MeV pions in thin silicon. The distribution for thicker absorbers is less skewed but never Gaussian.

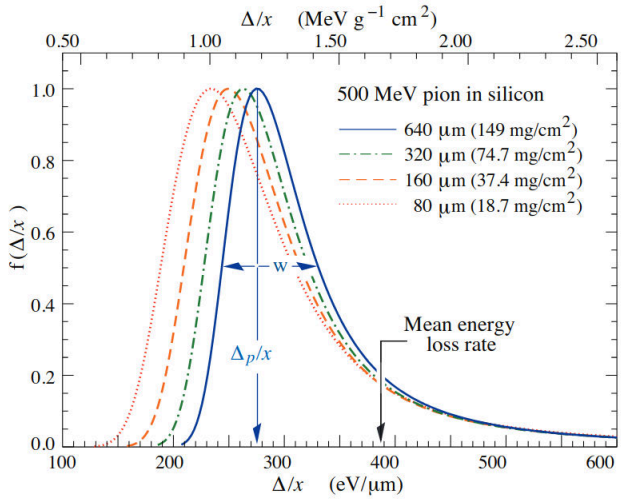

Figure 2.1: Straggling functions in silicon for 500 MeV pions. At the most probable value, Δ_p/x, the distribution is normalised to one. The width, w is the full width at half maximum. Taken from [12].

Bremsstrahlung

Bremsstrahlung, or braking, radiation refers to the photon emission of a decelerating charged particle. Within matter, Bremsstrahlung radiation typically occurs due to the deflection of a charged particle by another charged particle within the medium, i.e. an electron of an atomic nucleus. The energy loss via Bremsstrahlung occurs with a probability of:

$$P \approx Z^2 \frac{E}{m^2}, \tag{2.2}$$

where E is the particle's energy, m is its mass and Z is the charge number of the medium. The radiation length, X_0, is a property characteristic of the medium. X_0 is defined as the average distance over which a high-energy electron experiences an energy loss of $1/e$ due to Bremsstrahlung. X_0 can be approximated as [14]:

$$X_0 = \frac{716.4\,\mathrm{A}}{Z(Z+1)\ln\left(\frac{287}{\sqrt{Z}}\right)}\mathrm{g/cm^2}, \tag{2.3}$$

where A and Z are the atomic and charge numbers of the medium respectively.

Multiple Coulomb scattering

A charged particle traversing through matter will occasionally collide elastically with a nucleus, changing its trajectory due to Coulomb scattering. Each elastic collision deflects the trajectory of the charged particle by a small angle. For multiple small-angle deflections, the distribution of the scattering angle is roughly Gaussian

but with longer tails. The width of the Gaussian for the central 98% of the angular distribution, θ_0, is given by [12]:

$$\theta_0 = \frac{13.6}{\beta cp} z\sqrt{x/X_0}\,[1+0.038\ln(x/X_0)]\,, \tag{2.4}$$

where p, βc and z are the momentum, velocity and charge number of the charged particle respectively, whilst x is the distance of the medium traversed by the charged particle.

2.1.2 Photons

Photons mainly interact with matter in one of three ways:

Photoemission

A photon traversing through matter may transfer its energy to an atom, exciting an atomic electron to a higher shell. In the case of semiconductors, the atomic electron can be excited to the conduction band. The excited electron has an energy, E, of:

$$E = E_\gamma - E_b, \tag{2.5}$$

where E_γ and E_b are the initial photon and binding energy respectively. The vacancy left by the excited electron is filled by a higher energy electron, resulting in the emission of a photon.

Compton scattering

A photon can be scattered elastically on a shell electron, losing an amount of energy dependent on the scattering angle, θ. The energy of the photon after the elastic scattering, E', is given by [15]:

$$E' = \frac{E}{1+\frac{E}{m_e c^2}(1-cos(\theta))}, \tag{2.6}$$

where E is the initial photon energy and m_e is the mass of the electron.

Pair production

If the incident photon has an energy greater than $1.022\,\mathrm{MeV}$ ($2 \times m_e$), then an electron-positron pair can be produced in the Coulomb field of an atomic nucleus.

The photon needs to be in the Coulomb field of an atomic nucleus to satisfy momentum conservation via the recoil of the nucleus. The positron quickly annihilates with an electron producing two photons in opposite directions.

2.2 Basic principles of silicon devices

Silicon is a chemical element with the electronic configuration of $1s^2 2s^2 2p^6 3s^2 3p^2$. When N silicon atoms are isolated, each electron exists in one of two discrete energy levels, where the energy levels are said to be N-fold degenerate. If N isolated silicon atoms are brought closer together spatially, each silicon atom will share its four valence electrons, forming covalent bonds, with four neighbouring silicon atoms resulting in a diamond cubic lattice structure. In this case, the two discrete energy levels split into N closely spaced energy levels, as shown in Fig. 2.2. As $N \to \infty$, the N discrete energy levels are more akin to continuous bands which span a few eV. In the case of semiconductors, the bands are separated by a gap of the order a few eV, where no electron can occupy, known as the band gap [16].

For semiconductors at 0K, the lower energy band is entirely full, named the valence band, whilst the higher energy band is entirely empty, named the conduction band. At room temperature, some electrons have enough energy to become ionised, leaving the valence band, crossing the band gap and entering the conduction band, making the semiconductor weakly conductive due to free electrons and holes. For silicon, the band gap is 1.12 eV whilst the ionisation energy is 3.6 eV, roughly three times larger. The discrepancy is due to most $\left(\frac{2}{3}\right)$ of this energy being lost to the propagation of phonons [17]. For insulators, the band gap is much larger, such that the probability of an electron occupying a state in the conduction band is zero. For conductors, the conduction and valence bands may overlap.

Intrinsic semiconductors contain no impurities in comparison to the number of thermally generated electrons and holes. An example of an intrinsic semiconductor would be a pure silicon crystal[1]. Conversely, extrinsic semiconductors do contain impurities. The electrical properties of a semiconductor can be purposefully altered by adding another chemical element (a dopant) in a process known as doping [18]. Since truly intrinsic semiconductors are difficult to obtain, and their properties are likely to be altered by the doping process, intrinsic semiconductors are rarely used

[1] In practice, it is extremely difficult to obtain the purity required for a truly intrinsic semiconductor.

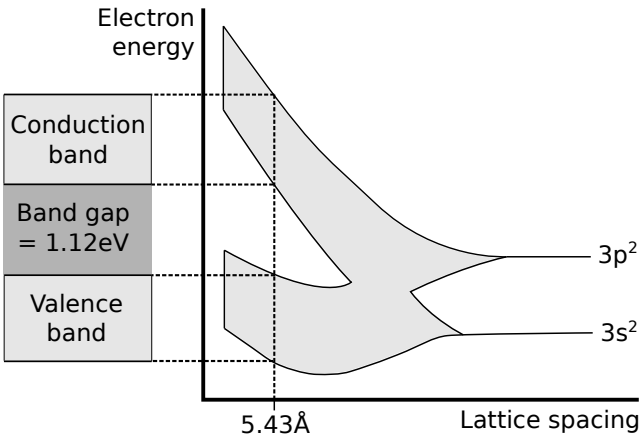

Figure 2.2: Energy levels of silicon atoms within a silicon crystal as a function of lattice spacing. 5.43 Å is the lattice constant of a silicon crystal [16].

in semiconductor devices. An extrinsic semiconductor can be doped with a donor, to create an excess of electrons in the conduction band, called an n-type semiconductor, or with an acceptor, to create an excess of holes in the valence band, called a p-type semiconductor.

For silicon, suitable donor atoms have five valence electrons (i.e. phosphorus) where the fifth electron has a low ionising energy, meaning four covalent bonds are formed with the silicon atoms, whilst the fifth electron resides in the conduction band. A suitable acceptor atom contains three valence electrons (i.e. boron), such that bonding with four silicon atoms results in a hole in the valence band.

The absorption of a photon by an electron in the valence band will excite that electron into the conduction band if the energy of the absorbed photon is greater than the band gap. If the photon energy is lower than the band gap, an electron can be excited into the conduction band, provided that the silicon semiconductor crystal contains intermediate energy levels due to the existence of crystal impurities [19].

2.2.1 Charge carrier transport in semiconductors

Free charge carriers, electrons in the conduction band and holes in the valence band, are not associated with any particular lattice site so are therefore essentially free particles. The mean kinetic energy of free charge carriers is $\frac{3}{2}$ kT, so their mean velocity is $O(10^7\,\mathrm{cm/s})$ at room temperature. Due to lattice vibrations, impurities and defects, their mean free path is $O(10^{-5}\,\mathrm{cm})$ [19]. The transport of the charge carriers through the medium is affected by the application of an external electric

field, known as drift, and the inhomogeneity of the distribution of charge carriers, known as diffusion [20].

Drift

The average displacement of a charge carrier due to random motion within a semiconductor that is not under the influence of an external electric field is 0. However, if an external electric field is applied, the charge carriers have an average displacement determined by the electric field. The drift velocity is dependent on the charge carrier mobility of electrons and holes, μ_e and μ_p, and the strength of the electric field, E. The net average drift velocity is given by:

$$v_e = -\mu_e E, \tag{2.7}$$

$$v_p = \mu_p E. \tag{2.8}$$

These relationships hold for electric fields small enough such that the velocity change caused by the acceleration due to the electric field is small with respect to the thermal velocity. At larger electric field values, the linear relationship breaks down. The velocity increases, leading to an increase in collisions. The velocity reaches saturation, v_{sat}. At room temperature, v_{sat} for silicon is $O(10^7)$ cm/s.

The mobility is a measure of the ease of carrier motion within an electric field. Mobility is given by:

$$\mu = \frac{q\tau}{m}, \tag{2.9}$$

where q is the electric charge of the carrier, τ is the mean free time and m is the mass of the carrier. τ is dependent on temperature, electric field strength, doping concentration and the number of lattice imperfections. An important quantity for silicon detectors is the resistivity. This is a measure of the strength with which a material opposes electric current. The resistivity, ρ, of detector grade silicon is $O(1\,\mathrm{k}\Omega)$ and is given by:

$$\rho = \frac{1}{q(\mu_p n + \mu_e p)}, \tag{2.10}$$

where n and p are the density of free electrons and holes respectively.

Diffusion

In the absence of an electric field at temperatures greater than 0K, charge carriers have a random thermal motion, resulting in no net movement. An inhomogeneous distribution of charge carriers within a semiconductor crystal gives rise to the effect of diffusion, whereby there is a higher probability for a charge carrier to move from the higher concentration to the lower concentration rather than in the opposite direction. The net movement of charge carriers via diffusion produces a current, the diffusion current. The diffusion current density is given by:

$$J_n = q\mu_n nE + qD_n \nabla n, \tag{2.11}$$

$$J_p = q\mu_p pE + qD_p \nabla p, \tag{2.12}$$

where $D_{n/p}$ is the diffusion coefficient, related to mobility by the Einstein equation:

$$D_{n/p} = \frac{kT}{q}\mu_{n/p} \tag{2.13}$$

Recombination and generation

A generation process occurs when electrons have enough energy to move from the valence band to the conduction band, creating equal numbers of electrons and holes in each band. Phonon, photon and charged particle interactions can supply this energy. A recombination process occurs when an electron moves from the conduction band to the valence band, i.e. electron-hole annihilation. If thermal equilibrium is disturbed, when $pn \neq n_i^2$, recombination and generation processes occur to restore the system to equilibrium, i.e. $pn = n_i^2$.

2.2.2 PN junction

A cornerstone of silicon tracking detectors is the pn junction. A pn junction in a semiconductor is created when a region of p-type is next to a region of n-type. Free electrons in the n-type semiconductor diffuse across the junction and combine with holes in the p-type semiconductor. Conversely, free holes in the p-type semiconductor diffuse across the junction and combine with electrons in the p-type semiconductor. Recombination of the electrons in the n-type produces a net posi-

tive space charge, whilst recombination of holes in the p-types produces a negative space charge. Thus, a potential difference exists between the two semiconductors, called the built-in voltage, V_0. The electric field associated with V_0 prevents further diffusion of charge carriers between the two regions. This leaves a region in the centre of the pn junction depleted of free charge carriers known as the depletion region[2]. An illustration of a pn junction is shown in Fig. 2.3.

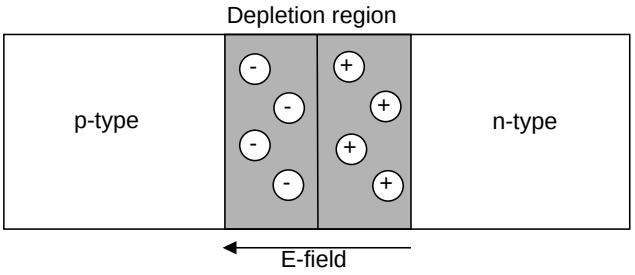

Figure 2.3: Illustration of a pn junction. When a p-type and n-type region are next to one another in a semiconductor, a region devoid of charge carriers is created, known as the depletion region [19].

Reverse bias

When a particle ionises the depletion region, free charge carriers are generated, resulting in an electric current. However, typically, $|V_0|$ is too small to create a depletion region larger than $O(\mu\text{m})$, meaning the active volume of a pn junction is too small for use as a particle detector. The width, W, of the depletion region can be increased with a change in doping concentration, or via the application of a bias voltage across the junction. A reverse bias, where the negative terminal is connected to the p-type, of magnitude V_{bias}, increases W. W depends on $|V_0|$ and the concentration of impurities. W is given by [21]:

$$W = \sqrt{\frac{2\varepsilon_r\varepsilon_0}{q}\left(\frac{1}{N_A}+\frac{1}{N_D}\right)\cdot(V_0 - V_{bias})}, \tag{2.14}$$

where N_A and N_D are the concentrations of acceptor and donor ions respectively and ε_r is the relative permittivity of the material.

In the case where $V_0 \ll V_{bias}$ and there is a highly asymmetric doping, i.e. $N_A \ll N_D$ or $N_D \ll N_A$, then W is given by:

[2] Also known as the space charge region.

$$W \approx \sqrt{\frac{2\varepsilon_r\varepsilon_0}{q}\frac{1}{N_{sub}} \cdot |V_{bias}|}, \tag{2.15}$$

where N_{sub} is the doping concentration of the weakly doped side, typically the substrate. The resistivity, described by Eq. 2.10, is therefore modified to:

$$\rho = \frac{1}{q\mu N_{sub}}, \tag{2.16}$$

where μ is the majority charge carrier mobility. Thus, W, can be expressed as:

$$W = \sqrt{2\varepsilon_r\varepsilon_0\mu\rho|V_{bias}|}. \tag{2.17}$$

Capacitance

A capacitance across the junction exists due to stored charge in the depletion region. The junction can be considered as a parallel plate capacitor where two conductive plates are separated by W. The capacitance associated with W is known as the depletion capacitance. Under a reverse bias, the depletion capacitance, C, of the pn junction can be approximated as a parallel plate capacitor with area A [19]:

$$C = A\frac{dQ}{dV} = \frac{A\varepsilon_r\varepsilon_0}{W} \approx A\sqrt{\frac{\varepsilon_r\varepsilon_0 qN}{2V_{bias}}}. \tag{2.18}$$

The depletion capacitance is inversely proportional to $\sqrt{V_{bias}}$.

Leakage current

Leakage current[3] refers to the steady current in a reverse biased diode without the presence of radiation. Leakage current is dominated by thermal generation of charge carriers. Leakage current per unit area, J_{vol}, has a strong temperature dependence [19]:

$$J_{vol} \propto T^2 e^{-E_g(T)/2kT}, \tag{2.19}$$

where E_g is the energy gap between the lower edge of the conduction band and the upper edge of the valence band. The temperature dependence of J_{vol} is such that J_{vol} doubles with every increase in temperature of about 8 K.

[3] Also know as dark current, reverse saturation current or volume generation current.

2.3 Front-end readout electronics

The front-end readout electronics are located on or close to the silicon active area. The analogue signal generated by the charged particle in the silicon sensor is fed into the front-end readout electronics and transformed into a signal that can be shipped off the detector. The basic functions of the front-end electronics are to amplify and shape the signal, digitise the signal, store the signal in a buffer and send the signal to the off-detector electronics. The front-end electronics can be split into two categories: the active area electronics and the periphery circuitry. The active area electronics consist of identical electronics present on each element in the sensor matrix, for instance each pixel in a silicon pixel detector. The periphery circuitry consists of the electronics common to some subgroup of the matrix, for instance a particular readout region, and host the electronics responsible for data processing and global signal propagation.

2.3.1 Active area electronics

The active area electronics can be categorised into analogue and digital circuitry. The analogue circuitry includes a charge-sensitive preamplifier, a pulse shaper, a discriminator and a test charge injector. The digital circuitry includes multiple event buffers.

Charge-sensitive preamplifier

The charge-sensitive preamplifier is essentially a charge-to-voltage converter. Charge liberated in the active area of the silicon sensor is amplified to provide a proportional voltage on the preamplifier output. Figure 2.4 shows the principle of a charge-sensitive preamplifier. It comprises an inverting amplifier with gain -A and a feedback capacitor C_f [20]. The voltage of the amplifier output is $-Av_i$, where v_i is the voltage of the amplifier input. In the case where the amplifier has an infinite input impedance, and therefore draws no current, then all current flows through C_f. The voltage difference across C_f is given by $v_f = (A+1)v_i$. The charge deposited on C_f is therefore $Q_f = C_f v_f = C_f(A+1)v_i$. Due to the infinite impedance of the amplifier input, $Q_f = Q_i$, so the dynamic capacitance of the amplifier input is given by:

$$C_i = \frac{Q_i}{v_i} = C_f(A+1). \tag{2.20}$$

The voltage output of the amplifier is $v_o = -AQ_i/C_i$, so the voltage output per unit input charge is given by:

$$\frac{dv_o}{dQ_i} = -\frac{A}{C_i} = \frac{A}{(A+1)C_f}. \tag{2.21}$$

For a gain $|A| \gg 1$:

$$A_Q = \frac{dv_o}{dQ_i} \approx -\frac{1}{C_f}. \tag{2.22}$$

The charge gain is determined by the feedback capacitor, an easily controlled component. The signal charge Q_s is split distributed amongst the sensor capacitance C_d and the dynamic input capacitance C_i. Hence, the ratio of measured charge to input charge is:

$$\frac{Q_i}{Q_s} = \frac{Q_i}{Q_d + Q_i} = \frac{C_i}{C_d + C_i} = \frac{1}{1 + \frac{C_d}{C_i}}, \tag{2.23}$$

resulting in the conclusion that the input capacitance must be large in comparison to the sensor capacitance in order to keep the ratio of measured charge to input charge close to unity.

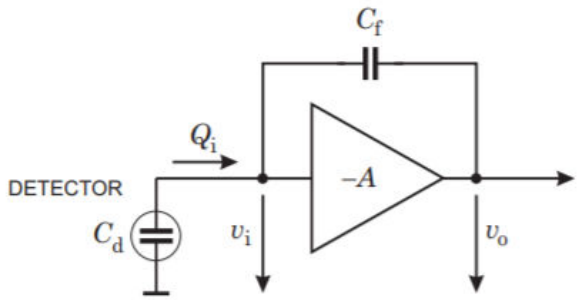

Figure 2.4: Principle of a charge-sensitive preamplifier. Taken from [22].

The power consumption of the charge-sensitive preamplifier must be kept low in order to reduce heat dissipation into the active area of the sensor, increasing the likelihood of thermal excitations of charge carriers.

Pulse shaping

Pulse shaping is necessary to limit the bandwidth of the signal to match the measurement time and to constrain the pulse width such that consecutive pulses do not overlap, i.e. to reduce pile-up. Pulse shaping is achieved through the use of high and low frequency band pass filters [23]. A high pass filter is applied first, resulting in a limitation on the pulse width. A low pass filter is subsequently applied, resulting in an increase in rise time. The effect is a higher rise time for a given pulse width. The bandwidth limitation results in a decrease of high and low frequency noise components such as sensor leakage current.

Discriminator

The amplitude of the amplified and shaped signal is compared to a preset threshold value with a discriminator circuit. If the signal is above the threshold value then the signal is classified as a hit. The threshold value must be high enough that the fake-hit rate of the detector is minimal but low enough that the detection efficiency is high. The threshold may need to be fine tuned across the detector due to variations in power supply and preamplifier gain. It is common that thresholds can be tuned locally across the detector, where the size of the area covered by the local tune depends on the detector technology utilised. For instance, a silicon pixel detector segmented into chips with individual power supplies would need local threshold tuning per chip. The response time of the discriminator is critical in high rate detection environments, such as the Large Hadron Collider (LHC).

Test charge injection

The controlled injection of test charges into the input of the preamplifier is crucial to verify the functionality of the active area circuitry. This is achieved through the use of a well calibrated injection capacitor with known voltage steps. The injection capacitor can also be used to tune the local thresholds for the discriminator circuitry. A test charge of known amplitude is injected into the input of the preamplifier and the threshold is tuned such that threshold uniformity is achieved across the detector.

Multi-event buffers

Multi-Event Buffer (MEB) are usually present in the active area electronics. Their role is to store a hit until it can be read out. In a high rate environment, multiple hits may occur within one read-out clock cycle. In this case, MEBs are necessary to prevent the loss of hit information. The depth of the buffer needs to be chosen based on the expected occupancy and read-out frequency. The readout of the MEBs can be completed in a different order depending on the desires of the user. Older events could be read out first, in a First In First Out (FIFO) configuration, or new events could be read out first, in a Last In Last Out (LILO) configuration.

2.3.2 Periphery circuitry

The main function of the periphery circuitry is managing the readout of data from a subgroup of active area elements. The periphery circuitry is also responsible for global propagation of signals to the active area elements under its control. Functions can include control of the injection capacitors, masking of active area elements and trigger distribution.

Readout from active area

A role of the periphery circuitry is to group together data from the active area elements within its control to forward to the off-detector electronics. This has the effect of reducing the number of data lines needed to ship data off-detector. Data will be grouped together with the address of the active area element. The readout of the active area electronics has some order, depending on the detector goals. For instance, reading out only elements that have detected a hit, a technique called zero-suppression, can be implemented to increase readout rate.

2.4 Overview of silicon tracking detectors

Silicon has a number of properties which make it a suitable material choice for a charged particle detector. Its small band gap leads to a large number of free charge carriers for relatively little energy loss of the charged particle (on average one electron-hole pair is liberated for each 3.6 eV deposited by a ionising particle [19]). The high carrier mobility of silicon allows the signal to be quickly extracted. Silicon

is mechanically stable which allows for complex structures to be manufactured, something which is aided by its extensive use in the electronics industry, meaning mature manufacturing technology can be exploited for use in high-energy physics.

Silicon detectors rely on the implementation of a pn junction in different geometries. The basic principles of a silicon detector can be described by a p-in-n sensor, where a p-type implant sits on the top of an n-type substrate. A depletion region is generated outwardly from the p-type implant. A charged particle traversing the depletion region ionises electron-hole pairs. Electrons move towards the p-type electrode which acts as a collection diode. The collected charge is amplified, the signal shaped and compared to a user defined threshold, before being sent off-detector by the chip periphery. A historical overview of silicon tracking detectors can be found in [24].

Silicon Strip Detectors (SSDs) were the first silicon detectors capable of measuring the position of a charged particle. An n-type substrate is implanted with numerous p-type strips with a typical pitch of $O(100\,\mu\text{m})$. The p-type strips can be placed on opposite sides of the n-type substrate, arranged in a grid, either perpendicular to one another, or with some stereo-angle. This approach allows two-dimensional position reconstruction with a cost-effective production. The main drawback is the production of ghost-hits when multiple hits are generated, making SSDs unsuitable for high event rates. Figure 2.5 illustrates how ghost hits arise. Generally n hits can generate up to n^2 crossing points, meaning up to $n^2 - n$ ghost hits.

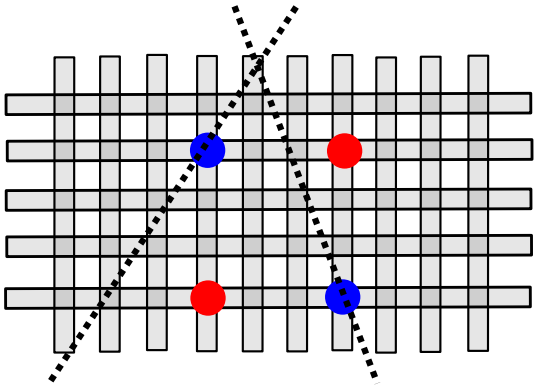

Figure 2.5: Illustration of ghost hits in silicon strip detectors. The blue circles represent genuine hits, red circles represent ghost hits and the black dotted lines represent the particle tracks.

Silicon Drift Detectors (SDDs) have a different principle of operation. Figure 2.6 shows an illustration of the cross section of a silicon drift detector. The cross

section is repeated with multiple n anodes and p strips. P-type implants on the top and bottom of an n-type substrate create a depletion region. Charge ionised in the depletion region drifts to the n-type diode where it is processed. The position of the crossing particle is determined in two dimensions. The drift time is used to determine the position in x, whilst the centroid of the Gaussian distribution of charge in the anodes is used to determine the position in z.

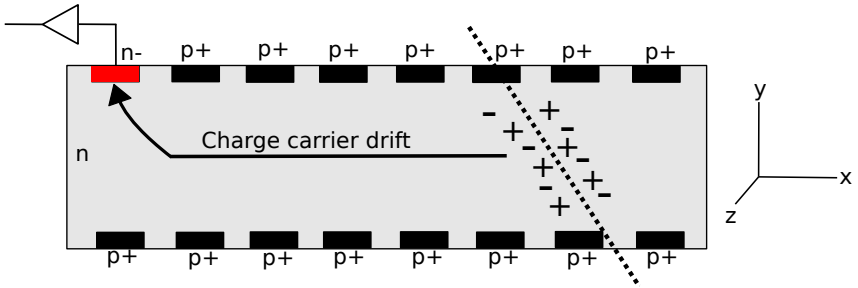

Figure 2.6: Illustration of the cross section of a silicon drift detector. A depletion region is created by p-type implants in a n-type substrate. Charge ionised by a traversing particle drifts to an n-type diode where it is collected.

2.4.1 Silicon pixel detectors

Silicon Pixel Detectors (SPDs) segment the sensor in two dimensions, creating a two dimensional matrix of square or rectangular sensing elements, called pixels. In high-energy physics, the pitch of each pixel tends to be small, $O(10-100\,\mu\text{m})$. SPDs emerged in high energy physics when vertex detectors needed to measure short lived particles in high multiplicity collisions. The high spatial resolution offered by SPDs means they can identify the secondary decay vertex of short lived particles, and the high granularity of SPDs allows them to cope with high particle multiplicities.

The readout circuitry is segmented in a similar fashion to the pixels, where each pixel has its own readout circuitry. The advantage to this approach is that ghost hits are not present, unlike strip detectors. The front-end active area readout circuitry is implemented in one of two ways. Traditionally, a separate silicon chip would contain the readout circuitry, and be bump-bonded to the sensor chip. Sensors with this design are known as hybrid pixel detectors. More recently, readout circuitry has been included in the sensor chip. Sensors with this design are known as Monolithic Active Pixel Sensors (MAPS).

Each pixel in an SPD has a very low capacitance, $O(0.1\,\mathrm{pF})$, due to the small pixel area extended over a thin, $O(100\,\mu\mathrm{m})$, layer of silicon. The low sensor capacitance, C_d, is a key benefit to using SPDs, since it permits fast signal shaping with a low noise. Another advantage to SPDs is the high radiation tolerance afforded by the small size of the pixels. Due to the volume dependence of leakage current, small pixels have small leakage current. Irradiation of silicon with $10^{15}\,\mathrm{cm}^{-2}$ particles increases the leakage current density to $\approx 30\,\mu\mathrm{Acm}^{-2}$ [19]. The small active area of each pixel leads to a smaller increase in leakage current per pixel compared to the larger active area of other technologies, such as SSDs.

Hybrid pixel sensors

Hybrid pixel detectors are produced in a similar way to SSDs. However, the p-type strips implanted in the n-type substrate are further subdivided across the length of each strip, forming pixels. Figure 2.7 shows a schematic of a hybrid pixel sensor. The pixel implants are analogous to the p-type strips in an SSD. Each pixel implant is bump-bonded to its respective active area readout electronics, the front-end channel. The size of the front end channel in the electronics chip must be identical to the size of the pixel implant in the sensor.

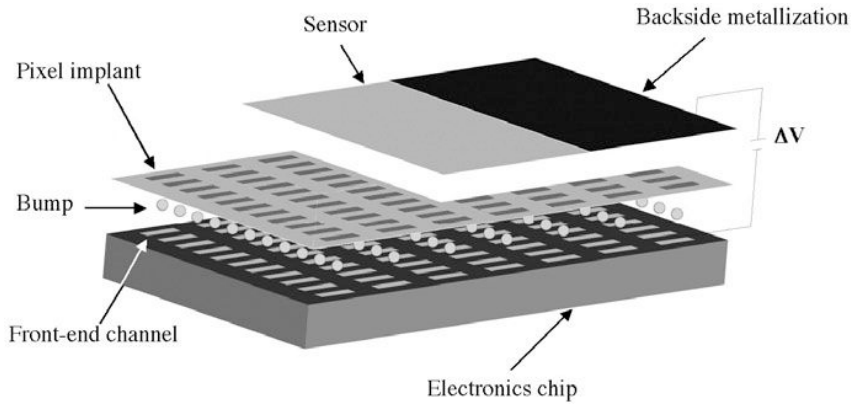

Figure 2.7: A schematic of a hybrid pixel detector. The sensor is bump-bonded to the readout chip. Taken from [19].

Monolithic Active Pixel Sensors

MAPS incorporate the electronics in the readout chip into the same silicon wafer as the sensors. This eliminates the need to produce many high density interconnections

between the sensor and readout chips, reducing the minimum pixel size, material budget and production cost. The capacitance of each pixel is also reduced slightly, increasing the signal to noise ratio. The improved granularity and reduced material budget contribute to improved vertexing and tracking capabilities. The improved material budget also presents a limitation of MAPS. Approximately 80 electron-hole pairs are liberated per µm in silicon, meaning thinner sensors generate a lower current, limiting the Particle IDentification (PID) capabilities. However, in the case of the ITS2, no PID capabilities are needed.

MAPS are typically organised into three layers: a p++ substrate, a p- epitaxial layer and a number of implants known as wells. Figure 2.8 shows a schematic of a basic MAPS design. For the ALPIDE chip, the epitaxial layer is 25 μm thick.

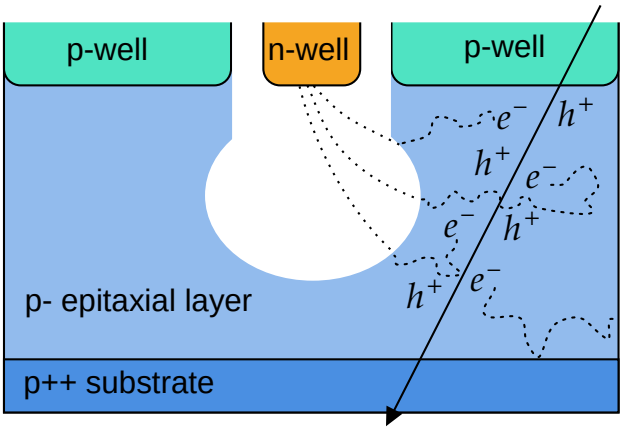

Figure 2.8: A schematic of a basic MAPS design. Charge ionised in the epitaxial layer diffuses to the depletion region (white). Readout circuitry sits on the p-wells (not shown). p-well/p- and p-/p++ junctions act as reflective barriers. Some electrons diffuse to neighbouring pixels, an effect known as clustering, or charge sharing.

The p++ substrate acts as a mechanical support, whilst the p-epitaxial layer is used as the sensitive volume. The wells can be n-type or p-type. N-type wells, or n-wells, act as collecting diodes for electrons, whilst p-type wells, or p-wells, host NMOS circuitry. Metal layers are bound to the wells, responsible for signal propagation off-pixel and forming parts of the in-pixel circuitry. The n-well diode is in contact with the p- epitaxial layer, creating a pn junction. A small voltage applied to the n-well collection diode increases the volume of the depletion region.

Electrons ionised in the epitaxial layer diffuse until they reach the depletion region where they are collected by the n-well diode, or until they recombine. Electrons ionised in the p-wells or substrate which diffuse into the epitaxial layer can also be collected by the n-well diode.

A key feature of the ITS2 is its use of a MAPS design, known as ALPIDE, throughout the detector. The ALPIDE chip is discussed in Sec. 4.4.3.

2.4.2 Noise in MAPS

Noise in MAPS can be split into two categories: Temporal Noise (TN) and Fixed Pattern Noise (FPN). FPN is the spatial variation of pixel output values under uniform stimuli. It is systematic and can be filtered offline. TN has various sources throughout the front-end electronics including:

- **Shot noise** is caused by fluctuations in DC current flow through the depletion region. The discrete nature of charge causes statistical fluctuations in the number of charge carriers which cross the edge of the depletion region. In pixel detectors, the main contribution to shot noise is leakage current.
- **Thermal noise**[4] occurs due to thermally induced movement of charge carriers. It increases with temperature.
- **Flicker noise**[5] affects MOSFETS and is caused by charge carriers being randomly trapped and released in the transistor channel by defects and contaminants.

In MAPs, shot noise increases linearly with leakage current, whilst thermal and flicker noise increase quadratically with the sensor capacitance. TN noise causes fluctuations in the output signal of the amplifier. Sometimes these fluctuations are great enough that the pixel charge threshold value is overcome, resulting in an identified hit by the discriminator. This leads to the detection of a noise hit or, as will be refereed to throughout this thesis, a fake-hit.

2.5 Accessing the quality of tracking detectors

This section will describe concepts necessary to access the quality of a tracking detector.

2.5.1 Single point resolution

A tracking detector of course needs to measure the position of a track. The position resolution depends on the pitch, p, between sensing nodes. For a binary readout, the position resolution in one dimension, σ_x can be expressed as:

[4] Also known as Johnson or Nyquist noise.
[5] Also known as 1/f noise.

$$\sigma_x = \frac{p}{\sqrt{12}}. \tag{2.24}$$

Equation 2.24 is valid when the readout is binary, the hit position is taken to be the centre of the pixel and there is no charge sharing. Charge sharing can improve the position resolution, since the position can be known with a higher precision by calculating the centre of mass of the cluster.

2.5.2 Transverse momentum resolution

Figure 2.9 shows the geometry of a particle traversing a region with a magnetic field, B. The normalised transverse momentum resolution depends on p_T , L, B and the resolution of the sagitta, σ_s (the derivation can be found in Appendix A):

$$\frac{\sigma_{p_T}}{p_T} \propto \frac{p_T}{BL^2}\sigma_s. \tag{2.25}$$

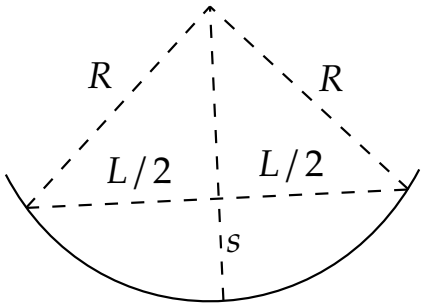

Figure 2.9: Geometry of a particle moving in a magnetic field. The solid line represents the particle track.

The transverse momentum resolution improves linearly with the magnetic field strength and sagitta resolution, whilst it increases quadratically with the arc length. As p_T increases, transverse momentum resolution decreases linearly.

2.5.3 Pointing resolution

Figure 2.10 shows the decay topology of a D^0 ($c\bar{u}$) into a kaon and a pion. The secondary decay vertex is the point at which the kaon and pion tracks are closest to each other. The D^0 flight line then must be the track directly connecting the secondary and primary vertices. The pointing angle is the angle between the reconstructed momentum of D^0 and the D^0 flight line.

Having a good pointing resolution means secondary vertices can be accurately, if at all, distinguished from the interaction point. This is especially important for low transverse momentum, p_T, short lived particles like the heavy flavour hadron, Λ_c (udc).

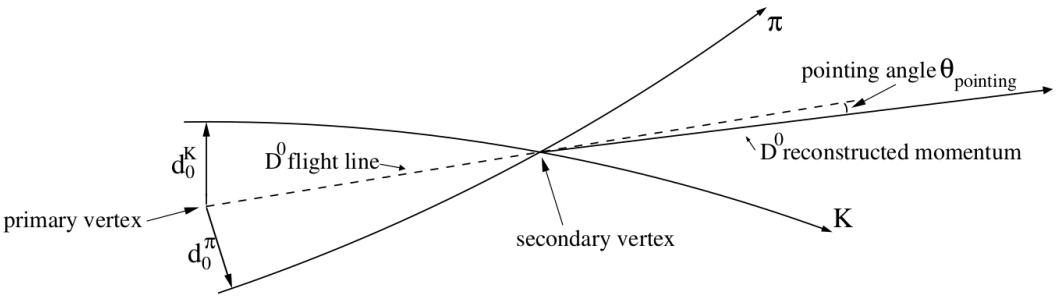

Figure 2.10: Decay topology of the D^0 heavy flavour meson into a pion and a kaon. d_0^K and d_0^π denote the kaon and pion impact parameters respectively. Image from [25].

The pointing resolution depends on the following three variables: the intrinsic spatial resolution, the radius of the innermost detector layer and the material budget.

2.5.4 Impact parameter

The impact parameter is the distance of closest approach from the retraced track and interaction point (d_0^K and d_0^π in Figure 2.10). The impact parameter resolution gives a measure of how accurately tracks are reconstructed. It is affected by the number and spacing of detector layers, the material budget and the pointing resolution.

2.5.5 Tracking efficiency

The tracking efficiency is the efficiency with which a real track can be reconstructed from some signal produced by the tracker. The tracking efficiency is a function of both the hardware of the tracker and the track reconstruction software.

Concerning the hardware of the tracker, the tracking efficiency is dependent on the threshold value and the noise of the detector. The threshold value specifies the amount of charge which must be deposited in the active volume of a pixel for that pixel to readout a hit. If the threshold value is too low, then noise can be more easily classified as a hit, reducing the tracking efficiency by contaminating the data with fake hits.

The tracking efficiency of high p_T events is larger than for low p_T events. Low p_T particles are affected more strongly by energy loss mechanisms as they traverse through the active volume of the detector. For this reason, the chance of a real track being reconstructed accurately is lower for a low p_T particle.

Chapter 3:
Quark Gluon Plasma

This chapter will give a brief overview of the physics relevant to the ALICE experiment, written for a non-expert in QGP physics. A selection of observables of the Quark-Gluon Plasma (QGP) will be discussed with a brief outline of the most recent results. Finally, the characteristics of an Inner Tracking System (ITS) optimal for heavy-ion collisions will be discussed.

3.1 The Standard Model

The current understanding of particle physics is anchored in the Standard Model (SM), which describes the fundamental particles and the interaction between them. The fundamental particles are split into fermions (with spin $\frac{1}{2}$) and bosons (with integer spin).

Bosons have either spin 1 (gauge bosons) or spin 0 (scalar bosons). Gauge bosons are force-carriers, mediating the interaction of particles via three of the four fundamental forces of nature (gravity proves so far to be elusive to the SM). The massless photon mediates the electromagnetic interaction, whilst the W and Z bosons (W^-, W^+, Z^0) mediate the weak interaction and gluons mediate the strong

Header photograph shows the bottom half of the Outer Barrel being lowered 56 m by crane to the bottom of the ALICE cavern. © CERN. Taken from [26].

interaction. A particle will interact via the strong force if it has colour charge, with the electromagnetic force if it has electric charge and with the weak force if it has weak isospin. There is only one spin 0 boson in the SM, the Higgs boson[1]. This is a massive particle responsible for the bare mass generation of all other fundamental massive particles.

The 12 fermions are further split into two categories, quarks and leptons. There are six quarks in the SM, which interact via the strong and electroweak forces. The remaining six fermions are all leptons. These are the electron, muon and tau and the neutrinos, the electron neutrino, the muon neutrino and the tau neutrino. The electron, muon and tau have electroweak charge and the neutrinos have only weak isospin.

3.1.1 QCD

The advent of bubble chambers and spark chambers in the 1950s led to the discovery of a plethora of new hadrons, in a period of particle physics known as the 'particle zoo era'. Independently in 1964, Gell-Mann and Zweig culled the population of the particle zoo by proposing that hadrons were not fundamental particles, but were made up two or three spin $\frac{1}{2}$ particles with fractional electric charge of three different types named up, down and strange [27, 28]. These particles were collectively named aces by Zweig but the name of quarks given by Gell-Mann stuck. From 1967 to 1973, deep-inelastic electron scattering experiments at SLAC confirmed the existence of point-like particles within a proton [29, 30]. The charm quark was predicted in 1970 [31] and discovered four years later [32, 33]. The beauty and top quarks were discovered in 1977 [34] and 1995 [35] respectively .

The part of the SM which describes the strong force is called Quantum Chromodynamics (QCD). The gauge bosons of QCD are called the gluons. Unlike the photon (the gauge boson in Quantum Electro-Dynamics (QED)) the gluons are self-coupling, carrying colour and anti-colour charge. Figure 3.1 shows the self-coupling of gluons possible in QCD, as well as a quark-gluon interaction.

[1]The ten year anniversary of the Higgs discovery by ATLAS and CMS is on the 4th July this year (2022).

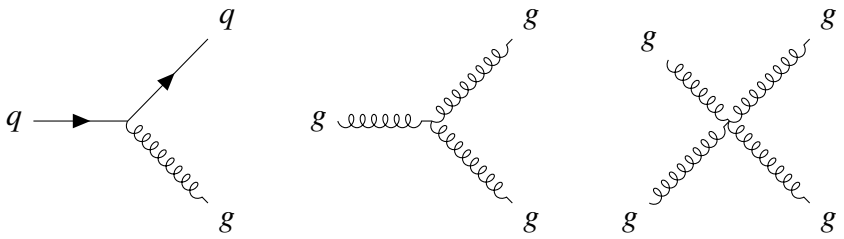

(a) qg interaction. **(b)** gg interaction. **(c)** gg interaction.

Figure 3.1: Feynman diagrams of the interaction modes within QCD.

Colour charge

Quarks have a quantum number called colour charge, originally hypothesised in 1964 by Greenberg [36] to explain why the Pauli exclusion principle is not violated in certain observed hadrons, such as Δ^{++}, which is composed of three up quarks which appear to have identical quantum numbers. The naming convention of the quantum number of colour dictates that colour can be red, green or blue (or anti-red, anti-green or anti-blue). Hadrons have a net colour charge of 0. Baryons are composed of three quarks and mesons are formed by one quark and and anti-quark, for example the D meson which is made up of a c (anti)quark and a different (anti)quark. In the case of a baryon, a net colour charge of 0 is achieved with the combinations of red, green and blue quarks. In the case of a meson, a net colour charge of 0 is achieved with a single quark of any colour and its corresponding anti-quark (with the corresponding anti-colour).

Running coupling

The running coupling in QED and QCD refers to the strength of the interaction as a function of the energy scale of the interaction.

In QED, the photon is not self-coupling. An electron in QED is consistently surrounded by virtual electron anti-electron pairs. As such the charge felt by a test charge depends on its distance, r, from the electron. At a larger r, more electron anti-electron pairs exist in the space between the electron and the test charge, screening the charge of the electron. As such, the running coupling in QED, α, decreases with increasing r.

Similarly in QCD, quark anti-quark pairs screen colour charge. However, the gluon is self-coupling, meaning gluon loops contribute to the overall colour charge, having an anti-screening effect. The interplay of screening and anti-screening ef-

fects has a distinct effect on the running coupling in QCD, α_s. Instead of decreasing in strength with an increase in r like the case of α in QED, α_s increases with increasing r. An increase in r corresponds to a decrease in energy scale, Q. The relationship between $\alpha_s(Q^2)$ as a function of Q is shown in Fig. 3.2.

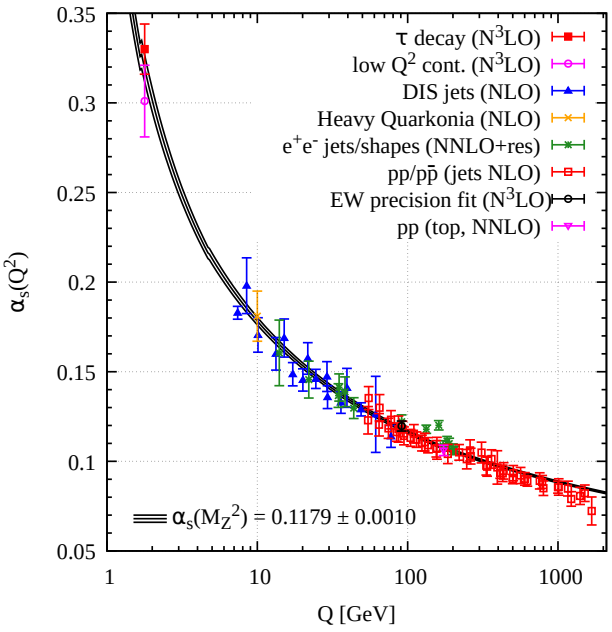

Figure 3.2: Summary of the measurements as of 2021 of the coupling constant, $\alpha_s(Q^2)$ as a function of energy scale Q. The brackets indicate the order of perturbation theory used, i.e. (NLO) is next-to-leading-order. Taken from [12].

As r decreases (a high energy probe), α_s decreases, whilst as r increases (a low energy probe) α_s increases. The behaviour at small r results in asymptotic freedom, whilst the behaviour at large r is called colour confinement. The discovery of asymptotic freedom by Gross, Politzer and Wilczek in 1973 [37, 38, 39, 40] won the Nobel prize in physics in 2004 [41].

Chiral symmetry breaking

The 'bare masses' of quarks are their masses in the absence of a chromodynamic field. The kinematic effects of the coupling of quarks to the electroweak (electromagnetic plus weak-β decay fields) and gravitational fields determines their response to the application of those fields. In the case where no chromodynamic field is present, the coupling between the matter fermions (leptons or quarks) and the field-bosons (photon, $W^{\pm}$ and Z^0 bosons) are sensitive to the spin-projection, or helicity, of the matter fermions. In these circumstances, the electroweak field is said to be 'maximally parity-violating'.

Fermion (e^-, μ^-, τ^-, their respective neutrinos and the quarks) coupling is Left-Handed (LH) (helicity -1) preferred, whilst anti-fermion (e^+, μ^+, τ^+, their

respective anti-neutrinos and the anti-quarks) coupling is Right-Handed (RH) (helicity +1) preferred. This behaviour of opposite handedness of fermions and anti-fermions is known as 'chiral symmetry'. The factor by which the strength of RH fermion couplings are reduced with respect to the strength of LH fermion couplings is proportional to the inverse of the mass squared of the fermion. Therefore, neutrinos and anti-neutrinos are almost always LH and RH respectively, electrons and anti-electrons are mostly LH and RH respectively, muons less so, and taus the least of all charged leptons.

Quarks are colour confined so are always subjected to some chromodynamic effects, which decline as asymptotic freedom is approached. The bare masses of quarks are therefore taken to be those corresponding to their electroweak interactions, including their handedness. The bare mass of up and down quarks are $2.16^{+0.49}_{-0.26}\,\mathrm{MeV}$ and $4.67^{+0.48}_{-0.17}\,\mathrm{MeV}$ respectively [12]. When confined within a proton, each of the up and down quarks contribute roughly 300 MeV, approximately $1/3$ of the proton mass. How is this mass generated?

Massless fermions possess chiral symmetry but massive fermions do not. This symmetry, although approximate for quarks who have a small mass in the deconfined phase, must be broken as quarks gain their inertial mass. A sponge in a dry, gaseous medium of water vapour will remain reasonably dry. If the temperature is lowered and the water condenses, the sponge will absorb more water and become heavier. Mass is generated when the water changes phase and a symmetry is broken [42]. Similarly, in the deconfined phase, quarks have a low mass (analogous to unsaturated sponges). When approximate chiral symmetry is broken (analogous to the temperature lowering, changing the phase of water), the quark is 'dressed' with gluons, increasing its mass (analogous to liquid water saturating the sponge).

The mass increase of quarks in QCD is known as approximate chiral symmetry breaking [42]. In the approximately symmetrical phase, quarks move freely and once this approximate symmetry is broken, gluons surround the quarks, 'dressing' them and giving them inertial mass. The gluons then restrict the quarks to existing in colourless combinations, i.e. hadrons. The transition to an approximately chirally symmetric state of matter can be achieved experimentally, as detailed in Sec. 3.2.

3.2 Quark-Gluon Plasma

Above a critical energy density of $0.7\,\mathrm{GeV/fm^3}$ [43], the measured degrees of freedom in excited hadronic matter strongly suggests it is composed of quarks and gluons rather than nucleons. Colour confinement no longer occurs within the spatial confines of a hadron, but rather within the spatial confines of the medium. This state of matter is known as the Quark-Gluon Plasma and is thought to have existed microseconds after The Big Bang, during the Quark Epoch [44].

3.2.1 Formation of the QGP

The transition from hadronic to partonic matter above a critical energy density can be thought of in a similar way as the Mott transition[2]. Increasing the density of hadronic matter, say a box of mesons, above a threshold will result in colour charge within a particular meson being screened such that the two original quarks within that meson no longer feel an interaction between them. Within the QGP, each of the original quarks is immediately surrounded by many other quarks, all a similar distance away and all equally energetically favourable to bind. As such, a specific quark could move from one side of the box to another by successive bindings with different quarks [45]. This is known as the Flip-Flop Mechanism as quarks traverse the medium with successive 'flips' of partner with associated 'flops' on the parts of the partners [42].

The result is a state of matter where quarks are no longer colour confined to the spatial limits of a hadron. Rather they are confined to the spatial limits of the QGP.

3.2.2 QGP evolution in relativistic heavy-ion collisions

Figure 3.3 shows the light cone, illustrated in two dimensions, of two nuclei approaching each other and colliding. During the first moments of the collision, from $t = 0$ until $t = \tau_0$, the system consists of hard processes like quark-pair production, as well as secondary collisions between partons. At $t = \tau_0$, local thermal equilibrium is reached, forming a QGP. The strong internal pressure gradients generated by the geometry of the collision cause a rapid expansion of the QGP. The low vis-

[2]The Mott transition occurs when the density of atoms is high enough for screening effects to occur on a smaller scale than the bonding potential between the nucleus and atomic electrons, resulting in a transition from insulator to metallic material.

cosity of the medium causes the QGP to expand collectively, in a way described by relativistic hydrodynamics. This collective expansion is known as flow.

As the QGP expands it cools. Eventually, the temperature of the QGP reaches the critical temperature, T_c, below which the energy density is low enough for hadronisation to occur. The medium thus undergoes a phase transition, from the QGP to a hadron gas. The medium continues to expand and cool and at some point inelastic collisions no longer occur, leading to some fixed ratio of hadron species. The point at which this happens is named chemical freeze-out. Elastic collisions[3] of hadrons continue to occur until the energy density becomes low enough that they stop. This point is known as kinetic freeze-out and is the moment when the hadron momenta become fixed. Final state particles are then detected experimentally [46].

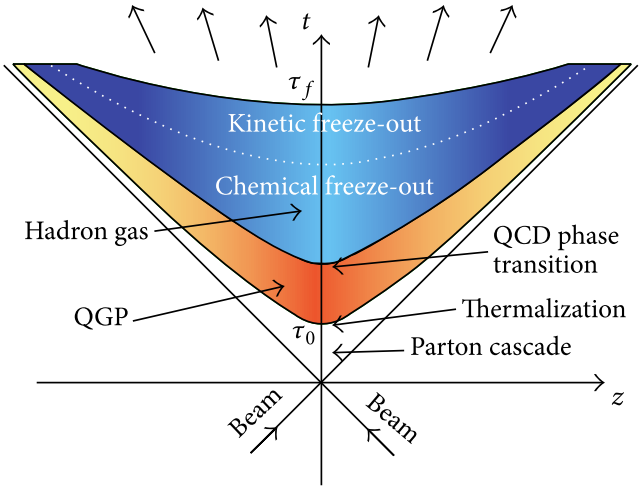

Figure 3.3: Illustration of the space-time evolution of a relativistic heavy-ion collision. See text. Taken from [47].

3.3 Selected QGP observables

In order to detect the creation of the QGP and measure its properties in the laboratory, experimental probes, or observables, of the QGP must be clearly defined. QGP observables can be split into two categories. Hard probes originate from initial nucleon-nucleon interactions. Some hard-probes interact with the strongly interacting medium, such as jets, which provides information of the medium. Other hard

[3] where the initial and final state particles are the same. In contrast to inelastic collisions where the initial and final state particles are different.

probes do not, such as prompt photon emission, and therefore provide information on the initial collision conditions. Soft probes originate from the medium itself and provide information on the bulk properties of the QGP. This section will outline a selection of these QGP observables and give a brief outline of current results.

3.3.1 Charged particle multiplicity

The rapidity and pseudo-rapidity distributions of primary charged particles are used to characterise the multiplicity. Pseudo-rapidity, η, is defined as [44]:

$$\eta = -\ln\left(\tan\left(\frac{\theta}{2}\right)\right), \tag{3.1}$$

where θ is the angle between the particle three momentum, $\bar{p}$, and the positive direction of the beam axis. As a function of $\bar{p}$:

$$\eta = \frac{1}{2}\ln\left(\frac{|\bar{p}| + p_z}{|\bar{p}| - p_z}\right), \tag{3.2}$$

where p_z is the momentum component along the beam axis. At high momentum where the particle rest mass is negligible, the pseudo-rapidity converges to rapidity, y, defined as:

$$y = \frac{1}{2}\ln\left(\frac{E + p_z}{E - p_z}\right), \tag{3.3}$$

where E is the total energy of the particle. Since E requires the knowledge of the particle's rest mass, pseudo-rapidity is more easily measured because it does not require particle identification.

Figure 3.4 shows the charged particle pseudo-rapidity density measured as a function of $\sqrt{s_{NN}}$ for central Pb-Pb and Au-Au collisions in comparison to pp, p$\bar{\text{p}}$, p-A and d-A collisions. The charged particle density, $\langle dN_{ch}/d\eta \rangle$, is divided by the average number of participating nucleon pairs in order to directly compare collisions systems with a different number of participating nucleons. The increase in $\langle dN_{ch}/d\eta \rangle$ with increasing $\sqrt{s_{NN}}$ is greater with A-A collisions in comparison to pp collisions. p-A and d-A collision data lie on the pp curve. This indicates the behaviour seen for A-A collisions cannot be due only to the larger number of participants.

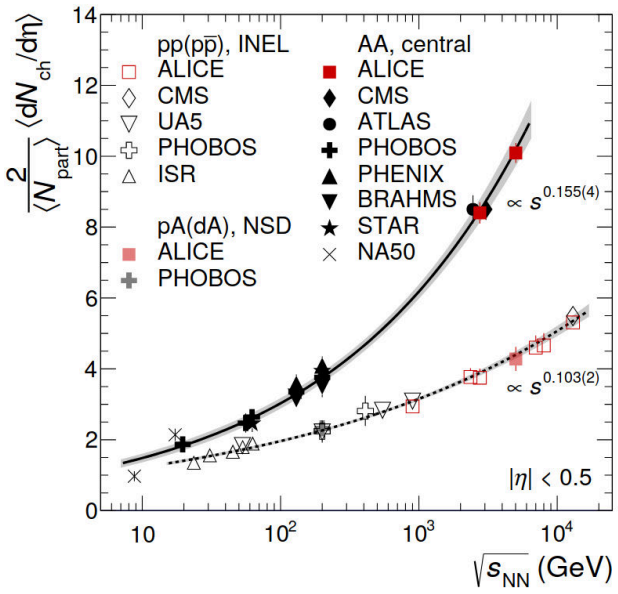

Figure 3.4: Values of $\frac{2}{\langle N_{part} \rangle} \langle dN_{ch}/d\eta \rangle$ as a function of $\sqrt{s_{\mathrm{NN}}}$ for central Pb-Pb and Au-Au collisions. Inelastic (INEL) pp and $\mathrm{p\bar{p}}$ collisions, as well as non-single diffractive (NSD) p-A and d-A collisions are shown. Taken from [48].

3.3.2 Direct photons

Direct photons, defined as photons that do not originate from a hadron decay, are produced in all stages of the QGP evolution. They are not strongly interacting, so traverse the medium unaffected [49]. Direct photons provide information on the conditions at the time the photon was produced. Prompt direct photons, produced from hard primary collisions, provide information on parton distributions in nuclei. Thermal direct photons, produced throughout the lifetime of the QGP, provide information on the temperature and space-time evolution of the medium. Prompt direct photons dominate at high-$\mathrm{p_T}$ ($p_T > 5$ GeV/c), whilst thermal direct photons dominate at low-$\mathrm{p_T}$ ($p_T < 4$ GeV/c) [50].

ALICE measured the effective temperature of the QGP to be $T = 297 \pm 12(\mathrm{stat}) \pm 41(\mathrm{sys})$ MeV [51], which is greater than the expected deconfinement value. The value was gained by studying the direct photon spectra for central (0–20% centrality) collisions, in the low-$\mathrm{p_T}$ range of $0.9 < p_T < 2.1$ GeV/c at a collision energy of $\sqrt{s} = 5$ GeV. The effective temperature reflects the temperature of the QGP averaged over its lifetime.

3.3.3 Collective flow

The QGP collectively expands and cools over time, eventually becoming cold and dilute enough to hadronise. The collective expansion of the medium is called col-

lective flow. Collective flow is an indication for a strongly interacting medium, indicating the presence of multiple interactions in the medium between its constituents, with collective flow increasing as the system approaches thermal equilibrium. Collective flow is an observable capable of providing information on the equation of state and transport properties of the system.

Heavy-ion collisions are characterised by their centrality, defined by the impact parameter. The impact parameter, b, is defined as the displacement from a central nuclei collision in the direction of a certain plane. In the case of $b = 0$, the heavy-ion collision is perfectly central. Radial flow occurs in central heavy-ion collisions, characterised as an isotropic azimuthal distribution in particle production.

If $b \neq 0$, the collision consists of participant nucleons, nucleons that collide, and spectator nucleons, nucleons that do not collide. An illustration of a non-central heavy-ion collision is shown in Fig. 3.5, where the blue regions depict spectator nucleons and the red, almond-shaped region depicts participant nucleons. The impact parameter is defined in the x direction.

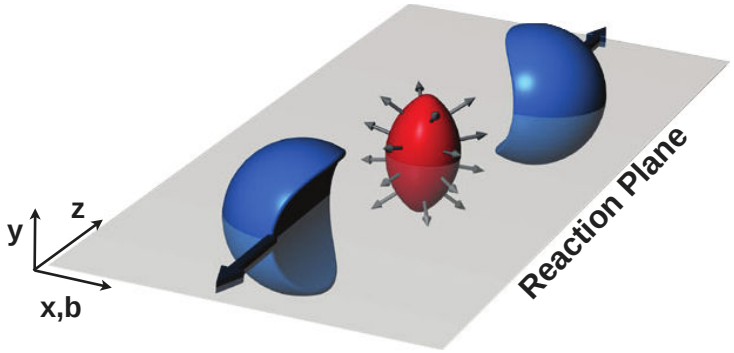

Figure 3.5: Two nuclei collide asymmetrically, forming an almond shaped region of participants (red), whilst the spectators do not interact. Elliptic flow occurs as a result. The reaction plane is defined as the spatial anisotropy with respect to the *xy* plane. The impact parameter, b, is defined as the displacement from a central nuclei collision in the x direction. Taken from [5].

Asymmetries of the system in non-central nuclei collisions lead to an azimuthal anisotropy in particle production, a flow phenomenon known as anisotropic flow. Anisotropic flow occurs due to azimuthally asymmetric pressure gradients, causing a correlation between the azimuthal angle of the produced particles and the reaction plane. Anisotropic flow is a characteristic trait of collective behaviour.

Anisotropic flow can be broken down into two flow phenomena, directed flow and elliptic flow. Directed flow occurs when there is a difference in the particle

production parallel and anti-parallel to the impact parameter. In this case, there is a preferred direction in particle emission.

Elliptic flow occurs if there is a difference in particle production parallel and perpendicular to the impact parameter. Referring to Fig. 3.5, the number of particles produced in the xz plane is greater than the number of particles produced in the y direction. The initial transverse energy profile is elliptical in shape. Due to the larger pressure gradients along the shorter axis of the ellipse (in the reaction plane), the acceleration is higher. This leads to the elliptical transverse energy profile evolving over time, tending to a circular distribution [5].

Elliptic flow has been measured by a number of experiments, including STAR at RHIC [52] and ALICE at the Large Hadron Collider (LHC) [53]. Figure 3.6 shows the elliptic flow, v_2, as a function of collision energy measured by different experiments. v_2 is the second coefficient in the Fourier decomposition of the azimuthal angle distribution of particles. It is a measure of the strength of the elliptic flow. v_2 is positive in the case of an excess of particles in the xz plane with respect to the y direction. v_2 is negative when there is an excess of particles produced in the y direction compared to the xz plane.

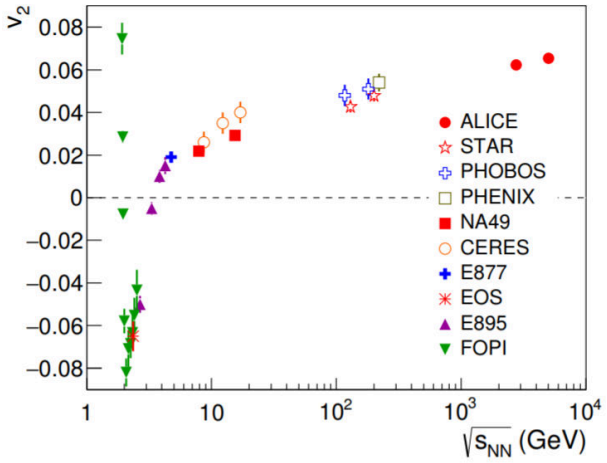

Figure 3.6: Integrated elliptic flow (v_2) for the 20-30% most central Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02\,\text{TeV}$ measured by ALICE, compared to elliptic flow measurements at lower energies with similar centralities taken by other experiments. Taken from [53]

3.3.4 Jet quenching

The nuclear modification factor is used to compare the change in particle yield over variances in transverse momenta seen in heavy-ion collisions with proton-proton collisions. The nuclear modification factor for heavy-ion collisions, R_{AA}, is defined as follows:

$$R_{AA} = \frac{dN_{AA}/dp_T}{\langle N_{coll} \rangle \, dN_{pp}/dp_T}, \tag{3.4}$$

where $N_{AA/pp}$ is the yield of the observables yield in AA or pp collisions respectively, p_T is transverse momenta and $\langle N_{coll} \rangle$ is the average number of binary nucleon-nucleon collisions for the centrality considered. $R_{AA} < 1$ indicates a suppression in the yield of the observable in AA collisions with respect to pp collisions, whilst $R_{AA} > 1$ signifies an enhancement.

Hard parton scattering occurring in the initial interaction of relativistic hadronic collisions produces cascades of hadrons, known as jets. Energy lost through gluon radiation and elastic scattering of the partons comprising the jets in the QGP is known as jet quenching. The amount of energy lost depends on the path length. Figure 3.7 illustrates the production of jets with and without the presence of the QGP. Jet quenching suppresses high-momentum particles produced in relativistic heavy-ion collisions therefore their R_{AA} is expected to be less than one.

Figure 3.8 shows R_{AA} of charged particles as a function of transverse momentum measured by ALICE for p-Pb collisions, Pb-Pb collisions at $0-5\%$ centrality and $70-80\%$ centrality. The energy density of p-Pb and p-p collision systems is not large enough for significant jet quenching effects to be expected. For $p_T > 2$ GeV/c, R_{pPb} is consistent with one, indicating that the suppression in R_{AA} for $p_T > 2$ GeV/c is not an initial-state effect, but instead a characteristic of the QGP created in heavy-ion collisions [54].

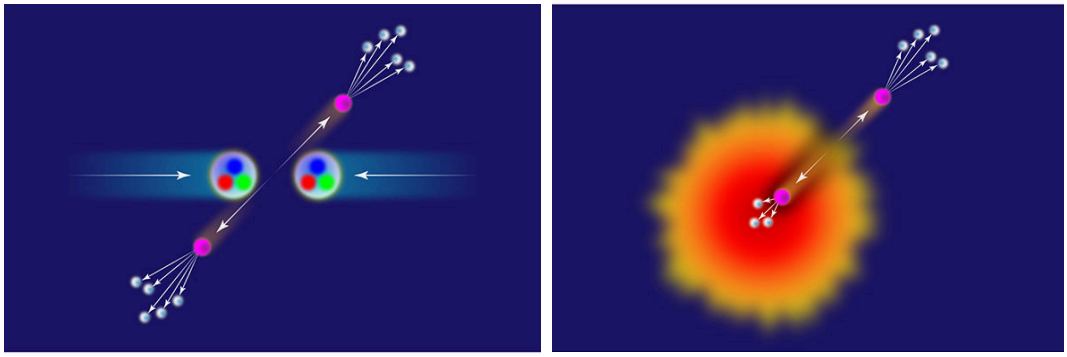

Figure 3.7: Jets can form in hadronic collisions. Partons (pink) hadronise forming directional jets of hadrons (white). In the absence of a QGP (left), momentum conservation leads to equal magnitude and opposite direction jets. In the presence of the QGP (right), jets lose energy to the medium, resulting in a leading jet and a quenched jet. Taken from [55].

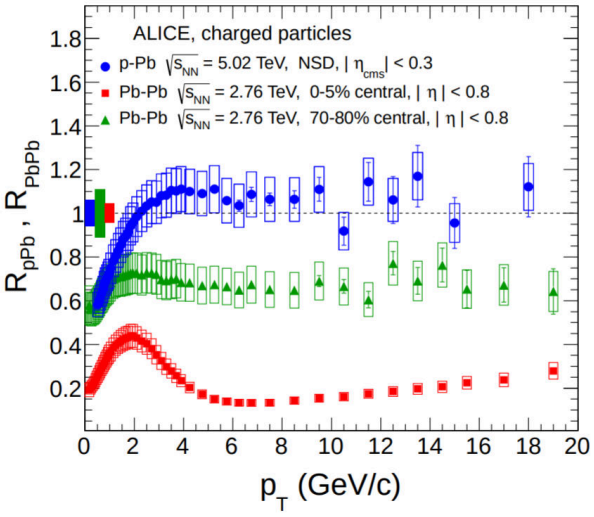

Figure 3.8: Nuclear modification factor of charged particles as a function of transverse momentum. Data shown are from p-Pb collisions at $\sqrt{s_{NN}} = 5.02\,\text{TeV}$, and Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76\,\text{TeV}$ for both $0\% - 5\%$ and $70\% - 80\%$ centrality. Taken from [54].

3.3.5 Heavy-flavour physics

Heavy-Flavour (HF) physics concerns the study of the properties of hadrons containing charm or beauty quarks, as well as using heavy-flavour hadrons to probe the collision system properties and development. The bare masses of charm and beauty quarks, 1.3 GeV and 4.2 GeV respectively [12], are significantly higher than the temperature at any stage of the heavy-ion collisions. Therefore, heavy-flavour quarks are mostly produced via hard scattering before the formation of the QGP [56]. Since the decay of charm and beauty quarks occurs on a timescale larger than the lifetime of the QGP, it can be assumed that the abundance of HF hadrons is constant from the initial collision to the hadronisation of the QGP. Light-flavour hadrons, however, can be created within the QGP so their abundance is not constant with time. The kinematic distribution of HF hadrons can reveal their interaction history, when thought of as Brownian particles within the QGP. Two main open questions regarding HF hadrons remain, concerning their thermalisation and hadronisation, as well as their in-medium energy loss and its mass dependence.

Hadronisation of heavy-flavour quarks and quarkonium recombination

A quarkonium is a flavourless meson comprised of a heavy quark and its own antiquark. Charmonium ($c\bar{c}$), bottomium ($b\bar{b}$) and toponium ($t\bar{t}$) are examples of quarkonia. The ground states of each of these quarkonia are known as the J/Ψ, Y' and θ mesons respectively. Due to the high mass of the top quark, the θ meson is not expected to be observable in nature, since the top quark decays before a

bound state can form. Hence, quarkonium generally refers only to charmonium and bottomonium.

Quarkonium states are bound by energies comparable to the mean energy of the QGP, $O(100)MeV$, therefore they are likely to dissociate. For a J/Ψ meson, the ground state of $c\bar{c}$, in the QGP, colour screening separates the c and $\bar{c}$. The QGP will subsequently cool towards the hadronisation temperature. The c (or $\bar{c}$) will have a higher probability of combining with a light quark to hadronise into a D (or $\overline{D}$) meson, than recombining with a c (or $\bar{c}$) to form a J/Ψ meson. Therefore the production of the J/Ψ meson will be suppressed, compared to pp collisions, by the presence of the QGP. This suppression of charmonium as a signature of the formation of the QGP was first proposed by Matsui and Satz [57].

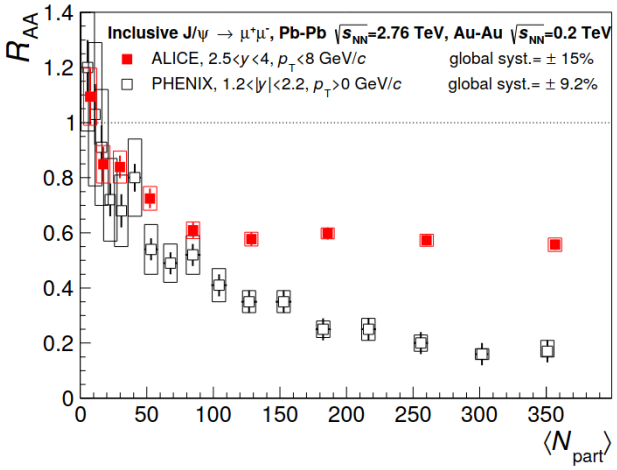

Figure 3.9: Inclusive J/Ψ R_{AA} as a function of the number of participant nucleons measured in Pb–Pb collisions with ALICE at $\sqrt{s_{NN}} = 2.76\,\mathrm{TeV}$, in comparison to Au–Au collisions measured with PHENIX at $\sqrt{s_{NN}} = 0.2\,\mathrm{TeV}$. Taken from [58].

The nuclear modification factor for J/Ψ as a function of the number of participant nucleons is shown in Fig. 3.9 for Pb-Pb collisions measured with ALICE at $\sqrt{s_{NN}} = 2.76\,\mathrm{TeV}$, in comparison to Au–Au collisions measured with PHENIX at $\sqrt{s_{NN}} = 0.2\,\mathrm{TeV}$. In both cases J/Ψ suppression is visible. However, a larger suppression effect is measured by PHENIX, at a lower centre-of-mass energy. This is due to quarkonium recombination[4], whereby statistical recombination of c and $\bar{c}$ from unrelated hard scatterings who occupy the same momentum and position space contribute to the charmonium yield [59]. The c and $\bar{c}$ production at LHC energies is expected to be a factor of ≈ 10 times larger than at RHIC energies, leading to a larger contribution to charmonium yield via quarkonium recombination.

If quark recombination plays a significant role in the hadronisation of charm quarks, an enhancement of the charmed baryon, Λ_c^+ (udc), relative to D^0 (cū) pro-

[4] Also known as coalescence.

duction is expected [60]. Within this narrative, up and down quarks occupying the same momentum and position space coalesce, forming a diquark (du). Considering the Λ_c^+ as a combination of a diquark and a charm quark, results in the conclusion that Λ_c^+ production would be enhanced.

Figure 3.10 shows the Λ_c^+/D^0 ratio measured by ALICE in central ($0-10\%$ centrality) and mid-central ($30-50\%$ centrality) Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV compared to pp collisions. A 3.7σ enhancement for Λ_c^+/D^0 in central Pb-Pb collisions with respect to pp collisions is observed in the p_T region of 4-8 GeV/s [61]. The Λ_c^+/D^0 ratio as a function of p_T can also be affected by flow, since baryons and mesons are affected differently by flow effects. It is currently unclear if the 3.7σ enhancement is due to quark recombination or flow effects.

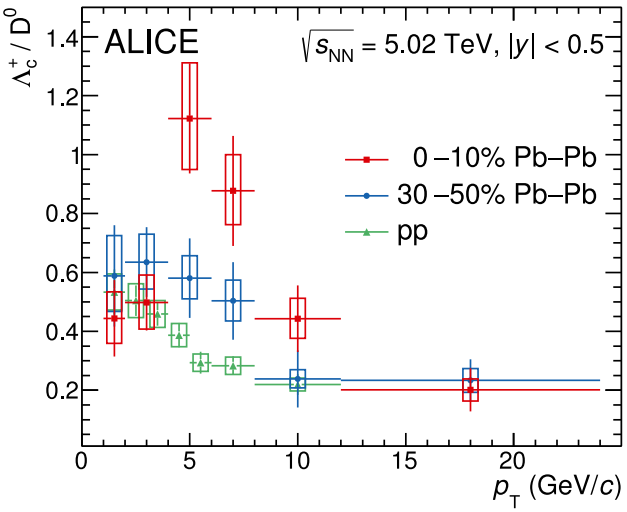

Figure 3.10: Λ_c^+/D^0 ratio measured by ALICE in $0-10\%$ and $30-50\%$ centrality Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV compared to pp collisions at the same energy. A 3.7σ enhancement for Λ_c^+/D^0 in central Pb-Pb collisions with respect to pp collisions is observed in the p_T region of 4-8 GeV/s. Taken from [61].

In-medium energy loss of heavy-flavour quarks

The energy loss of a HF quark depends on the quark momentum and the temperature and size of the medium. Energy loss occurs primarily through radiative or collisional energy loss. Collisional energy loss describes energy loss via elastic scattering with the constituents of the medium and is more dominant at low particle momentum. Radiative energy loss describes inelastic scattering with the medium constituents with the emission of a gluon or photon and dominates at high particle

momentum.

Phenomenologically, R_{AA} can be used to constrain the transport properties of the medium. Furthermore, the elliptic flow can be used to constrain the dependence of the path length on the in-medium energy loss. The momentum dependence of in-medium energy loss can be constrained by measuring the R_{AA} of the $\mathrm{p_T}$ distributions of the D and B meson over a wide momentum range. Figure 3.11 shows the R_{AA} of prompt D-mesons measured by ALICE as a function of $\mathrm{p_T}$ for Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02\,\mathrm{TeV}$ within the centrality classes of $0-10\%$, $30-50\%$ and $60-80\%$. The suppression increases as from the least to the most central Pb-Pb collisions. The stronger suppression for central collisions compared to peripheral collisions is due to the increased density, size and lifetime of the QGP. The R_{AA} for p-Pb collisions at the same centre-of-mass energy is also shown. This is consistent with unity, indicating the suppression seen in the Pb-Pb system is due to the formation of a hot and dense QGP.

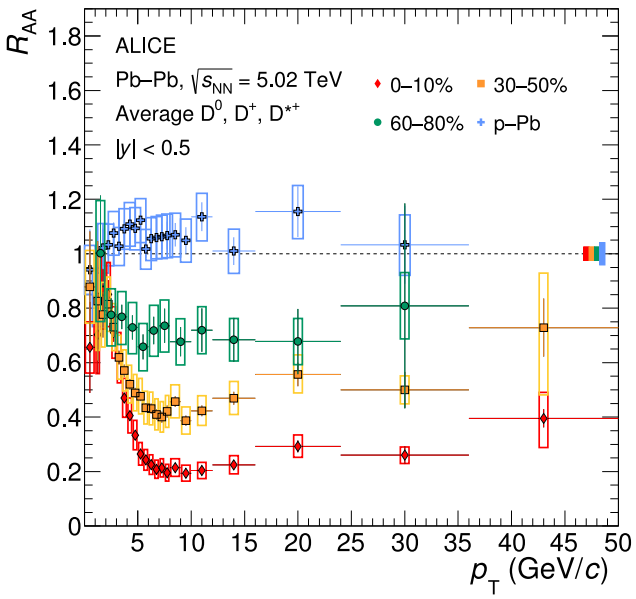

Figure 3.11: Prompt D-meson R_{AA} measured by ALICE as a function of $\mathrm{p_T}$ for Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02\,\mathrm{TeV}$ within the centrality classes of $0-10\%$, $30-50\%$ and $60-80\%$, as well as p-Pb collisions at the same centre-of-mass energy. Taken from [62].

3.4 Inner tracking system design for QGP physics

The design of the ALICE detector is of course driven by the physics requirements and the experimental conditions created in ultra-relativistic heavy-ion collisions by the LHC. The extremely high charged particle multiplicity of heavy-ion collisions, which can be up to three orders of magnitude higher than in pp collisions of the same nucleon-nucleon centre-of-mass energy [63], represents one of the most demanding constraints on the design of the ALICE ITS. High charged particle multiplicities demand a tracking system with a high granularity and readout rate. Furthermore, a number of rare heavy flavour decays are difficult to identify over the combinatorial background at the trigger level. To achieve satisfactory statistics, a large minimum-bias data sample must be obtained, exaggerating the need for an ITS to have a high readout rate.

The range in transverse momentum for QGP observables spans more than three orders of magnitude. Jet physics is typically on scales of well over $100\,\mathrm{GeV/c}$, whilst collective effects and low $\mathrm{p_T}$ heavy-flavour hadron decays are observed on a scale of typically tens of MeV/c. Measurements of observables across a wide momentum range dictate the necessity of an ITS with a low material budget to minimise multiple scattering at low $\mathrm{p_T}$.

Low $\mathrm{p_T}$ charm and beauty hadrons decay quickly (for instance the $c\tau$ of Λ_c is $60.7\,\mu\mathrm{m}$, two times smaller than that of D^0), motivating the need to maximise the primary and secondary vertexing capabilities, achieved by minimising the impact parameter resolution. This can be achieved by constructing an ITS with a low material budget, a low intrinsic spatial resolution and an innermost tracking layer with the smallest possible radius.

The designs of the previous ALICE ITS and upgraded ALICE ITS will be outlined in Chapter 4. The effects of the upgraded ITS on the measurements of QGP observables will be discussed in Sec. 4.4.2.

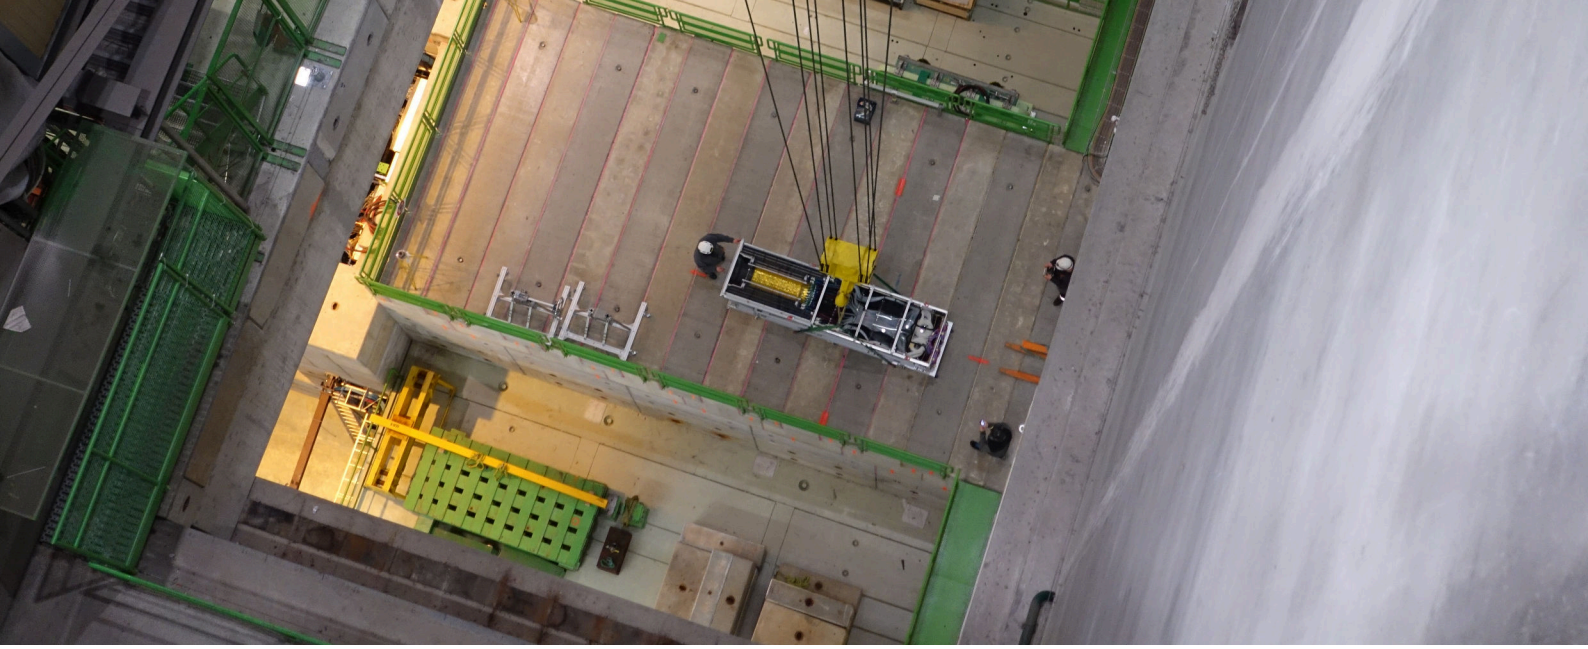

Chapter 4:
ALICE upgrade over LS2

This chapter will give an overview of the LHC and ALICE, with a focus on the sub-detectors upgraded over LS2. The upgraded ITS, ITS2, will be discussed in detail.

4.1 The LHC

The Large Hadron Collider (LHC) at CERN is a 26.7 km circular hadron accelerator and collider situated on average 100 m underground below the bordering regions of France and Switzerland near Geneva. The LHC is the last acceleration step of the CERN acceleration complex, capable of accelerating protons up to a centre of mass energy of $\sqrt{s} = 13.6\,\text{TeV}$ and lead nuclei up to a centre of mass energy per nucleon of $\sqrt{s_{NN}} = 5.02\,\text{TeV}$. Lead nuclei, whose atomic number is 82 and whose mass number is 208, are accelerated to a centre of mass energy of 13.6 TeV, so their centre of mass energy per nucleon is $\frac{82\times 13.6\,\text{TeV}}{208} \approx 5\,\text{TeV}$. The LHC is designed to accelerate 2808 proton or 592 lead nuclei bunches per beam. The beams are deflected around the LHC with 1232 superconducting dipole magnets, cooled to 1.9 K with 120 tonnes of liquid helium, and producing a magnetic field of 8.4 T. Quadrupole magnets control the transverse beam dimensions whilst finer beam corrections are

Header photograph shows the bottom half of the Outer Barrel after its descent via crane to the bottom of the ALICE cavern. © CERN. Taken from [2].

provided by higher order multipole magnets. The beams traverse ultra-high vacuum conditions where pressures are $O(10^{-9}\,\text{mbar})$.

The luminosity is a measure of the collision frequency per unit area. The LHC luminosity is defined as:

$$\mathscr{L} = f\frac{N_1 N_2}{4\pi\sigma_x\sigma_y}, \tag{4.1}$$

where f is the beam crossing frequency, N_1 and N_2 are the numbers of particles in each colliding bunch and σ_x and σ_y describe the transverse beam profiles in the horizontal and vertical planes.

In LHC Runs 3 and 4, from 2022 - 2030, the LHC aims to deliver an integrated luminosity of $13\,\text{nb}^{-1}$ for Pb-Pb collisions, corresponding to an interaction rate of 50 kHz.

4.2 Overview of ALICE in LHC Runs 1 and 2

ALICE is optimised to study relativistic heavy-ion collisions. The high particle multiplicities produced in relativistic Pb-Pb collisions require ALICE to be highly granular. The need for an efficient and robust tracking system capable of maintaining a reliable performance with up to 8000 charged particles per unit of rapidity motivated the use of a Time Projection Chamber (TPC) as the main tracking detector [64], with the caveat of a limited readout rate of 1 kHz during Runs 1 and 2, due to the TPC drift time of 100 µs [65]. In addition, bulk observables such as elliptic flow require the capability to detect low-p_T particles, necessitating ALICE to have a minimum material budget. The material budget of the central region of ALICE is the lowest of the four main LHC experiments. A magnetic field of 0.5 T, a compromise between momentum resolution, low momentum acceptance and tracking efficiency, improves tracking of low p_T particles in comparison to ATLAS and CMS, which use magnetic fields of 2 T and 3.8 T respectively [66, 67].

Figure 4.1 shows the ALICE apparatus used in LHC Runs 1 and 2. The ALICE sub-detectors can be split into three groups: the central barrel, the forward detectors and the muon arm. The central barrel consists of the Inner Tracking System (ITS), the TPC, the Transition Radiation Detector (TRD), the Time Of Flight detector (TOF), the High-Momentum Particle Identification Detector (HMPID), the ElectroMagnetic Calorimeter (EMCal), the Di-jet Calorimeter (DCal), the photon

spectrometer and the Charged Particle Veto detector (CPV). The central barrel subdetectors are used mainly for tracking, vertexing and Particle IDentification (PID).

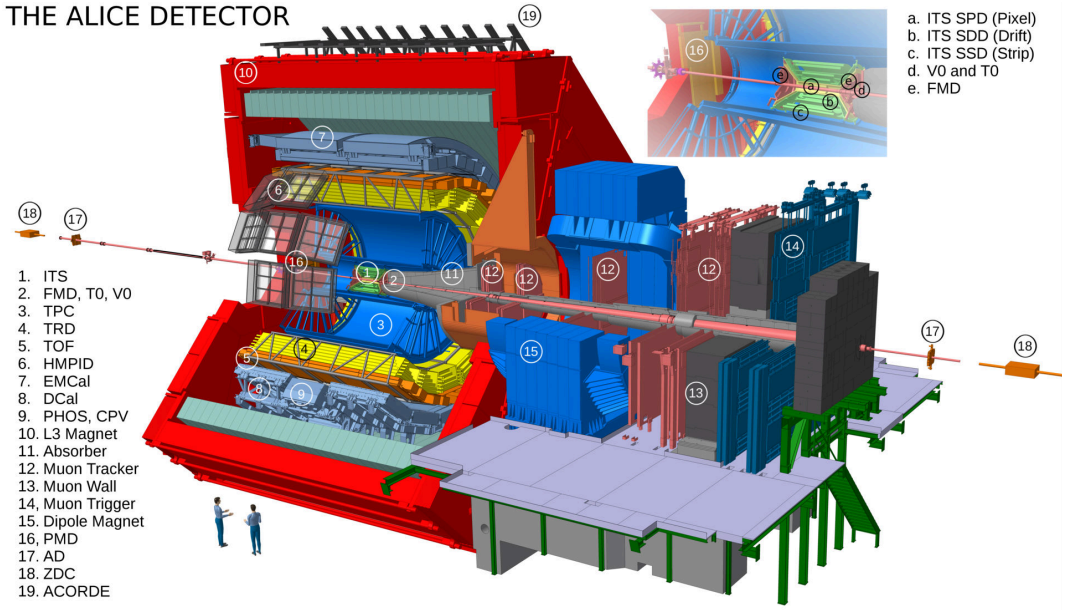

Figure 4.1: The ALICE apparatus used in LHC Runs 1 and 2. From the ALICE figure repository.

4.2.1 Inner Tracking System

The ITS surrounds the beam-pipe at the interaction point. In Runs 1 and 2 of the LHC, the ITS consisted of six cylindrical layers using three different kinds of silicon technology: two innermost layers of Silicon Pixel Detector (SPD)s, the two middle layers of Silicon Drift Detector (SDD)s and two outermost layers of Silicon Strip Detector (SSD)s. A schematic of the ITS used in LHC Runs 1 and 2 can be seen in Fig. 4.2

The main requirements of the ITS used in Runs 1 and 2 were a primary interaction vertex localisation with a resolution better than 100 µm, a secondary vertex reconstruction of hyperon, D and B meson decays, the tracking and particle identification of particles with a transverse momentum below 200 MeV/c or particles traversing the dead zones of the TPC, and the improvement of momentum and angular resolution for particles reconstructed by the TPC [68].

The innermost two layers of SPDs were fundamental for primary vertex localisation and secondary vertex reconstruction. Pixel detectors were chosen to cope

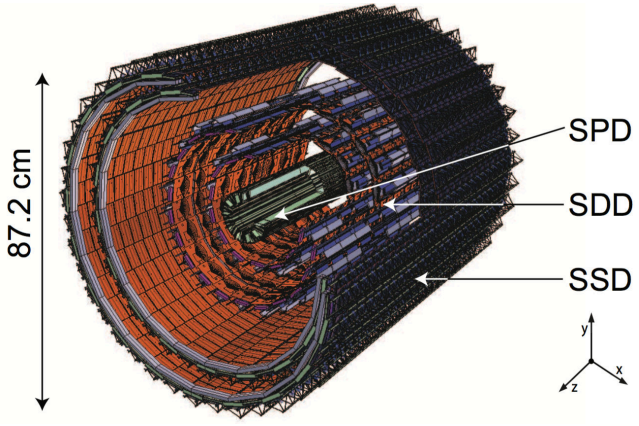

Figure 4.2: Schematic of the ITS used in Runs 1 and 2. Two layers of SPD surround the beam-pipe. Two layers of SDD encompass the SPDs. Two layers of SSD surround the SDD.

with the high track density (up to 100 tracks/cm^2 for the innermost two layers for Pb-Pb collisions at $\sqrt{S_{NN}} = 5.5\,\text{TeV}$ [69]), since they provide a high-granularity with a high signal-to-noise ratio.

The SDD and SSD layers had charge readout enabling the ITS to be used for stand-alone particle identification in the non-relativistic region via dE/dx. The middle layers were subject to a charged particle density of $O(10\,\text{cm}^{-2})$ so were constructed with SDD technology which has a good granularity and allows the measurement of energy loss [70]. The outer layers were subject to a charged particle density of $O(1\,\text{cm}^{-2})$ so the use of SSD technology was appropriate, allowing energy-loss measurements.

The SPD modules were comprised of silicon hybrid pixel detectors, where the 200 µm thick sensor and 150 µm front-end electronics chip were manufactured on two separate silicon wafers and bump bonded together. The front-end chip contained the readout circuitry for the 8192 pixels in the sensor chip. Each pixel had the dimensions of 50 µm $\times$ 425 µm in r$\phi \times$ z. The intrinsic resolution was 12 µm and 100 µm in rϕ and z respectively [71]. The material budget was 1.14% X_0 per layer. The readout circuitry consisted of a pre-amplifier with leakage current compensation and a discriminator with threshold tuning. Pixel readout was binary, where a signal exceeding the threshold is stored as a logical one in an on-pixel four hit deep front end buffer.

The SDD modules used a high-resistivity, 300 µm thick, n-type silicon wafer. A central cathode divided the sensitive area into two drift regions along the bending

direction (rϕ), where electrons move in opposite directions. The wafer was segmented on both surfaces with 292 parallel cathode strips with a pitch of 120 µm which deplete the detection volume and generate a drift field ($\approx 500\,\mathrm{V/cm}$). The drift field directs electrons liberated in the active area towards 256 collection anodes with a 295 µm pitch aligned to the beam axis [72]. The material budget was 1.13% X_0 and 1.26% X_0 for the inner and outer middle layers respectively.

The SSD modules consisted of 768, 300 µm thick, double-sided strip sensors connected to two HAL25 front-end chips via aluminium-kapton microcables. The strips had a pitch of 95 µm and a P-N-side stereo angle inclination of 35 mrad, allowing the SSD to provide a stereo view as well as a good ghost hit rejection [70]. The material budget was 0.83% X_0 per SSD layer.

4.2.2 Time Projection Chamber

The TPC surrounds the ITS. It is the primary tracking and PID detector of ALICE. It provides charged particle momentum measurements with vertex determination, track separation and PID via dE/dx with a resolution of 5%. The TPC can operate across a $\mathrm{p_T}$ range of $0.1\,\mathrm{GeV/c}$ to $100\,\mathrm{GeV/c}$, where the transverse momentum resolution is about 5% for $\mathrm{p}_T < 20\,\mathrm{GeV/c}$ and about 35% for $\mathrm{p}_T = 100\,\mathrm{GeV/c}$ when used without the ITS in stand-alone mode. When used with the ITS and the TRD, the transverse momentum resolution increases by a factor of 5-7, dependent on the $\mathrm{p_T}$ [68]. The material budget is 3.5% X_0.

Figure 4.3 shows a sketch of the TPC. The TPC has an inner radius of 0.85m, an outer radius of 2.5m and a length of 5m, resulting in a total active volume of $88\,\mathrm{m}^3$. A cylindrical field cage houses a central high-voltage electrode, dividing the active area into two halves and providing a drift field of $400\,\mathrm{V/cm}$. The drift volume is filled with a mixture of Ne, CO_2 and N_2 with a ratio of 16:2:1. This gas mixture was chosen because the Multiple Coulomb Scattering (MCS) is low, the diffusion characteristics are good and the positive ion mobility is high to clear ions from the drift volume relatively quickly. Electrons ionised from a particle traversing the gas mixture can be transported a distance of up to 2.5 m to the endplates, with a drift velocity of up to $2.7\,\mathrm{cm/\mu s}$.

Each of the two endplates contains 18 azimuthally segmented readout chambers, further segmented radially into inner and outer readout chambers, resulting in 72 in total. The radial segmentation is necessary to deal with higher track densities

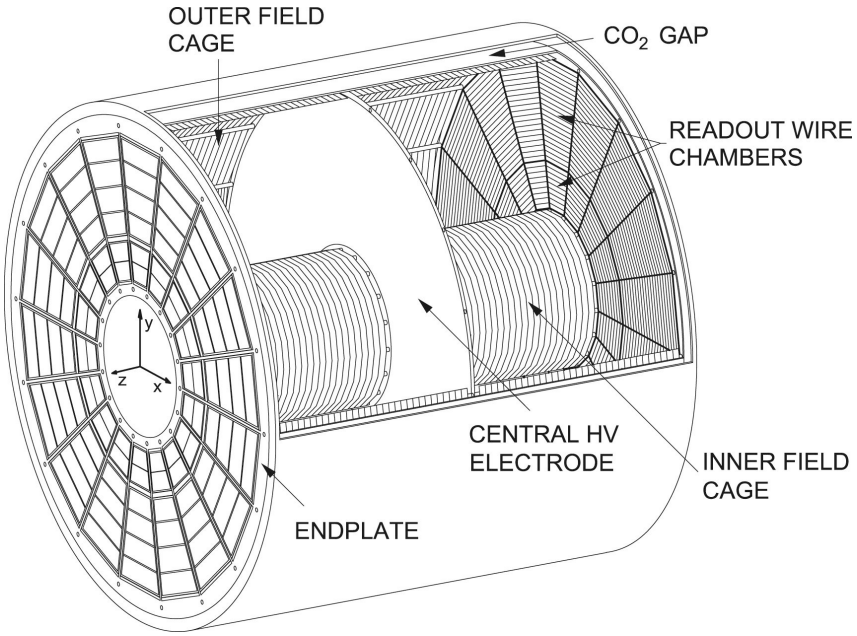

Figure 4.3: Sketch of the ALICE TPC. The drift volume is filled with a gas mixture. Charged particles ionise the gas mixture, where the liberated electrons drift to the MWPC readout chambers located in the endplates. Taken from [65].

towards the centre. The readout chambers used in the TPC for LHC Runs 1 and 2 were based on MWPC technology with gated readout.

Charged particles traversing the drift chamber ionise the gas mixture. Ionised electrons drift to the MWPCs in the end plates, where their x, y arrival point is measured. The z coordinate (along the axis of the field cage) is determined by the time for charge to drift from the ionisation event to the MWPC readout chambers, providing a full 3D track reconstruction.

The MWPCs consist of an array of cathode pads. Each of the segmented readout chambers houses 159 pad rows in the radial direction. Three pad sizes are used depending on the radial position, where the smaller pads are used towards the centre of the readout chambers: $4 \times 7.5\,\mathrm{mm}^2$ pads are used for the inner readout chamber, whilst $6 \times 10\,\mathrm{mm}^2$ and $6 \times 16\,\mathrm{mm}^2$ are used in the outer readout chamber for radii below and above 198.6 cm respectively [73]. A neon based gas mixture fills the chamber such that traversing particles ionise the atoms in the gas. Electrons move to the nearest anode further ionising the Neon gas on their journey. The charge amplification via further ionisation is known as a Townsend avalanche.

Ions that return to the drift volume accumulate and cause distortions to the drift field. This can occur in high multiplicity environments where the ion collection in

the MWPC becomes saturated. To keep the ion charge density below that created by primary ionisation, the MWPCs contain an active ion gating grid, used to collect ions from the amplification region whilst preventing them from returning to the drift volume. In the open gate mode, the wires in the gating grid are held at V_G, allowing electrons from the drift volume to enter the amplification region within the MWPC. In closed gate mode, the wires in the gating grid are biased with a bipolar field, $V_G \pm \Delta V$, preventing ions from drifting into the drift volume.

4.3 ALICE upgrade strategy during LS2

To allow high precision measurements of rare probes at low p_T the ALICE upgrade aims to both prepare for a significant increase in interaction rate delivered by the LHC, as well as improve the vertex reconstruction and tracking capabilities, in particular at low p_T . An integrated luminosity of $13\,nb^{-1}$ will be delivered during Run 3 for Pb-Pb collisions, which is a factor 100 larger for minimum-bias data compared to Runs 1 and 2 for minimum bias events. The resulting interaction rate of Pb-Pb collisions will be 50 kHz and the ALICE strategy is to read out all Pb-Pb interactions.

Over LS2 the following upgrades were completed:

- The beam pipe radius was reduced from 29.8 mm to 19.88 mm to allow the inner-most layer of the ITS to be moved closer to the interaction point.
- An entirely new ITS and Muon Forward Tracker (MFT) were installed.
- The TPC readout chambers were upgraded to increase the readout rate.
- The readout electronics of the TRD, TOF, PHOton Spectrometer (PHOS) and Muon spectrometer were improved to cope with the higher data rate.
- The online and offline systems (O^2 project) were upgraded to sustain the increase in data volume to 1 TByte/s.

The upgrades of the TPC and MFT will be briefly described in Secs. 4.3.1 and 4.3.2 respectively, whilst the upgrade of the ITS, the topic of this thesis, will be described in depth in Sec. 4.4.

4.3.1 Upgraded Time Projection Chamber

The increase in the Pb-Pb interaction rate to 50 kHz represents a challenge for the LHC Runs 1 and 2 TPC, due to its limited readout rate. The readout rate of the MWPC based readout chambers was limited by the active ion gating grid. The opening time of the gating grid is limited to 200 µs by the time needed to collect all the ions in the amplification region. Combined with the drift time of 100 µs of electrons in the drift chamber, the readout rate of the TPC was limited to about 3.3 kHz.

The upgrade of the TPC over LS2 focused on enabling a continuous, untriggered readout system, excluding the grating grid as an option. In order to increase the TPC readout rate, the MWPC based readout chambers have been upgraded to Gas Electron Multiplier (GEM) based readout chambers. The GEM comprise 50 µm thick insulating Polyimide foil with 2 - 5 µm thick copper coated surfaces. A regular pattern of holes is present on the foil formed via photo-lithographic processing. The holes have an inside diameter of 50 µm, an outside diameter[1] of 70 µm and a pitch of 140 µm. A potential difference of 200 - 400 V (depending on the gas mixture used) applied between the metal layers creates a field strength of 50 kV/cm within the holes [74].

The foil sits between a drift volume and an amplification region, where both regions are filled with gas. The potential difference between either side of the foil accelerates electrons from the drift region into the amplification region, causing a Townsend avalanche in the amplification region. The electrons then drift to a collection electrode. The majority of positive ions drift back to the drift volume side of the GEM foil. Some ions flow back into the drift volume. The ion backflow must remain below 1% in order to limit the effects of the drift field distortions to within the drift volume region 10 cm away from the GEM [73]. Since no active ion gating grid is required, a continuous, untriggered readout is possible, where the readout rate is limited only by the drift time of 100 μs, to 10 kHz.

4.3.2 Muon Forward Tracker

The MFT is a silicon tracking detector covering the pseudo rapidity region of $-3.6 < \eta < -2.5$. This pseudorapidity region is not covered by other LHC detectors, mak-

[1] The holes in a GEM foil are wider on the surface and thinner in the centre. The outside and inside diameters refer to the diameters on the surface and in the centre of the hole respectively.

ing the measurements unique and complementing those carried out by the ITS at mid-rapidity. It is constructed using a Monolithic Active Pixel Sensors (MAPS) design known as ALPIDE, which also forms the basis for the upgraded ITS and is discussed thoroughly in Sec. 4.4.3. The MFT consists of five disks, where each disk is subdivided into two half-disks. Each half-disk fits above and below the beam pipe. Each half-disk has a detection plane on the front and back sides, where the sensors on the front and back overlap to ensure hermeticity.

The position of the MFT relative to the upgraded ITS can be seen in Fig. 4.4.

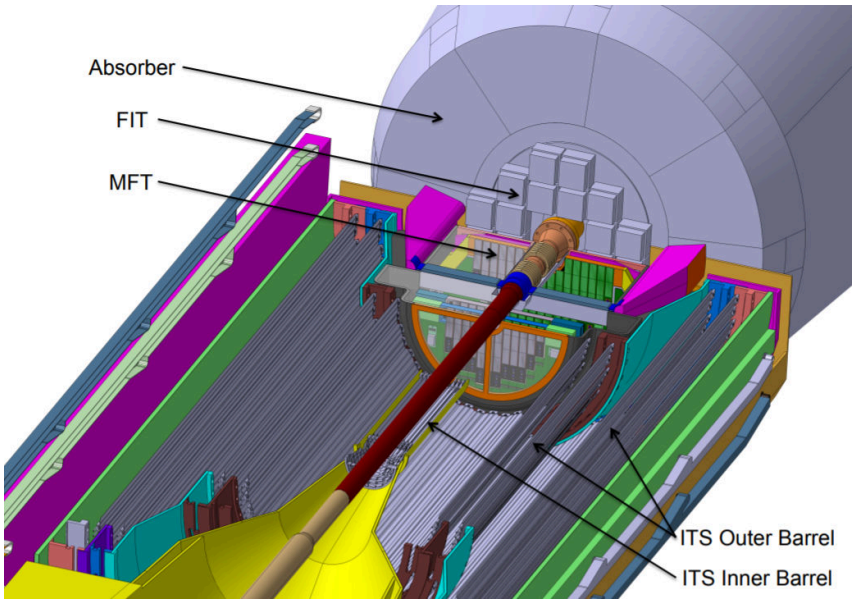

Figure 4.4: The positions of the MFT relative to the upgraded ITS. Taken from [75].

4.4 The ALICE ITS upgrade

During the LHC Long Shutdown 2 (LS2), the ALICE ITS was replaced by an entirely new system. This section gives an overview of the motivations for the ITS2 and the technology and components used.

4.4.1 ITS upgrade objectives

The objectives of the ITS upgrade (ITS2) can be split into the following four categories [25]: impact parameter resolution improvement, tracking efficiency and

p_T resolution improvement at low p_T , readout rate improvement and ease of maintenance improvement.

The pointing resolution will be improved by a factor of three and six in the rϕ and z coordinates respectively, as shown in Fig. 4.5 (left). This will be achieved by moving the innermost layer 16 mm closer to the beam-pipe, reducing the pixel pitch from 50 µm × 425 µm to approximately 30 µm × 30 µm and reducing the material budget to 1.1% X_0 per layer for the outer layers and 0.35% X_0 per layer for the inner layers.

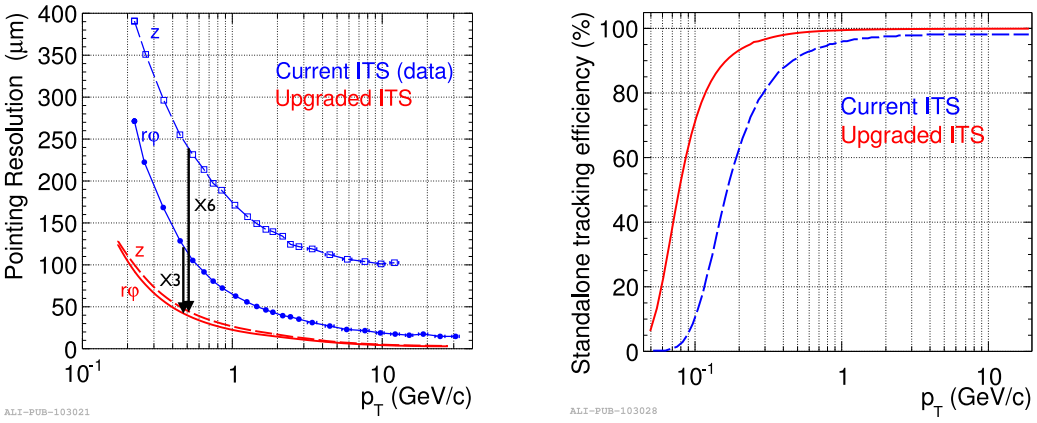

Figure 4.5: Expected performance plots comparing the previous ITS and ITS2 obtained from Monte Carlo simulations. Pointing resolution is shown on the left, whilst stand-alone tracking efficiency is shown on the right. Taken from [25].

The tracking efficiency and p_T resolution at low p_T will be improved, as shown in Fig. 4.5 (right), by using seven detection layers instead of six, and by using exclusively MAPS technology, leading to an improved single hit resolution throughout the entire ITS.

The readout rate will be increased from 500 Hz in Pb-Pb and 1 kHz in pp collisions to 50 kHz in Pb-Pb and 200 kHz in pp collisions. The increase in readout rate will allow the readout of all events in Pb-Pb collisions, deprecating the need for low level triggers that are inefficient due to the large combinatorial background [10].

ITS2 will sit inside an assembly cage that can be removed relatively easily during the yearly LHC shutdown periods for maintenance.

ITS2 consists of seven layers of a Complementary Metal Oxide Semiconductor (CMOS) MAPS design known as ALice PIxel DEtector (ALPIDE) [76]. A schematic of the ITS2 is shown in Fig. 4.6. The innermost three layers of the ITS2 are known as the Inner Barrel (IB), whilst the outermost four layers are known as

the Outer Barrel (OB). The OB is further segmented into the Middle Layers (ML) and Outer Layers (OL) which are the innermost two layers of the OB and the outermost two layers of the OB respectively. The dimensions and expected occupancies of ITS2 are shown in Table 4.1. The pseudorapidity region of $|\eta| < 1.22$ for 90% of the most luminous area is covered by the tracker, which has an active region of roughly $10\,\mathrm{m}^2$, segmented in 12.5 billion pixels.

The key feature of ITS2 is the use of ALPIDE chips, which are discussed in section 4.4.3.

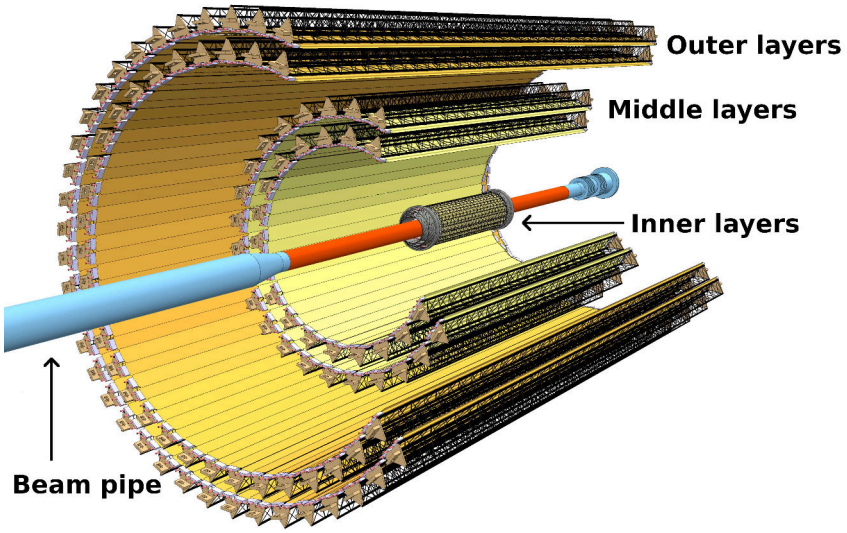

Figure 4.6: Layout of the ITS2. The central three concentric layers are known as the Inner Barrel, while the middle two layers are known as the Middle Layers and the outer two layers are known as the Outer Barrel. The outermost layer has a radius of 40.5 cm and a length of 147.5 cm. See Tab 4.1 for the dimensions of the ITS2. Taken from [25].

Table 4.1: Dimensions and expected occupancy for each layer in ITS2. Occupancy (Occ.) assumes an active area of $4.3425\,\mathrm{cm}^2$ [25].

Layer	0	1	2	3	4	5	6
Stave type	IB	IB	IB	ML	ML	OL	OL
No. staves	12	16	20	24	30	42	48
No. chips	108	144	180	2688	3360	8232	9408
Radius (mm)	23	31	39	194	247	353	405
Length (mm)	271	271	271	843	843	1475	1475
Active area (cm^2)	421	562	702	10483	13104	32105	36691
Hit density (cm^{-2})	72.84	47.78	34.50	2.042	1.418	0.602	0.601
Occ. (pixels/chip)	316	208	150	8.87	6.16	2.61	2.61

4.4.2 Performance increase with ITS2

Full Monte Carlo simulations were developed to study the performance of ITS2 and published in the Technical Design Report [25]. The HIJING [77] event generator was used to simulate central Pb-Pb collisions at $\sqrt{s_{NN}} = 5.5\,\text{TeV}$ with an impact parameter of $0 < \text{b} < 5\,\text{fm}$. For Heavy-Flavour (HF) performance studies, the HIJING events were enhanced with heavy-quark pairs generated with PYTHIA6 [78]. A complete GEANT4 simulation of ITS2 was produced, including transport through the detector, as well as tracking and vertex reconstruction. The average pixel noise distribution and fake-hit rate were extracted from ALPIDE test beam measurements and implemented in the pixel response simulation. The improvement in pointing resolution of a factor of three and six in the rϕ and z coordinates respectively, shown in Fig. 4.5, will be crucial for low p_T measurements, including measurements of heavy-flavour hadrons used to quantify in-medium energy loss and heavy-quark thermalisation.

HF baryons: Λ_c and Λ_b

The $\Lambda_\text{c} \rightarrow \text{pK}^-\pi^+$ decay is a challenging measurement given the short decay length of the Λ_c, $c\tau \approx 60\,\mu\text{m}$. The measurement is important for measuring the $\Lambda_\text{c}^+/\text{D}^0$ ratio to understand the importance of quark coalescence on heavy-quark hadronisation, as discussed in Sec. 3.3.5. With the ITS used for LHC Runs 1 and 2, the Λ_c could be studied with a significance[2] of about 5 in pp collisions, and a poor significance in Pb-Pb collisions due to the large combinatorial background.

The improved pointing resolution of the ITS2 will increase the significance of Λ_c measurements to about 8 for Pb-Pb collisions at $\text{p}_T = 2\,\text{GeV/c}$ for the centrality class of 0 - 20%. The improvement in significance and S/B is shown in Fig. 4.7.

Figure 4.8 shows the enhancement of the $\Lambda_\text{c}/\text{D}^0$ ratio as a function of p_T for Pb-Pb collisions at $\sqrt{S_{NN}}$ in the 0 - 20% centrality class, assuming an integrated luminosity of $10\,\text{nb}^{-1}$. The statistical uncertainties for pp collisions were assumed to be negligible compared to Pb-Pb collisions.

Given ALICE achieves its target luminosity, a previously unachievable measurement of the Λ_c baryon will become possible with good precision down to a p_T of $2\,\text{GeV/c}$, permitting the study of the importance of quark coalescence in heavy-flavour hadronisation.

[2] Significance is defined as $S/\sqrt{S+B}$, where S and B are the signal and background respectively.

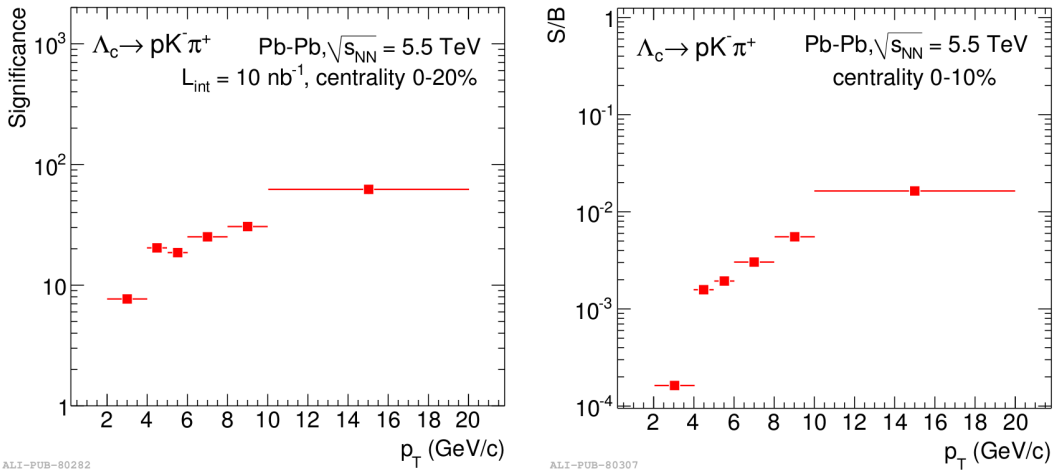

Figure 4.7: Expected significance (left) and S/B (right) for the Λ_c baryon as a function of p_T for Pb-Pb collisions at $\sqrt{S_{NN}}$ in the 0 - 20% centrality class, assuming an integrated luminosity of $10\,\text{nb}^{-1}$. See the text for a description of the MC simulation performed. Taken from [25].

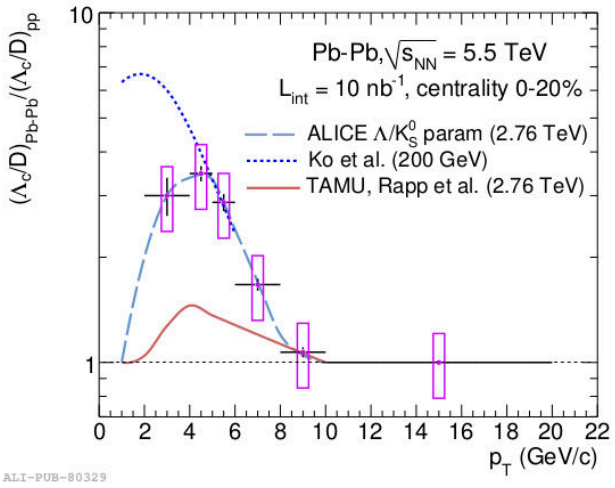

Figure 4.8: Enhancement of the Λ_c/D^0 ratio as a function of p_T for Pb-Pb collisions at $\sqrt{S_{NN}}$ in the 0 - 20% centrality class, assuming an integrated luminosity of $10\,\text{nb}^{-1}$. The expected statistical uncertainties are drawn with vertical lines, whilst the expected systematic uncertainties are shown with boxes. The points are drawn on the ALICE Λ/K^0_S double ratio. Model calculations are shown, Ko et al. (200 GeV) [79], and TAMU, Rapp et al. (2.76 TeV) [80]. Taken from [25].

4.4.3 ALPIDE

The ALPIDE chip is a Monolithic Active Pixel chip designed for the ALICE ITS2 and manufactured by TowerJazz with their 180 nm CMOS imaging process. Figure 4.9 shows a schematic of an ALPIDE chip. The ALPIDE chip measures 15 mm × 30 mm and consists of 512 × 1024 pixels. ALPIDE is divided into an active matrix and a digital periphery, where the periphery is positioned on a (1.2×30) mm^2 edge. The chips are thinned down to 50 µm and 100 µm for the IB and OB respectively. Each pixel measures 29.24 µm × 26.88 µm and contains a n-well collection node, an analogue front-end and a digital readout. The analogue front-end consists of an amplifier, signal shaper and discriminator, whilst the digital read out contains a 3-bit Multi-Event Buffer (MEB).

A deep p-well shields the n-well, allowing the use of PMOS transistors inside the pixel matrix and therefore full CMOS logic. This makes it possible to have an amplifier, signal-shaper, discriminator and multiple event buffers in-pixel. The front-end stays continuously active and has a power consumption of 40 mWcm^{-2}.

A 2 µm diameter, low capacitance n-well diode used together with an epitaxial layer resistance of 1 kΩcm and a reverse bias voltage of -3 V contributes to a radiation tolerance of 270 krad Total Ionising Dose (TID)[3], 1.7×10^{12} MeV/n$_{eq}$ Non-Ionising Energy Loss (NIEL)[4], the expected dose after 10 years of operation in the ITS2 environment [81]. The n-well diode measures roughly 300 times smaller than the pixel size, which combined with the reverse bias and reduced capacitance of the n-well diode, increases the signal to noise ratio. This is important for good detection efficiency at low power consumption [82].

Operating principles

Figure 4.10 shows a schematic of signal flow within the ALPIDE front-end circuitry. A charged particle ionises electrons in the epitaxial layer, which move towards the N-well diode. As a result, the V_{PWELL} experiences a reduction in potential, lowering the potential of the pixel input node IN by ΔV_{IN}, $O(10\,\text{mV})$, within a charge collection time of $O(10\,\text{ns})$. The potential returns to the baseline voltage $V_{BASELINE}$ over the course of about 100 µs via the reset mechanism.

[3] TID is the energy deposited over time where 100 rad is defined as 1 J absorbed by 1 kg of material.

[4] NIEL is the energy loss to atomic displacements.

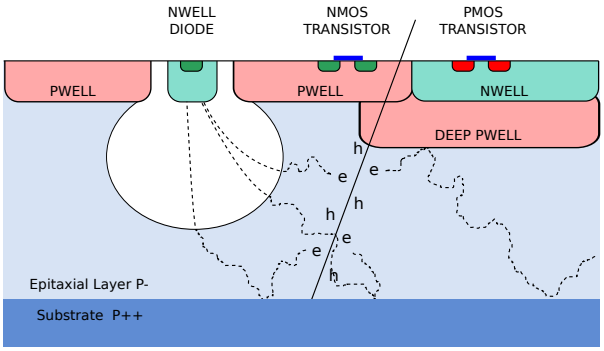

Figure 4.9: Cross-sectional schematic of the ALPIDE chip. A deep p-well shields an n-well allowing the implementation of NMOS and PMOS transistors, enabling full CMOS logic. Electrons ionised by a charged particle drift to the depletion region, where they are collected by an n-well diode. Taken from [25].

The amplifier and shaper deliver a pulse that has a peaking time of approximately 2 µs, and a total pulse width of less than 10 µs. The discriminator compares the pulse with a threshold configurable with the V_{CASN} and I_{THR} nodes. The V_{CASN} node has an exponential relationship with the threshold value, which increases as V_{CASN} decreases. The I_{THR} node has a linear relationship with the threshold value, which increases as I_{THR} increases. The discriminator changes the analogue input pulse HIT to a digital active-low signal HITb. The input to the hit buffer becomes enabled on the assertion of a STROBE. The address of the hit pixel is provided by the Address Encoder Reset Decoder (AERD).

The capacitor C_{inj} injects a test charge into the pixel by applying a negative voltage pulse via V_{PULSE}, in testing scenarios.

Data transmission

The pixels are arranged in 512 double columns which are read out by a combinatorial priority encoder with zero-suppression, meaning the only addresses sent to the chip periphery are those of hit pixels. No clock is propagated over the matrix, keeping the power consumption low at 40 nW for the front-end circuitry of each pixel. A single Region Readout Unit (RRU) reads out a group of 16 priority encoders sequentially. There exist 32 RRUs on each ALPIDE chip which read hits from the matrix in parallel. Hits are stored in a memory block per region, minimising latency and acting as de-randomising buffers. These data are transferred to a readout module sequentially via Low Voltage Differential Signal (LVDS) high-speed link or

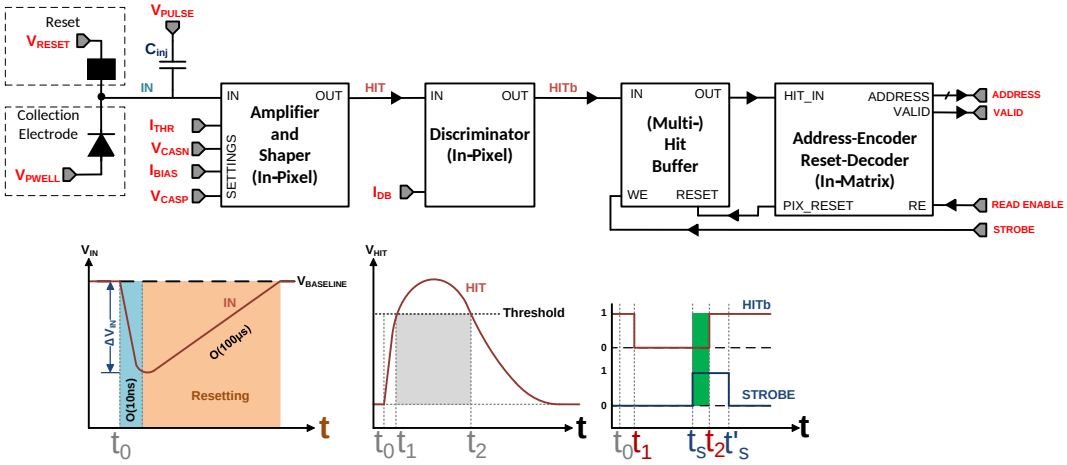

Figure 4.10: Schematic of the signal flow within the ALPIDE front-end circuitry. See text. Modified from [82].

parallel bus. The high-speed link and parallel bus provide throughputs of 960 Mb/s and 320 Mb/s for the IB and OB respectively. The line rates are 1.2 Gb/s and 400 Mb/s for the IB and OB respectively, where the difference in comparison to the throughput is from the 8b10b encoding overhead. 8b10b encoding is used to maintain a DC level of the differential signal close to 0V. For the IB, data are transferred from each chip to the off-detector electronics via LVDS high-speed link. Whilst for the OB, slaves transfer data to masters via parallel bus and masters transfer data to the off-detector electronics via high-speed link.

The Serial Data Transmission block takes care of the data transmission from master chips to the off-detector electronics. The Data Transmission Unit (DTU) consists of a Phase-Locked Loop (PLL), driver and serialiser. The serialiser receives the data from the digital periphery of the master chip and transmits it to the driver, which drives the transmission line. For the purposes of testing, a configurable pattern generator provides data to the input of the serialiser. A static-bit pattern, counter pattern or Pseudo-Random Bit Sequence (PRBS) pattern can be provided by the pattern generator.

The ALPIDE chip receives triggers and slow control via a differential port with a bitrate of 40 Mb/s. For the OB, masters relay the control and triggers to the slaves.

Operating modes

The ALPIDE chip can operate in continuous and triggered data acquisition modes. In continuous mode, a trigger assertion is managed internally by the ALPIDE chip

with a programmable time base. Continuous mode is illustrated in Fig. 4.11. The acquisition window remains open during the time between consecutive triggers. The comparator outputs are latched into the MEB during the acquisition window. An event on the boundary of consecutive acquisition windows can get recorded twice. The chip first reads out the newest frame request, rather than data already in the MEB. To ensure a free buffer in which to store a future hit, data read out in the latest frame undergoes deletion. In the case where the MEB becomes full, the chip interrupts the ongoing frame request to delete a hit from one of the MEB slices. The chip data packet trailer records the interruption.

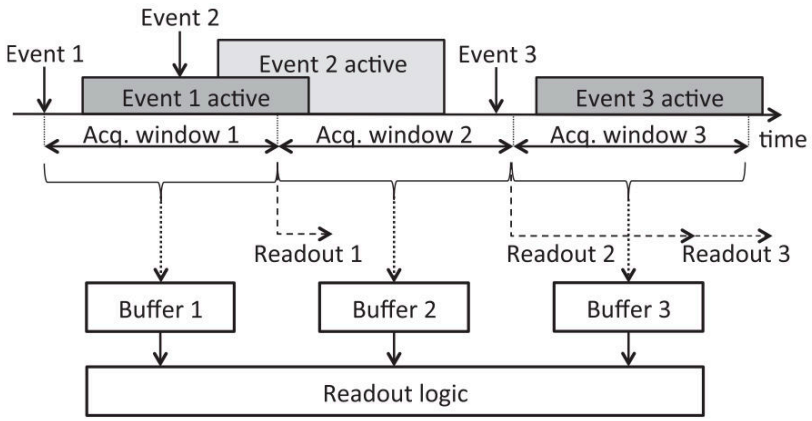

Figure 4.11: Continuous data acquisition mode. See text. Taken from [83].

In triggered mode, the trigger signal is connected to the ALICE Central Trigger Processor (CTP). An acquisition window is opened on the reception of each trigger. The comparator outputs are latched into the MEB during the acquisition window. Events closer in time than the shaping constant of the amplifier can get recorded twice. Hits already stored in the MEB get read out before new triggers. The reception of a trigger generates an acquisition window, during which the hits in the MEB are read out. Outside of this acquisition window, no hits are read out. If a new trigger is received whilst the MEB remains full, the generation of a new strobing interval does not occur. Instead, the chip data packet trailer records the occurrence and an empty chip data packet in the data stream indicates the trigger.

Triggered and continuous mode differ in the logic handling of the MEBs becoming full, as well as in the length of the acquisition window. Triggered mode uses a typical acquisition window of $O(0.01\,\mathrm{ns})$, whilst continuous mode uses a typical acquisition window of $O(1\,\mu\mathrm{s})$.

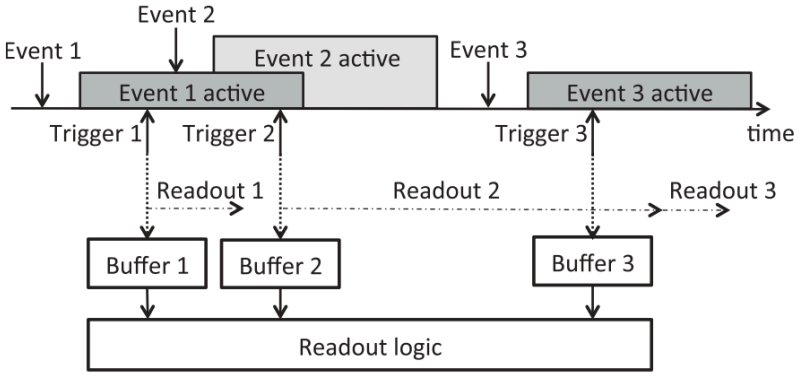

Figure 4.12: Triggered data acquisition mode. See text. Taken from [83].

4.4.4 ITS2 components

The IB, ML and OL Hybrid Integrated Circuit (HIC)s and staves, which transform the ALPIDE chips into the ITS2 detector, will be described in the following sections.

IB staves

Each IB stave consists of nine ALPIDE chips glued and wire-bonded to an Flexible Printed Circuit (FPC). Power, ground, clock, control and data are all propagated along the FPC, with separate data lines for each ALPIDE chip. A schematic of the readout architecture of an IB stave is shown in Fig. 4.13.

The silicon side of the ALPIDE chips is glued to a carbon fibre cold plate, under which cooling ducts lie. The carbon fibre space frame, glued to the cold plate, provides spatial support. A schematic of an IB stave is shown in Fig. 4.15 (left).

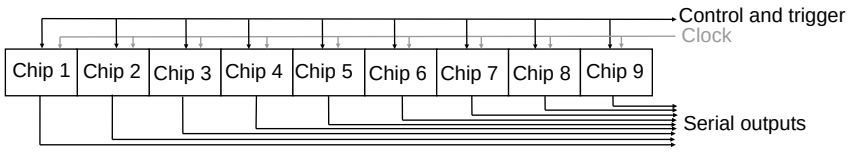

Figure 4.13: IB stave readout architecture [83]. See text.

OB Hybrid Integrated Circuits

Each OB HIC consists of 14 ALPIDE chips arranged into two rows of seven, glued and wirebonded to a FPC. The 14 ALPIDE chips consist of two master chips and

12 slave chips, where each row of the HIC consists of one master chip controlling 6 slave chips. The master chip forwards the clock and control from the off-detector electronics to its slave chips. The master chip also forwards the data from its slave chips to the off-detector electronics via high-speed link. A schematic of the OB HIC readout architecture is shown in Fig. 4.14.

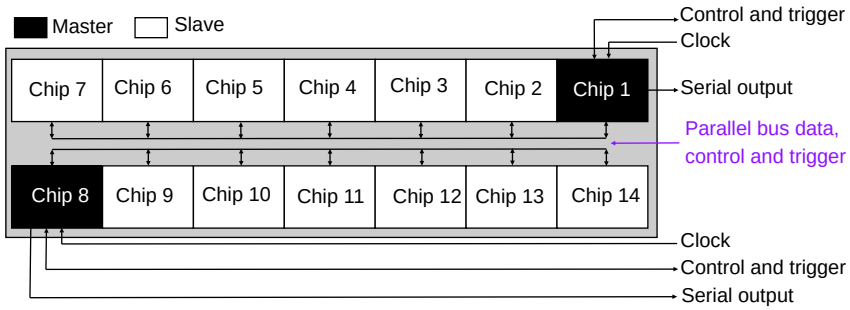

Figure 4.14: OB HIC readout architecture [83]. See text.

OB staves

Figure 4.15 (right) shows a schematic of an OB stave. An OB stave consists of a carbon fibre Space Frame (SF) and two electrically independent Half-Stave (HS). Each HS consists of a carbon fibre Cold Plate (CP), with a number of HICs glued to its surface, seven for OL staves and four for ML staves. The CP has one continuous water pipe, with one water inlet and one outlet. The breakdown of material budget per OB stave component can be seen in Tab. 4.2.

Each HS receives two differential clocks and two differential controls from the off-detector electronics. One of each is distributed to each side of each HIC. Each master has a dedicated high-speed link to the off-detector electronics. The clock, control and high-speed lines all pass through the FPCs. FPCs on adjacent HICs are bridge soldered together.

The power bus distributes the power to each HIC from the off-detector electronics, via the cross cables on each HIC. Three channels per HIC exist in the power bus and one ground line common to all HICs. The three channels are DVDD, AVDD and Bias.

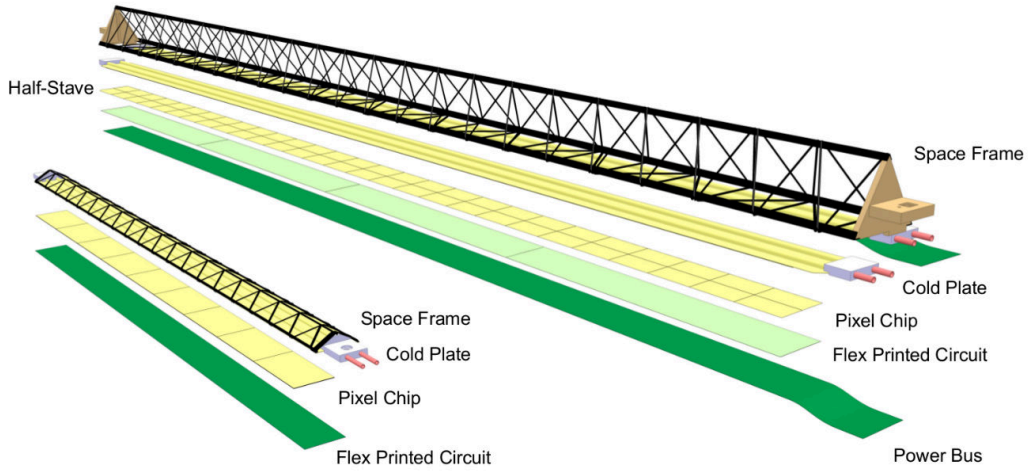

Figure 4.15: Schematic of an IB stave (left) and an OB stave (right). Taken from [25].

Table 4.2: Material budget breakdown of an OB stave.

Component	Material budget
ALPIDE	15%
FPC	50%
Cooling	20%
Space frame	15%

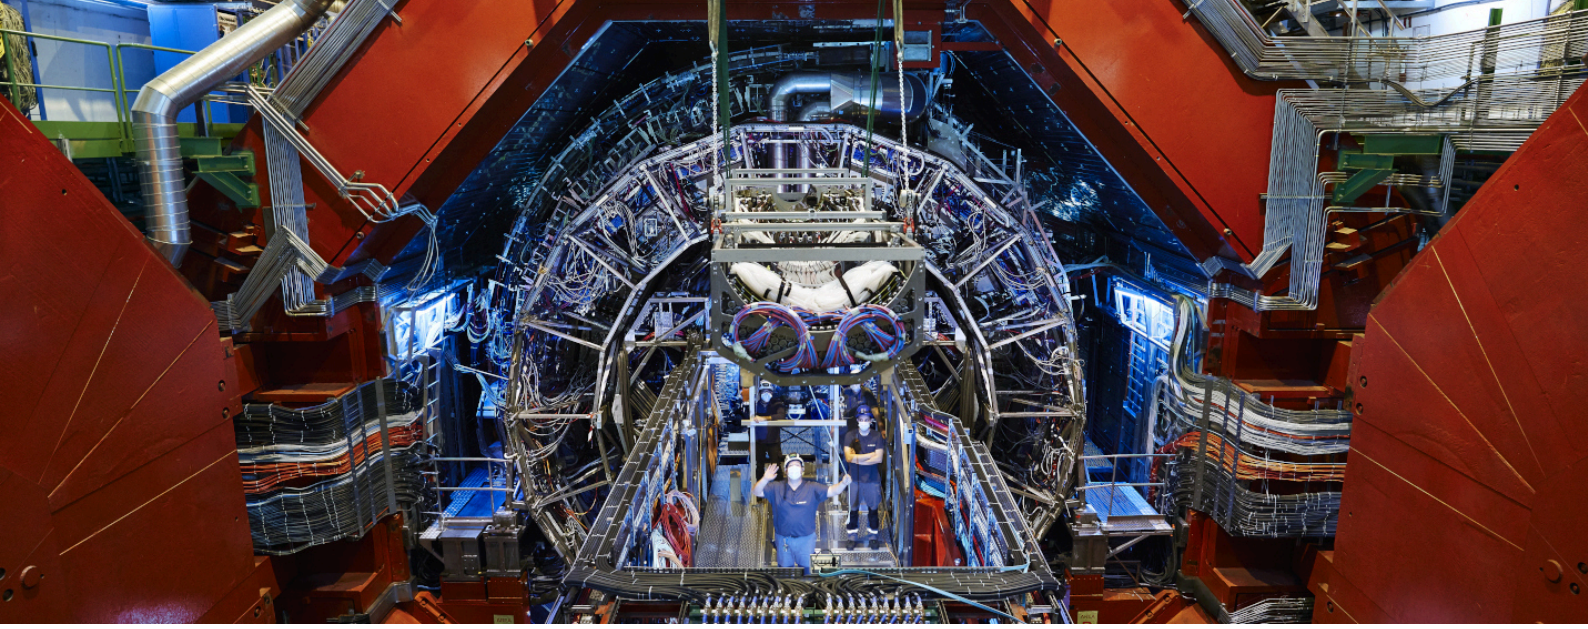

Chapter 5:
HIC and stave construction

The following chapter will discuss the HIC and stave construction in Sec. 5.2 and 5.3 respectively. The characterisation of HICs and staves will be discussed in Chapter 6.

5.1 Kybernetica database

Throughout the HIC and stave construction, all construction and characterisation test results were logged into a database created by Kybernetika (Košice, Slovakia). The Kybernetika database provides each activity and component with an ID, allowing the user to display the results for all components which participated in a given activity, as well as the results of all activities a given component participated in. The electrical and metrological results presented in Chapter 6 were uploaded to the Kybernetika database throughout the construction and characterisation procedures. The construction procedure is described in this Chapter.

5.2 Module construction

Chips are distributed from the chip test labs (CERN and Pusan) to one of four HIC assembly sites (Bari, Liverpool, Strasbourg and Wuhan). Chips are packaged in

Header photograph shows the bottom half of the Outer Barrel being lowered by crane onto the ALICE mini-frame. © CERN. Taken from [26].

Electro-Static Discharge (ESD) bags under vacuum. The Flexible Printed Circuit (FPC)s are shipped from their construction and test site (Trieste) on a transport plate in ESD boxes. After reception of the chips and FPCs, the chips are aligned in the ALICIA machine, whilst the glue mask is applied to the FPC. The FPC is then glued to the chips, after which the wire-bonds from the pads on the FPC are bonded to the pads on the chips. The process is described in more detail in the following sections.

5.2.1 Alignment of chips in the ALICIA

The ALICE Integrated Circuit Inspection and Assembly machine (ALICIA) is a custom machine built by IBS Precision Engineering [84] for the ITS2 project capable of positioning ALPIDE chips with an accuracy of $\pm$ 5 µm (3σ) with respect to reference features [85]. The machine uses high-resolution imaging to measure the position of each chip. The same imaging system is also used to check the quality of each chip, monitoring dimensions, cleanliness and checking for damage.

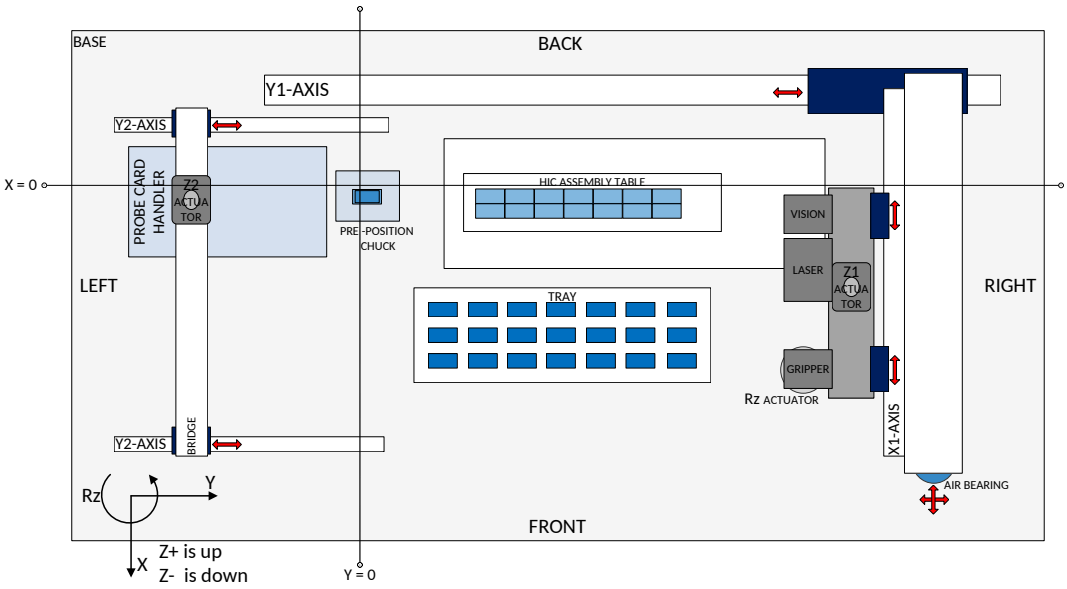

Figure 5.1: Schematic of the main components of the ALICIA. The gripper moves ALPIDE chips from the tray to the pre-position chuck, where the vision system is used to conduct QA. The gripper places ALPIDE chips on the HIC assembly table. Taken from [85]

A schematic of the ALICIA machine can be seen in Fig. 5.1. The vision system of ALICIA measures the position of the markers on the chips and of one of two sites where the chips are held via vacuum: the pre-position chuck and the HIC assembly

table. The pre-position chuck is used for Quality Control (QC) and initial alignment of the chips, whilst the HIC assembly table is where the HICs are deposited before gluing the FPC. The chip gripper is a vacuum handler used to manipulate the chip positions. The chip gripper is mounted to a high-precision frame which moves in X, Y and Z and rotationally in the XY plane (Rz).

The chips are distributed in mini-trays, where each mini-tray has a unique barcode. Four mini-trays are inserted into the ALICIA tray. The gripper is used to remove chips from the tray and place them on the pre-position chuck, where the chip orientation is measured and QC conducted via the vision system. Chips that meet the quality criteria are placed on the HIC assembly table, whilst chips that fail these criteria are placed back into the tray. The chip identity is known through the chip position in the mini-tray and the unique barcode on each mini-tray. The rejection or acceptance of a given chip is recorded in the Kybernetika database. The following criteria are used to determine if a chip should be rejected:

Defects on the chip edges and corners: A chip with an edge defect with a size of $< 10\,\mu$m is accepted and placed on the HIC assembly table. A chip with an edge defect with a size of $> 10\,\mu$m is rejected[85]. A chip with an edge defect is shown in Fig 5.2.

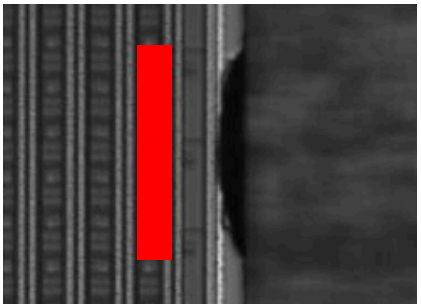

Figure 5.2: A damaged chip edge as seen through the ALICIA microscope. The damage (to the right of the red rectangle) compromises the seal ring which can be seen as the parallel lines on the edge of the chip.

Cleanliness of contact pad: The cleanliness of the contact pad on the chip is examined for dirt and damage. The surface cleanliness is defined as the proportion of the surface which is free from dust. If the surface cleanliness is $> 90\%$, the chip is accepted and placed on the HIC assembly table[85]. If the surface cleanliness is $< 90\%$, the chip is rejected.

5.2.2 FPC glue mask

Each FPC is transported in a carrier plate, where each carrier plate has a unique barcode and a label stating the FPC name. The FPC name has the form XY00ZZZZ. X can be either A or B to denote the tab type. A tab type of A denotes any HIC which is intended to be used in any position on the stave except for HIC 0, whilst a tab type of B is strictly for HIC 0. HIC 0 is the HIC that is closest to the RU and contains a data transmission tab on its FPC. Y can be either L or R to denote the position of the cross-cables, either left (L) or right (R). Z denotes four digits unique to the particular FPC. This FPC name is used throughout the construction process as an identifier for the FPC and, later, for the HIC.

The FPC is removed from the transport plate with a vacuum gripper. This vacuum gripper is then placed onto the working surface such that the underside of the FPC (without electrical connections) is facing up. The glue mask is placed onto the glue mask vacuum gripper before being lowered onto the FPC using dowel pins to achieve the correct alignment. Araldite 2011 is then applied onto the glue mask before being spread over the FPC using a tool to ensure the required thickness and uniformity of glue. Figure 5.3 shows a photograph of the application of glue to the glue mask adhered to an FPC.

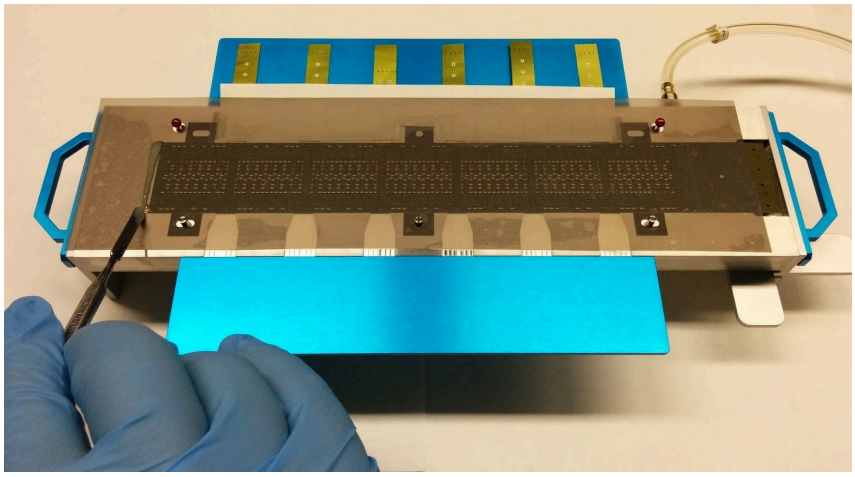

Figure 5.3: Photograph of glue being deposited on an FPC using a glue mask. Glue is applied in uniform lines across the mask before being spread evenly across the mask.

5.2.3 FPC and chip bonding

The FPC is connected to the ALPIDE chips mechanically, via glue, and electrically, via wire-bonds.

Gluing the FPC onto the chips

After the glue mask is removed, the FPC is ready to be glued onto the chips in the assembly tray of the ALICIA. The FPC is rotated with the FPC gripper such that the glue side faces down. The FPC gripper is lowered onto the chips using dowel pins for alignment and left for five hours to allow the glue to cure. After this time, the FPC gripper is removed from the FPC, leaving the FPC and chips on the assembly tray within the ALICIA.

The HIC is moved from the assembly tray with a HIC vacuum gripper after removing the vacuums applied to each chip in the assembly tray. The HIC is returned to the transport plate, which is moved underneath a microscope. The microscope is used to inspect the ALPIDE pads for glue contamination and correct alignment of the FPC above the pad. After successful inspection, the HIC is moved to the wire-bonding table.

Wirebonding the FPC to the chips

Wire-bonds are used to electrically bond the FPC to the chips. Three aluminium wire-bonds connect each FPC bonding pad to each chip bonding pad, where three wire-bonds are used for redundancy. The wire bonds are made using ultrasonic wedge bonding. Figure 5.4 shows a schematic of a wire-bond, whilst Figure 5.5 shows a photograph of the wire-bonding from an FPC bonding pad to the chip's bonding pad.

5.3 Stave construction

Fig. 5.6 and Fig. 5.7 show an overview of the workflow of the assembly and testing procedure of one OB HS and one OB stave respectively. This section outlines all procedures shown. The assembly procedure for ML and OL staves is identical apart from the number of HICs used and the dimensions of the cold plates and space frame.

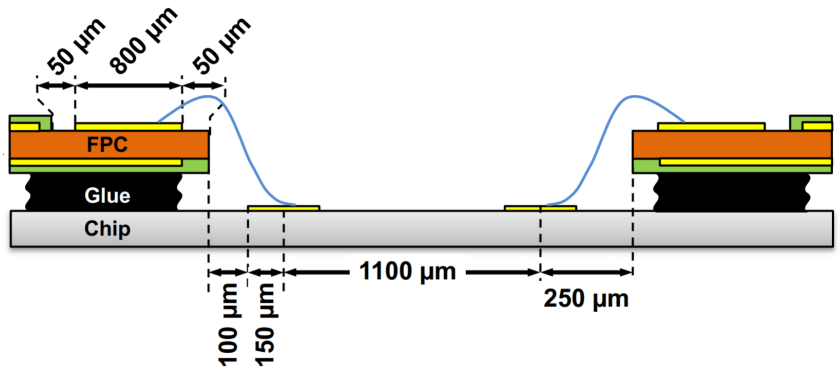

Figure 5.4: Schematic of the FPC and chip wire-bond connection. Taken from [86].

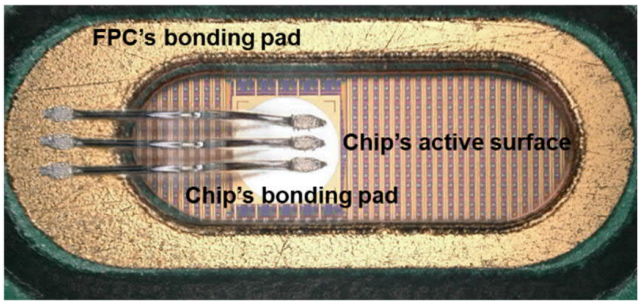

Figure 5.5: Photograph of the FPC and chip wire-bond connection. Taken from [86].

The stave construction was a joint effort across five international construction sites. The OL staves were constructed in Amsterdam (NIKHEF), Daresbury (STFC Daresbury), Frascati (LNF institute) and Turin (INFN institute). The ML staves were constructed solely in Berkeley (LBNL).

5.3.1 Coordinate-Measuring Machine

The stave assembly procedure involved the use of a Mitutoyo CRYSTA-Apex V9166 Coordinate Measuring Machine (CMM) to measure the position of four markers on each of the ALPIDE chips. The CMM is run in CNC mode to find the approximate location of the physical chip markers based on the nominal positions stored to file. The CMM is then run in manual mode to properly locate the ALPIDE marker. The CMM was also used to perform metrology measurements of the constructed staves, which is discussed further in Sec. 6.1. The coordinate system used by the CMM is shown in Fig. 5.8. The Y axis is parallel to the longest side of the stave, whilst the x axis is parallel to the shortest side of the stave. The z axis is perpendicular to the plane defined by the ALPIDE chips. The intrinsic resolution of the CMM is

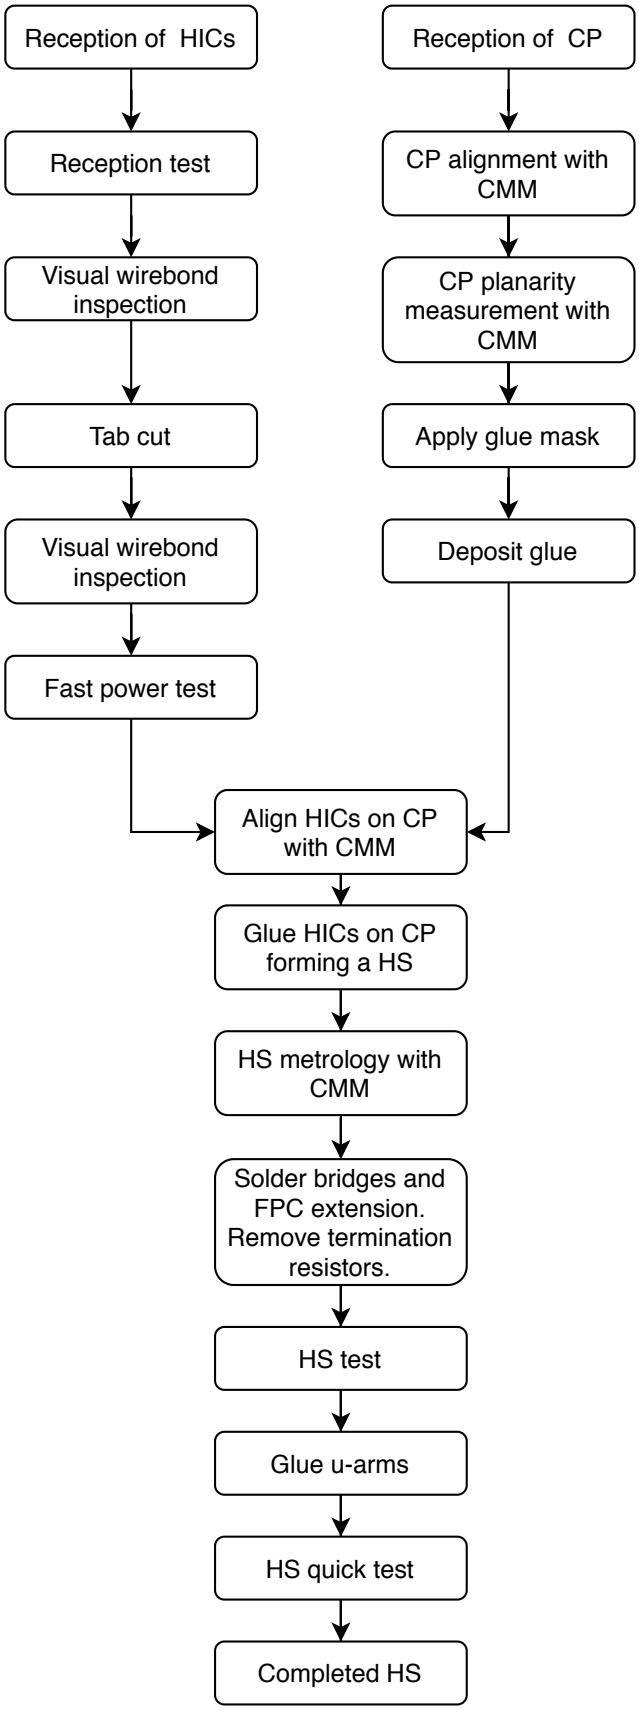

Figure 5.6: Workflow to produce one OB HS.

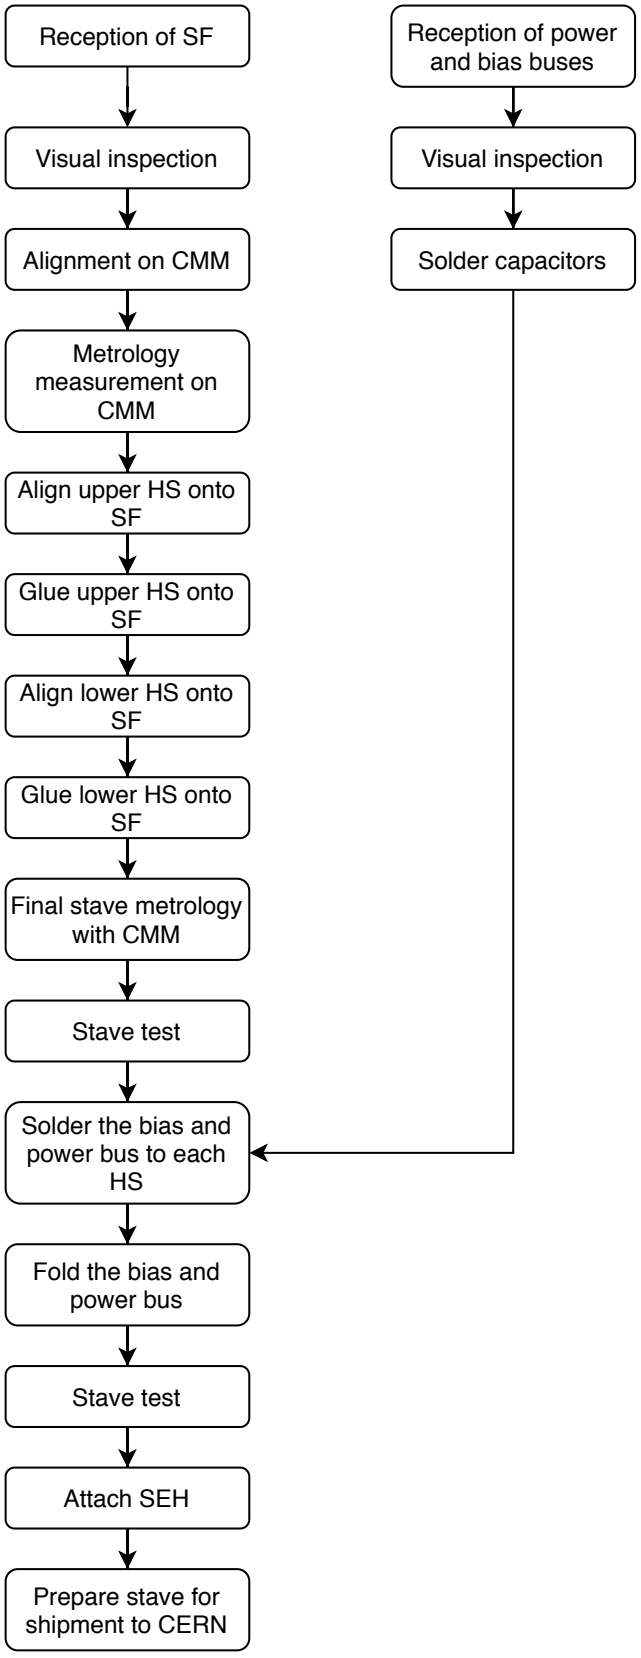

Figure 5.7: Workflow to produce one OB stave.

2 µm in the x and y and 10 µm along the z axis. The resolution in xy is dictated by the size of the chip marker shown in Fig. 5.8, whilst the resolution in z depends on the focus of the CMM camera. Results from the CMM metrology are stored in the Kybernetika database and discussed in Sec. 6.1.

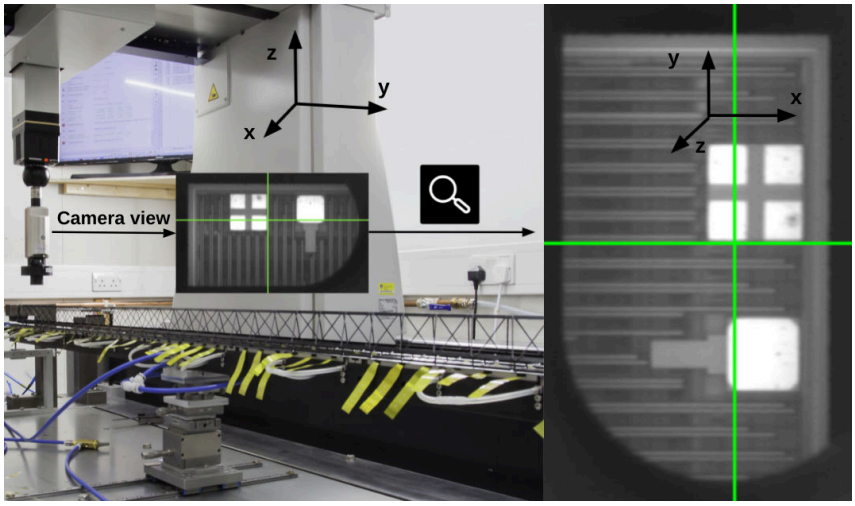

Figure 5.8: A single chip marker as seen with the CMM. The green cross must be aligned to the white cross on the marker to measure the position of the marker. This has a resolution of 2 µm in xy, whilst the focus of the camera contributes to the resolution of 10 µm in z. The coordinate system is shown for both the CMM room and CMM image.

5.3.2 Visual inspection

HICs are received by the stave construction sites after shipment from the HIC construction sites. The HICs are transported on a carrier plate shown in Fig. 5.9.

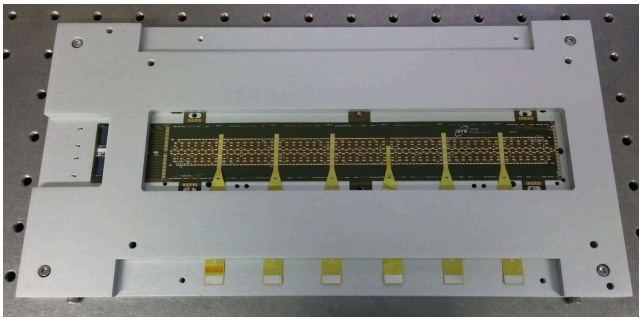

Figure 5.9: A HIC shown on a carrier plate after shipment from a HIC construction site to a stave construction site.

Using a microscope, the wire-bonds of each HIC are examined before and after the tab-cutting procedure.

5.3.3 Tab cutting

Each HIC has an FPC extension tab for electrical characterisation and 6 wings with which it attaches to the carrier plate. The extension tab and wings are removed before the installation of the HIC onto a HS, illustrated in Fig. 5.10. Once on the HS, a succession of bridge solder joints electrically connect adjacent HICs. The FPC extension and tabs are cut off using the following procedure.

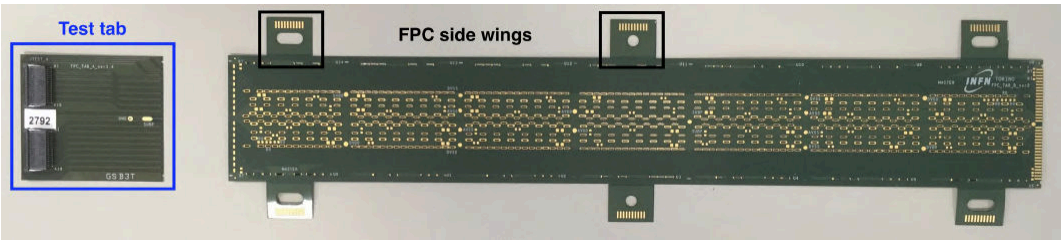

Figure 5.10: An FPC with the extension tab cut off. The six FPC wings are also removed during the tab-cut procedure. Figure from [87].

1. The HIC is lifted from the carrier plate to the blue plate of the module flipper using vacuum pads. Dowel pins are used to align the HIC which insert through the tabs on the HIC into the blue side of the module flipper, shown in Fig. 5.11. The red plate of the module flipper then lowers onto the HIC. The shape of the red side of the module flipper is designed such that the vacuum pads touch areas of the FPC where no wire-bonds exist, to ensure no damage occurs to the HIC. Both sides of the module flipper are then rotated 180° to the right, such that the HIC is FPC side down. The vacuum of the red side is applied while the vacuum of the blue side is removed. The blue side is then rotated back to the left, leaving the HIC attached to the red side with the FPC side down.

2. The HIC is transferred to the tab cutting machine using a module gripper.

3. The tab cutting machine, shown in Fig. 5.12, is used to cut the FPC extension and alignment tabs off each HIC. The HIC is moved to the edge of the guillotine with the translation table and a microscope is used to align the edge

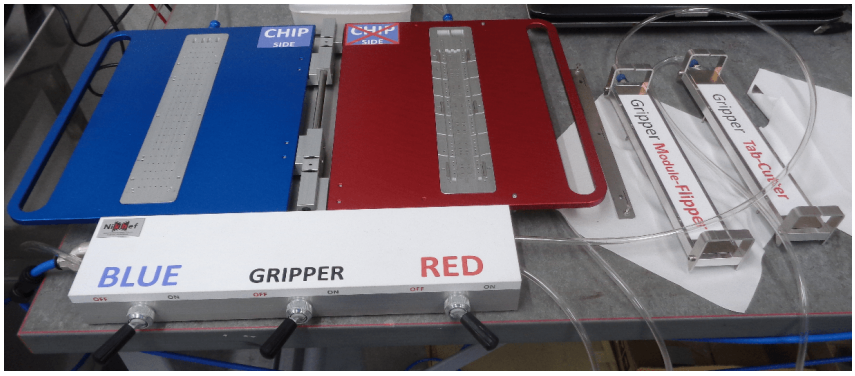

Figure 5.11: Module flipper machine. Each HIC is flipped from silicon side down to silicon side up using a system of rotating vacuum grippers. The 'Gripper Module-Flipper' and 'Gripper Tab-Cutter' seen on the right are used to move the HIC from the module flipper to the tab-cutter.

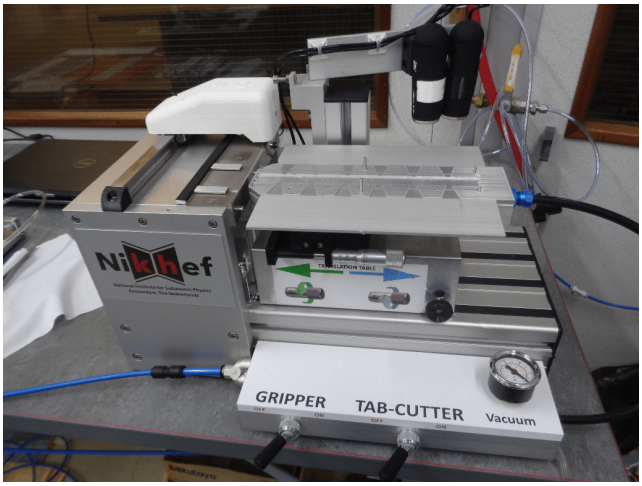

Figure 5.12: Tab-cutter machine. The HIC is placed on the translation table. The FPC extension tab is cut with the white guillotine (left) after the HIC is aligned using a microscope.

of the HIC to the guillotine edge. The silicon chips need to be between 50 and 60 microns from the edge for the FPCs to be properly aligned for the bridge soldering. Once in position, the guillotine edge is moved across the FPC extension.

4. The module gripper is used to move the HIC back onto the red side of the module flipper. A cutting guide is placed over the 3 wings on one side of the HIC. The cutting guide is aligned and held in place with dowels. A scalpel is used to remove the wings. This is repeated for the other side. Finally, the module flipper is used to flip the HIC back so it is FPC side up and the gripper is used to place it back onto the carrier plate.

Yield

A total of 2174 HICs underwent the tab-cutting procedure successfully across all stave construction sites. The yield of the process was 99.54%. Table 5.1 shows the mechanism of failure for the tab-cut procedure of all 10 damaged HICs. Three HICs were lost due to operator error, cutting a B-type HIC as an A-type HIC. The FPC extension of the A-type HICs must be cut leaving a $50-60$ μm gap from the chip edge. The FPC extension of the B-type HICs must be cut leaving a gap of at least 67 μm from the chip edge, in order to preserve the ground connection on the FPC. These three identical failure mechanisms occurred at three separate construction sites, showing it was an easy mistake for the operator to make.

Table 5.1: Mechanism of failure for HICs damaged in the tab-cut procedure.

HIC	Mechanism of Failure
AL002010	Vacuum lost, resulting in cracked chip
AR002859	Wrench dropped onto sensor
BL001550	Operator error - cut as A type HIC
AL001668	Point like crack on chip 6
AL001743	Point like crack on chip 6
BL000622	Operator error - cut as A type HIC
BR000087	Operator error - cut as A type
AR000504	Crack on chip 6
BR000314	3 wire-bonds broken on 2 chips
AL000527	Fast power test failed after cut, no obvious cause

5.3.4 Gluing of the HICs to the Cold Plate

The process for gluing the HICs to the cold plate is summarised below:

1. A glue mask is applied to the cold plate, such that once the HIC is lowered onto the CP, only the silicon is in contact with glue. In the meantime, Araldite 2011 is mixed and allowed to cure for roughly 40 minutes before being applied to the cold plate and glue mask using a pressurised glue gun. A glue scraper tool is then dragged over the glue deposition to ensure the glue application is of uniform thickness.

2. The HIC is lifted from the carrier plate using a module gripper and attached to the alignment station, shown in Fig. 5.13. The alignment station is used to put the HIC in roughly the correct position above the glue deposition by eye. An alignment program on the CMM is run which locates four markers on each corner of the FPC and uses them in conjunction with the cold plate measurement data to align the HIC properly. The alignment station is used to manipulate the HIC until the CMM confirms that the HIC is in the correct position on the cold plate. The HIC is then lowered onto the glue using the alignment station until pressure sensors indicate that it has been lowered enough for a good glue bond to be achieved.

3. Fine alignment tools are attached to the module gripper and base, shown in Fig. 5.14, to allow small adjustments to the HIC position to be made after lowering onto the glue. The CMM is then used to check the HIC position and the fine alignment screws are used to correct for the movement of the HIC during or after lowering onto the glue. Sometimes, if the glue has a low viscosity, the HIC may drift upon it, making the fine alignment screws essential.

4. Once the position of the HIC is within tolerance, the vacuum from the alignment station to the module gripper is removed. The other HICs are then glued using the same procedure. Each module gripper is left on the HIC overnight while the glue cures.

5. Once the glue has cured, the module gripper is removed and the CMM is used to verify and record the positions of the HICs on the cold plate, using markers on each ALPIDE chip.

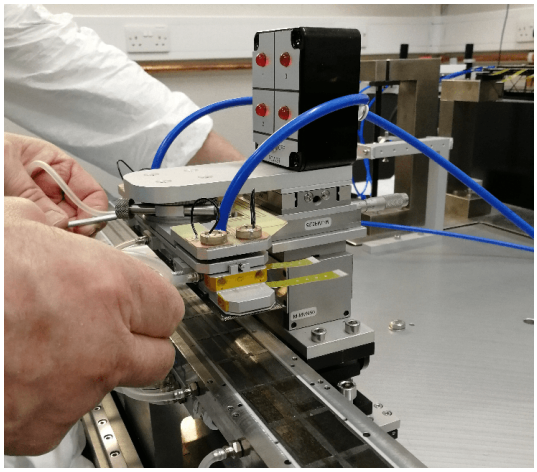

Figure 5.13: Aligning a HIC above the glue deposition on a cold plate. The HIC is held in place via a vacuum and lowered with control of its position in every plane. The black box with four LEDs is a pressure sensor used to determine when the HIC is in place.

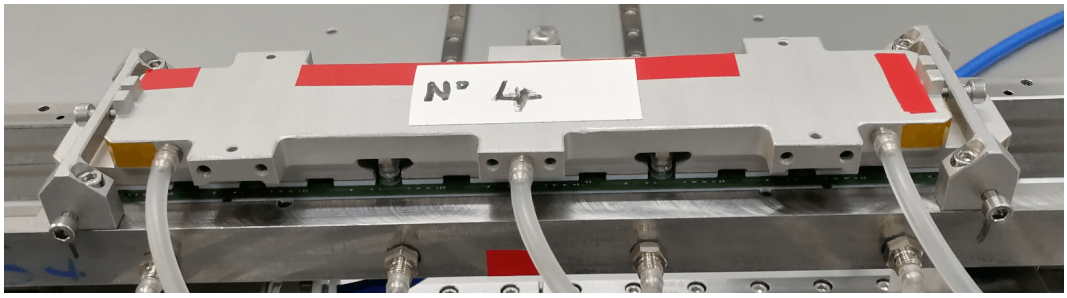

Figure 5.14: A single HIC after lowering onto the CP. The fine-adjustment screws can be seen on either side of the module gripper.

5.3.5 Soldering of the FPC extension and bridges

Once all HICs are glued to the cold plate, the assembly is moved from the CMM to the soldering table where the HICs are soldered together via bridges to allow data transmission. In addition, each HIC has two sets of three address resistors on the FPC. The relevant resistors are removed to give each HIC a unique address. The address resistors give each HIC a number from one to seven. The data lines, CLK and CTRL, will have resistances of around 125 kΩ if all bridges are soldered correctly and all termination resistors are removed properly. A deviation from this prompts another look at the above procedure.

5.3.6 Gluing of the U-arms

Eventually the power and bias buses need to be attached to the half stave. In anticipation of this, plastic U-arms are glued to the HS. Three are glued on the side of each HIC opposite the cross-cables, with the U-arm facing towards the cross-cables, as shown in Fig. 5.15. The power and bias buses will be tucked into the U-arms once the buses are folded over the HICs. This will be discussed in more detail in Sec. 5.3.9.

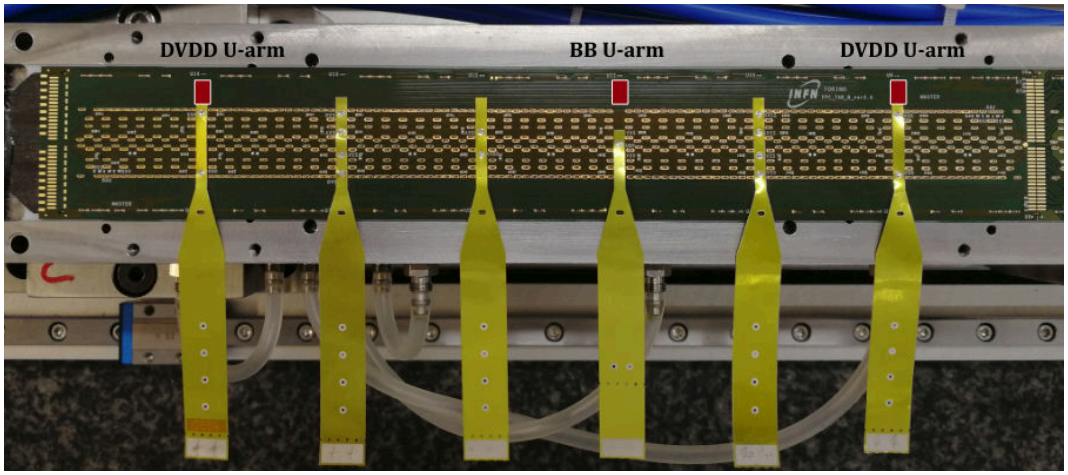

Figure 5.15: Placement of U-arms. Two U-arms are positioned above the DVDD cross cables and one is placed above the BB cross cable.

5.3.7 Gluing of the Half Staves to the Space Frame

The space frame sits on a set of rubies on the CMM. The HS needs to be flipped, such that the FPC side is facing down, before gluing to the space frame. This is achieved using the HS handling bar which is a set of vacuum pads covering the length of the whole HS. This is lowered onto the HS, and then the vacuum is switched off for the HS base and on for the handling bar. A metrology check is done after lowering the handling bar and switching the vacuums to see if the position of the HS alters due to the switching of the vacuums. Then, the handling bar is used to flip the HS and lower it onto an alignment station. Using the alignment station, the HS is moved underneath the SF and rotated and shifted in X and Y so that it is in reasonable alignment with the SF by eye. Then, a HS alignment program on the CMM is used which uses the FPC holes and known SF position data to properly align the HS below the SF. First it is aligned in the XY plane before being raised up to the SF. Adjustment grub screws are used at either end of the HS as well as the alignment station to support the HS in the Z axis. By repeating the alignment program multiple times and adjusting the height of the HS, the correct alignment can be achieved. Once aligned properly, the HS is glued to the SF at each SF U-leg, as shown in Fig. 5.16. After the glue has been left to harden for 12 hours, a final metrology scan is performed. The procedure is repeated for the second half stave.

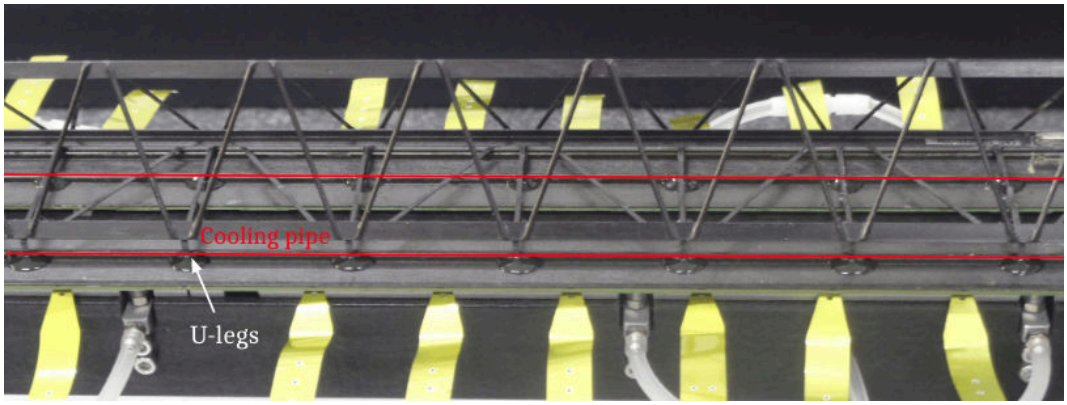

Figure 5.16: Gluing a HS to the SF. The U-legs on the SF are glued to the sides of the cooling pipes on the CP.

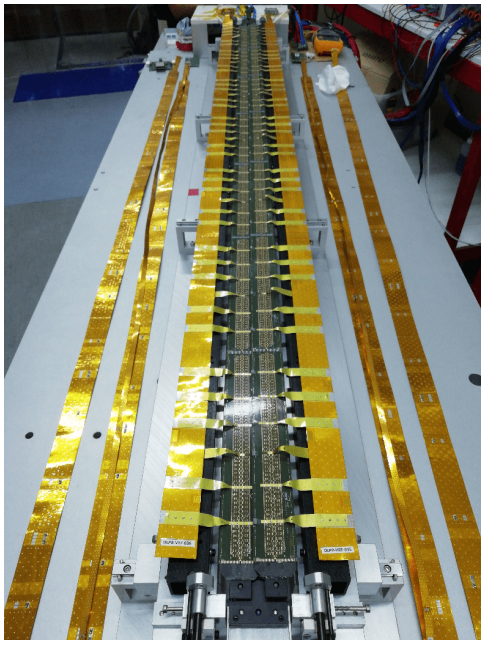

(a) A full stave with the power and bias buses soldered but not folded.

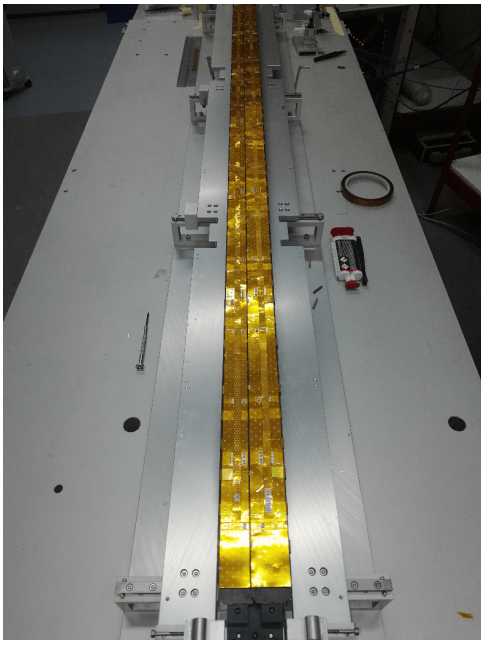

(b) A full stave with the power and bias buses folded.

Figure 5.17: One full stave before and after folding the power and bias buses.

5.3.8 Power and Bias Bus soldering

The stave is moved to the power bus folding stand. Once on this, the power and bias buses are soldered to the cross cables. Each power bus has individual aluminium lines to DVDD and AVDD of each HIC which are insulated with kapton. The power bus provides a common ground. The bias bus has a similar structure to the power bus, providing an aluminium bias line to each HIC, insulated with kapton. The power and bias buses are soldered to a filter-board which plugs into a breakout board and connects to the power system via copper wires.

5.3.9 Power and Bias Bus folding

The power and bias buses are folded over the top of the HICs and into the U-arms which were glued to the FPCs previously using a folding tool, as shown in Fig. 5.17. The folding tool uses vacuum pads on a manoeuvrable bar to adhere to the buses. The bar is then folded over the top of the HICs before being lowered down and pushed forwards such that the buses slide into the U-arms. The capacitors on the FPCs stop the power bus from making contact with the wire-bonds. It is very im-

portant to ensure the power bus is as flat as possible before folding to prevent kinks in the power bus from being able to touch the wire-bonds in between capacitors. Once the buses are folded, the wire-bonds are less easily damaged.

Once folded, more U-arms are glued to the outside edge of the HS to hold the power and bias buses in from both sides.

5.3.10 Shipment

OL staves are packaged into aluminium boxes which are loaded into vacuum bags, with silica gel. The aluminium box and vacuum bag are then lowered into a wooden box, onto a spring isolated support, designed to reduce vibration of the stave during transit. The wooden box is then loaded into a temperature controlled van and driven to CERN.

The ML staves were exclusively manufactured in Berkeley, so the transport to CERN involved travel by plane. Each of the ML staves was packaged in a clear perspex box and carried by a member of the Berkeley team via plane in an extra passenger seat [1].

[1] The project leader at Berkeley, Leo Greiner, claimed this was the 'most fantastic outreach' he had ever done[88].

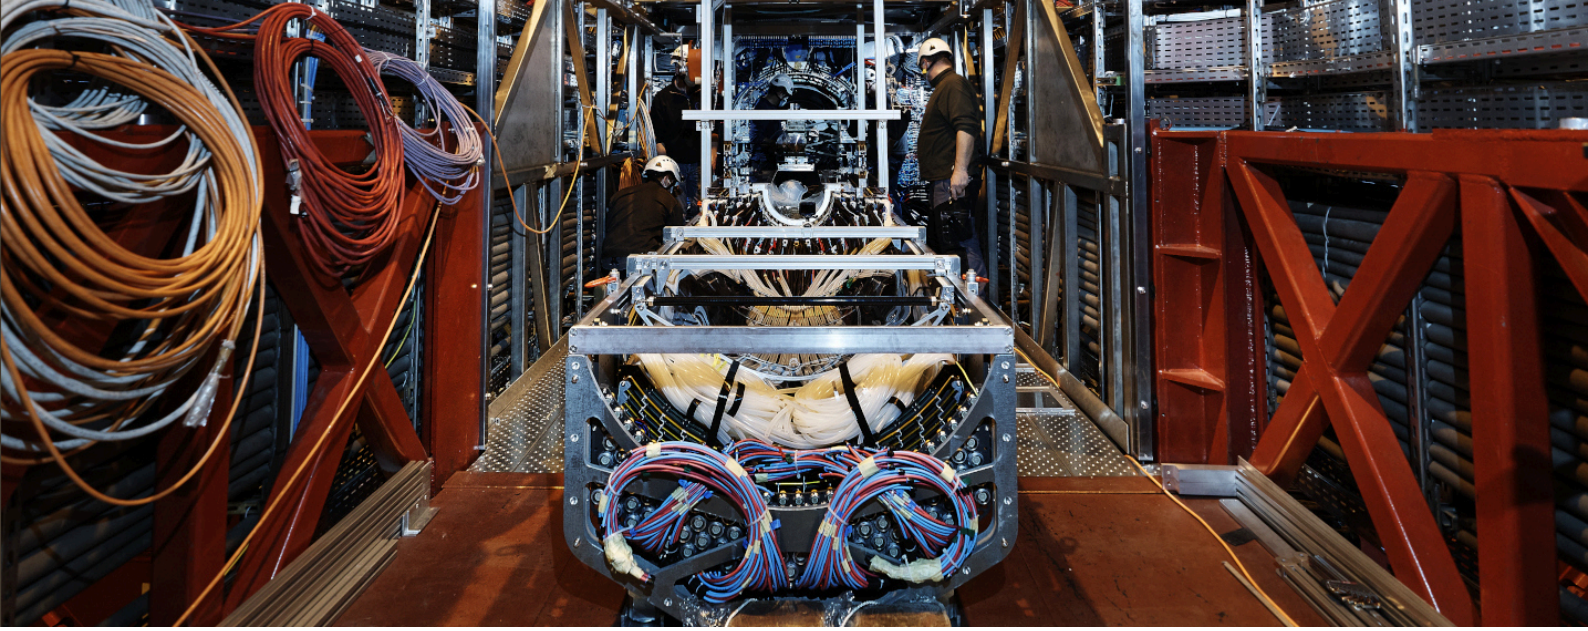

Chapter 6:

Stave characterisation

Section 6.1 will describe the metrological requirements the staves must satisfy in order to fit within the ITS2 barrel, as well as the metrological characterisation of the staves. Section 6.2 will describe the electrical characterisation of the staves.

6.1 Metrological characterisation

The staves will eventually be installed into the barrel geometry of the ITS2. As shown in Fig. 6.1, the staves will overlap, where the closest distance to one another is $O(1\,\text{mm})$. It is necessary to note, however, that each stave is not defined rigidly in space due to the flexibility of the power and bias buses. The tolerance on the absolute displacement from the measured chip marker to the nominal chip marker was decided to be 500 µm to ensure no staves touched once assembled into the barrel geometry.

A final metrological survey is performed by the CMM to determine the position of the HICs and chips on the stave. The procedure provides a valuable assessment of the precision of the placement of HICs onto the cold plate. The results will be used as starting parameters for the alignment of the detector with cosmic muons.

Each ALPIDE chip has four markers, one in each corner, which are found by the CMM and their positions measured. During the HS metrology, all these markers are

Header photograph shows the bottom half of the Outer Barrel on the ALICE mini-frame, en route to the ALICE detector. © CERN. Taken from [26].

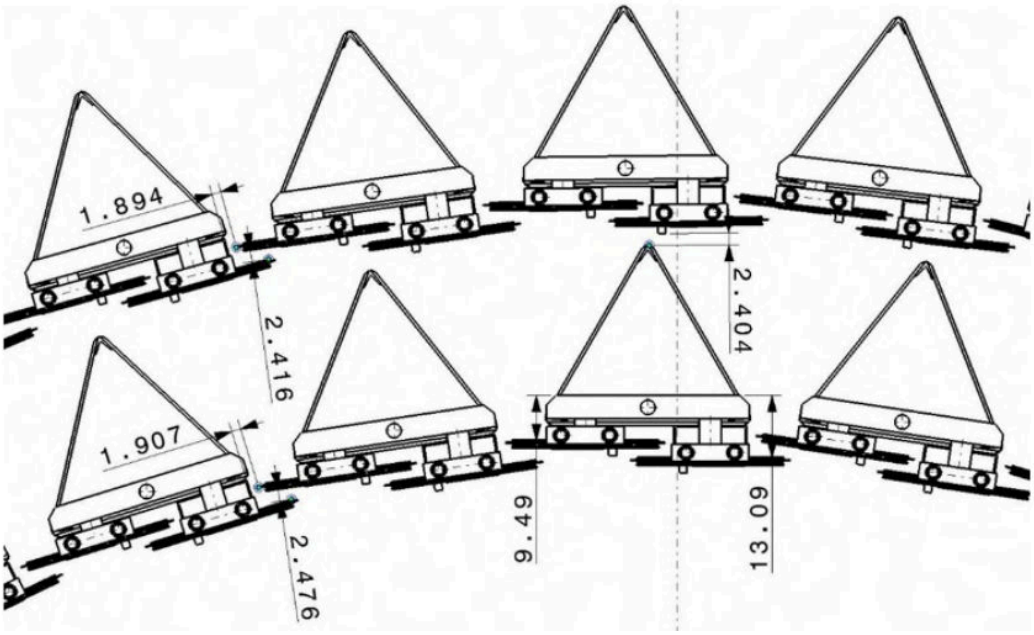

Figure 6.1: Cross-sectional view of some staves in layers 5 and 6. The overlap of two adjacent half staves from two separate staves is visible. The closest distance between these two half staves touching is 1.894 mm for layer 6 and 1.907 mm for layer 5.

visible. However, during the stave metrology, three markers per HIC are covered by U-arms on Half-Stave Lower (HS_L), whilst all markers on the inner edge of Half-Stave Upper (HS_U) are covered by HS_L. Therefore, for the OL staves, 56 and 154 markers are visible for HS_U and HS_L respectively. Whilst for the ML staves, 28 and 88 markers are visible for HS_U and HS_L respectively.

6.1.1 Nominal marker residuals for the OB

The difference from nominal marker position for x,y and z can be seen in Fig. 6.2 for all OB staves, fitted with a Gaussian using MINUIT [89]. The mean x residual was $(-5 \pm 0.3)\,\mu$m with a standard deviation of $(58.8 \pm 0.3)\,\mu$m, whilst the mean y residual was $(-11.6 \pm 0.2)\,\mu$m with a standard deviation of $(34.6 \pm 0.2)\,\mu$m. The standard deviation for the y residual distribution is almost half of that of the x residual distribution. This is likely to be due to the mechanical process of gluing the HICs to the CP. The fine-adjustment tools were more likely to allow the HIC to drift in the x axis than in the y axis before the glue had set. In both dimensions, the mean marker residual was shifted by $O(10\,\mu\text{m})$ in the negative direction. This offset could have a number of origins, but since the sample includes several different

construction sites and technicians, the offset is likely to be due to shared machinery rather than assembly characteristics, such as glue drift. The x residual will affect how much staves overlap once installed in the barrel geometry. The overlap of adjacent staves is $O(1\,\text{mm})$, shown in Fig. 6.1, so the offset of the x residuals of $O(10\,\mu\text{m})$ is not a cause for concern.

The mean z residual was $(68.5 \pm 0.5)\,\mu\text{m}$ with a standard deviation of $(92.4 \pm 0.4)\,\mu\text{m}$. The mean z residual indicates the staves were generally thicker than nominal in z. The uncertainty on the z position is $10\,\mu\text{m}$ rather than the $2\,\mu\text{m}$ in x and y due to the resolution being dependent on the camera focus.

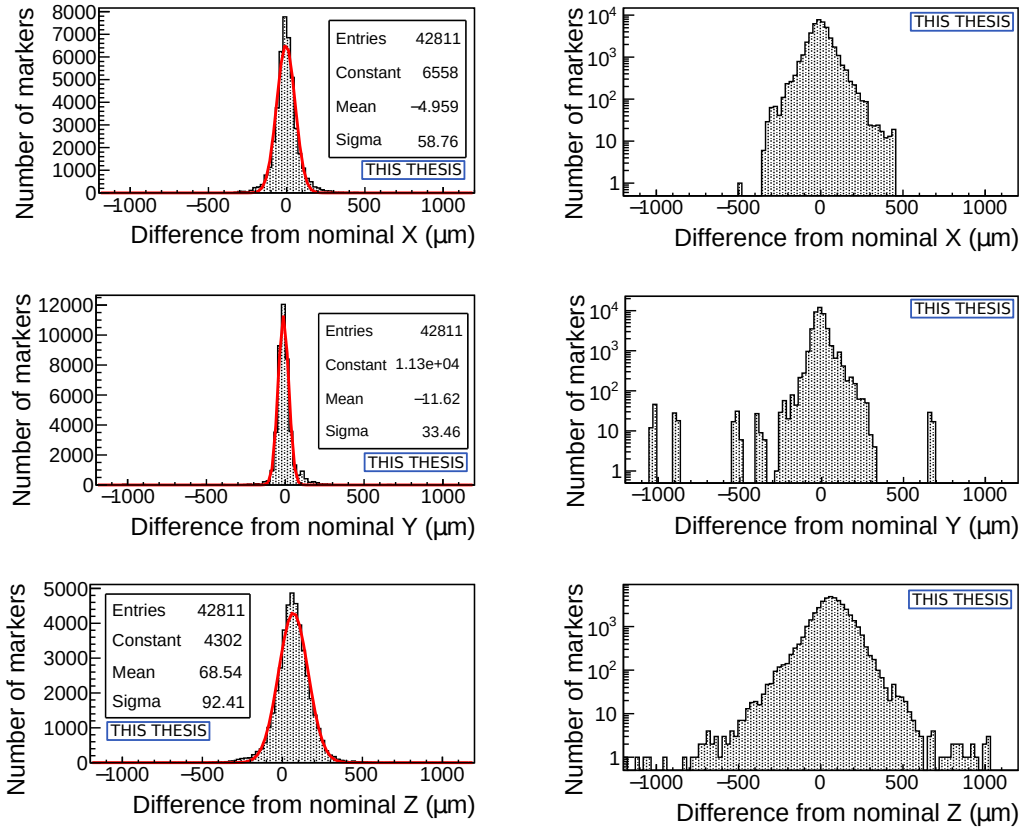

Figure 6.2: Difference from nominal x, y, z position and actual x, y, z position for all chip markers on all OB staves. The left column shows the results of a Gaussian fit. The right column shows the distribution on a logarithmic y scale to see the outliers.

Figure 6.3 shows the absolute displacement from nominal for all OB staves. The position of every marker is within the required $500\,\mu\text{m}$ tolerance.

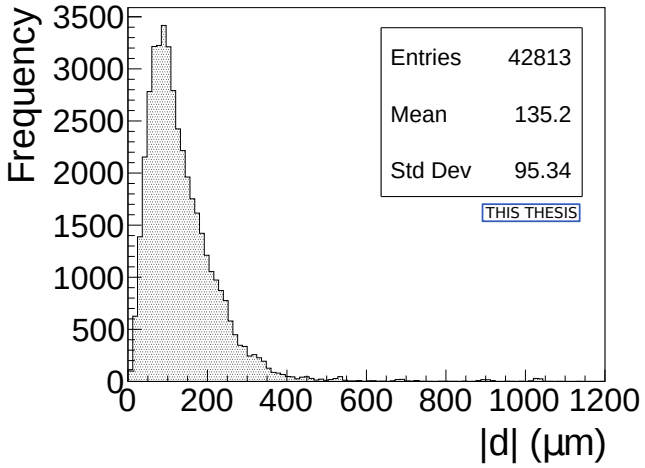

Figure 6.3: The absolute distance from nominal marker position to measured marker position for all chip markers on OB staves measured with the CMM.

6.2 Electrical characterisation

Section 6.2.1 will describe the electrical characterisation systems, whilst Sec. 6.2.3 will describe the specific tests conducted.

6.2.1 Overview of test system used during construction

Each test system used during construction incorporates a MOSAIC board and a power-board.

The MOSAIC board

The MOdular System for Acquisition Interface and Control (MOSAIC) board was designed to test ALPIDE chips, modules, and half-staves during the construction stage of ITS upgrade. The MOSAIC board was not designed to operate in a high radiation environment.

The MOSAIC board is built around a single FPGA[1], for which a block diagram is shown in Fig. 6.4. Two sets of 33 Low Voltage Differential Signal (LVDS) signal pins (plus one ground reference) are provided by two Robinson Nugent P50E-068-P1-SR1-TG connectors, housed in two FPGA Mezzanine Card (FMC) slots. The LVDS pins communicate via I^2C to a power-board, described in Sec. 6.2.1. High

[1] Xilinx XC7A200T-2 FFG1156C

speed I/O communication is handled by a Samtec Eye Speed HDI6-035 connector, providing 10 high-speed transceivers, running up to 6.6 Gb/s, for data transfer. Moreover, the Samtec connector provides two LVDS, used for data transfer, and two Multipoint Low Voltage Differential Signal (M-LVDS) bidirectional differential lines, used to monitor ALPIDE front-end electronics and provide control and pulse signals, as well as a 40 MHz clock. Control of the MOSAIC board, as well as data transfer from the MOSAIC board, is handled via Ethernet connection.

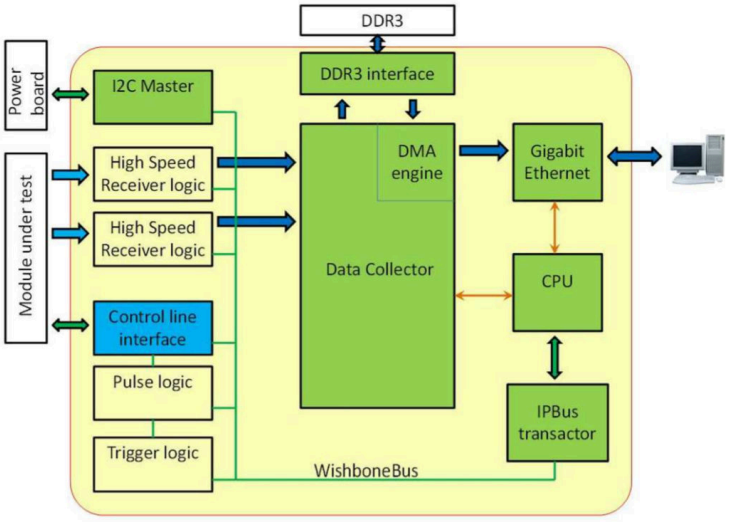

Figure 6.4: Block diagram of the FPGA firmware in the MOSAIC board. Taken from [90]. See text.

Data are transferred from the high-speed receivers to block buffers. A read request is flagged to transfer these data to the DDR3 memory at a sustained rate of 5 Gb/s. Each data block contains a header that describes the block, noting its size, source and flags from the source. Off-board communication protocols are handled by the on-board CPU, whilst data transfer is handled by the Direct Memory Access (DMA) engine in the data collector. Data is transferred off-board via Ethernet at 120 MB/s using 30% of the CPU load.

The wishbone bus handles the control and monitoring of the ALPIDE chips. User Datagram Protocol (UDP) and IP transactions are received via Ethernet from the PC and translated to wishbone commands by the IPBus Translator.

The power board

The power board can power up to 16 HICs via 32 1.8 V connections with a drive strength of 3 A, one analogue and one digital per HIC. Eight negative voltage outputs with a low current provide bias from a common source. A single power board is split into two power units, which contain 16 and 4 power and bias connections each. Between the power units, there is a water-cooled heat exchanger connected to an external cooling system. Each power unit is controllable via I^2C interface. Every power and bias connection can be individually enabled. Each power supply can be individually adjusted, whilst each bias supply can be adjusted on a power unit basis. Voltages, currents and temperatures can all be monitored. Over-current protection has an adjustable threshold on each power supply and over-temperature protection is provided on a power unit basis. A temperature measuring circuit is present on each power unit, allowing the readout of an external PT100 sensor. The PT100 sensors are only installed onto each HS once the stave is received in CERN however.

An over-current condition on one power channel results in the corresponding channel being switched off. A red LED near the channel regulator is illuminated to indicate the condition has been met. An over-temperature condition results in all channels on the corresponding power unit being switched off. A red LED located at the rear of the power unit indicates the over-temperature condition has been met. When no over-current or over-temperature condition has been met, and the power board is supplying power to the HICs, a green LED is illuminated next to each channel regulator.

The power board receives power from a variable power supply, in the case of the construction sites, and from a VME power crate, in the case of the final detector power system at CERN. Eight 3.3 V connections provide power to the power supply channels and one -5 V connection provides power to the bias channels. Each power unit has one I^2C connection and one breakout board connection. The breakout board provides power and bias to the detector, via 16 and 4 channels respectively. Two differential channels connect to the PT100 sensor, whilst one connection provides an infrastructure ground, connected to the carbon fibre space frame of the stave with kapton tape.

There are two operating modes of the power board, controllable via a switch, IB and OB mode. In IB mode, the second power unit is powered via the first power

unit, reducing the number of power connections needed to the power supply, four instead of eight. In this case, each power unit can power one IB stave. In OB mode, each power unit is powered independently, necessitating the use of eight input power connections. For OL staves, one power unit provides power to all 14 HICs in a HS, with two channels left over. For an ML stave, one power unit is enough to power both HSs, where each HS consists of eight HICs.

6.2.2 Specific test systems

Four individual test setups, incorporating the MOSAIC and power boards, are used to characterise HICs and staves. The test setups are listed below. The HIC test and endurance test systems were used at HIC construction sites. The fast HIC test and HS test systems were used at stave construction sites. The HS test system was also used at CERN, after the reception of shipped staves. An overview of which test types are performed during each test routine is shown in Table 6.1.

Table 6.1: Test types performed per test system.

Test type	HIC test	Endurance test	Fast test	HS test
IV scan	✓	✗	✓	✓
Register read-back	✓	✓	✓	✓
Power	✓	✓	✓	✓
DCTRL	✓	✗	✓	✓
FIFO	✓	✓	✗	✓
Digital	✓	✗	✗	✓
White frame	✓	✗	✗	✓
Threshold	✓	✗	✗	✓
Threshold tune	✓	✗	✗	✓
Noise occupancy	✓	✗	✗	✓

HIC test system Each HIC is tested individually within a HIC test box. The HIC test box ensures the HIC is not exposed to light during testing. The HIC is connected to the power-board via a power connector. This screws into the HIC carrier plate such that the spring-actuated connection pins make contact with the cross-cables. The Samtec eye-speed cables connect to the FPC extension of the HIC. The test setup in the LSDC is shown in Fig. 6.5a.

Endurance test system for HICs The endurance test crate for HICs is similar to the HIC test system, however, it is capable of holding ten HICs at once in two rows of five. The test setup in the Liverpool Semiconductor Detector Centre (LSDC) is shown in Fig. 6.5b.

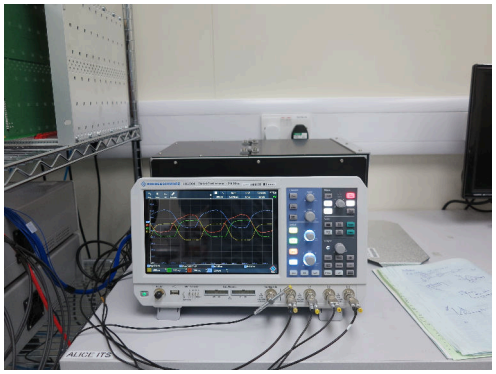

(a) An oscilloscope sits in front of the HIC test box in the LSDC. The oscilloscope is used for the DCTRL test.

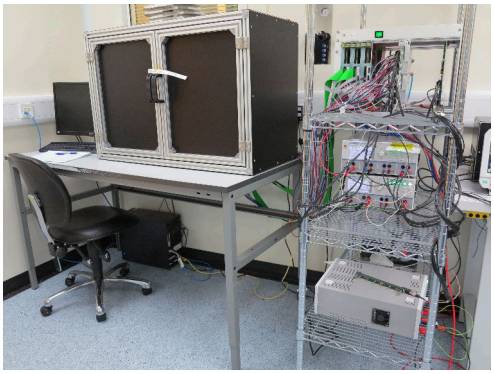

(b) The endurance test crate in the LSDC. The MOSAIC boards can be seen in the rack on the right of the photograph.

Figure 6.5: The HIC (a) and endurance (b) test setups in the LSDC.

Fast HIC test system The fast test system is used at stave construction sites, to quickly test a HIC after reception of that HIC, and after the tab cut procedure. The HIC is powered with the power connector and no data transmission lines are connected. Only tests to check the integrity of the powering of the HIC are performed.

Half-stave test system The HS test system is used for testing a HS without a power bus soldered, with a power bus soldered but not folded and with a power bus soldered and folded. The system is similar in all three cases. For the case where the power bus is not soldered, a power connector is used to power the HS. This power connector is similar to the one used for the HIC tests, however, it is simply longer to account for multiple HICs. When the power bus is soldered, the power is supplied via the power bus and breakout board. The breakout board connects to the power board.

6.2.3 HIC and stave tests

This section will describe the tests performed on each HIC and the corresponding grade given. Tests described in this section are described in chronological order.

Where applicable, test results extracted from the Kybernetika database are shown.

IV scan

An IV curve is plotted for each voltage channel: AVDD, DVDD and BB. A low impedance ($< 100\,\Omega$) indicates an electrical short between the voltage channels. A short can be further investigated using a thermal camera. Damage to the silicon can cause an electrical short and would result in an area of higher than normal temperature (a hotspot), visible with a thermal camera, as shown in Fig. 6.6.

For AVDD and DVDD, a range of 0 - 0.2 V in 20 steps of 0.02 V is scanned. For BB, a range of 0 - 4 V in 50 steps of 0.08 V is scanned. The channels have a compliance of 100 mA. All channels turn off if compliance is reached on any single channel. The resistance at each measured point is calculated using Ohm's law. The impedance of a particular channel is calculated as the sum of resistances for each non-zero current value, divided by the total number of measured points. A HIC fails the scan if any channel has an impedance less than $100\,\Omega$.

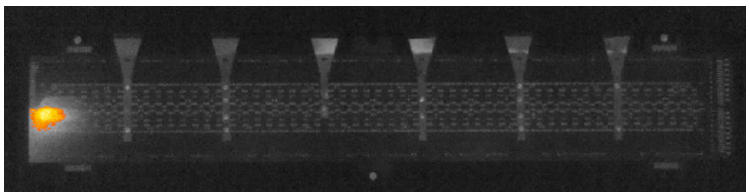

Figure 6.6: A powered HIC with damage to the silicon resulting in a hotspot visible using a thermal camera.

Initial register readback

While the HIC is powered, the IBIAS register of each chip is written to and read back from to test that each chip is correctly powered and DCTRL communication is possible. If the IBIAS register is set successfully, the current drawn by the HIC will increase. Knowing this, it is possible to determine if a chip has fully functioning DCTRL communication, just write, or just read.

Power test

The power test is started with the HIC left already powered from the initial register read-back test. The HIC is then powered down and the clock is stopped. Next, the HIC is powered, the clock turned on and the chips configured. The voltage drop

to the chips is corrected after each of these steps while the power consumption is recorded. Finally, an IV curve of the back bias is produced over a BB range 0 - 4 V. If the back bias current exceeds 15 mA, the HIC is classified as back bias not working.

DCTRL test

The DCTRL test measures the strength of the DCTRL drivers and ensures the driver strength increases linearly with DAC setting. The scan is performed on chips that transmit data to the off-detector electronics, therefore all chips on an IB HIC and the master chips on an OB HIC.

For each of the 16 driver settings, a measurement of the peak-to-peak, amplitude, rise and fall time are carried out. The rise and fall times are defined as the time taken to go from 20% to 80% of the peak height and vice-versa.

FIFO scan

The FIFO scan checks the integrity of the high-speed links. 16 bit wide FIFO buffers are located in the periphery circuitry of the ALPIDE chip which store the addresses of the hit pixels. The FIFO scan inputs four 16-bit numbers: 0x0000 (0000 0000 0000 0000), 0x5555 (0101 0101 0101 0101), 0xaaaa (1010 1010 1010 1010) and 0xffff (1111 1111 1111 1111) into the 16 bit FIFO buffers. The four hex numbers are then read out and compared to the input.

There are two modes of failure for this scan. In one, named a FIFO exception, the MOSAIC cannot read or decode the data sent from the chip. These typically indicate a communication problem possibly due to the setup or a weak DCTRL link of one of the master chips. The other failure mode, known as a FIFO error, occurs when the read-back is completed but the value differs from the written one, indicating a problem with a single chip. The entire FIFO scan is repeated with $\pm$10% of the nominal supply voltage. No HICs were found to show FIFO errors or exceptions after debugging any connection issues.

Digital scan

The digital scan injects digital hit signals into the front-end electronics of each pixel (i.e. the MEBs) and counts the number of hits recorded to test the responsiveness of

the electronics. 50 injections are performed over one row of pixels at a time. With this the following errors may be seen:

- 8b10b errors: counts the number of times MOSAIC flags a corrupt 8b10b word. Occurrences are most likely due to non-optimal driver settings.
- Timeouts: one event per ALPIDE chip is expected for each trigger sent by the MOSAIC. If the expected number of events is not received within a given time frame, a timeout is flagged by the MOSAIC. Timeouts can occur for several reasons: data corruption on the high-speed link (for example alongside 8b10b errors), triggers not received by the chips (for example if the chips were not successfully configured) or single chips blocking the readout of a HIC (for example a loud double column not masked correctly).
- Corrupt events: if an event is received correctly but cannot be decoded.
- Priority encoder errors: either a single pixel sent its address twice or the pixel addresses appear in the wrong order. This is caused by a fault priority encoder and is often fixed by masking the single contributing pixel.

No HICs had any timeouts or corrupt events. Figure 6.7 shows the number of bad pixels per HIC. Bad pixels are defined as pixels that are either noisy (have a number of hits greater than the number of injections), dead (do not respond to injections) or inefficient (have a number of hits lower than the number of injections). The mean number of bad pixels per HIC for all HICs was ≈ 5000, 0.07% of all pixels on a HIC.

Threshold tuning

As described in Sec. 4.4.3, each pixel within the ALPIDE chip has a capacitor, C_{inj}, capable of injecting a test charge into the sensing node, where the magnitude of the test charge is controlled by the V_{PULSE} node. V_{PULSE} is generated by internal DACs in the chip and gives a resolution of 256 steps over V_{REF} of 1.8 V, meaning augmentation of voltage step varies the test charge by $10\,e^-$.

Each pixel has a comparator circuit that compares the charge collected in the sensing node to a threshold value preset on a chip-by-chip basis. This threshold value is influenced by two on-chip DACs, V_{CASN} and I_{ITHR}. V_{CASN} has an exponential relationship with the threshold value and I_{THR} has a linear relationship.

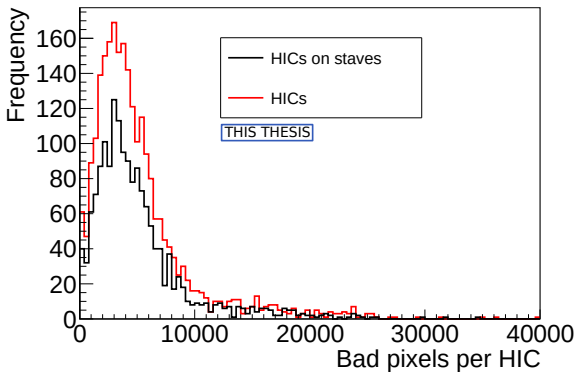

Figure 6.7: Bad pixels per OB HIC measured during the digital scan. HICs selected for staves are shown with a black line.

V_{PULSE} is chosen to pulse the active volume of the silicon with a charge equivalent to the desired threshold value. In the case of the MOSAIC scans, the desired threshold was $100\,e^-$. To tune the chip thresholds, first V_{CASN} is varied whilst I_{THR} and V_{PULSE} are kept constant. The optimum V_{CASN} value is chosen as the value for which the pixel firing probability is closest to 50%. Next, I_{THR} is augmented, whilst holding V_{CASN} and V_{PULSE} constant. Again, the most appropriate value of I_{THR} is where the pixel firing probability is closest to 50%.

Threshold scan

The purpose of the threshold scan is to verify the threshold values and to get a measure of the temporal noise. The threshold scan inputs 50 test charges, Q_{INJ}, of the same magnitude via C_{INJ}, for a range of V_{PULSE}, typically 50 voltages steps from $0\,e^-$ to $350\,e^-$. Once the charge injected into a pixel surpasses the threshold of the chip, the pixel will register a hit. For each value of Q_{INJ}, the number of times a pixel registers a hit, N_{HIT}, and the number of total charge injections, N, is measured. The response function, $r(Q_{INJ})$ of the front-end circuit is:

$$r(Q_{INJ}) = \frac{N_{HIT}}{N}. \tag{6.1}$$

In theory, $r(Q_{INJ})$ should be a step function, however, temporal noise produces an S-curve rather than a step function. For a functional pixel, $r(Q_{INJ})$ is described by the error function, $f(Q_{INJ})$:

$$f(Q_{INJ}) = \frac{1}{2}\left[1 + erf\left(\frac{Q_{INJ} - Q_{THR}}{\sqrt{2}\sigma}\right)\right], \tag{6.2}$$

where σ describes the temporal noise and Q_{THR} is the charge threshold, defined as the value of Q_{INJ} for which a pixel has a firing probability of 50%. Figure 6.8 shows the average σ per chip at a threshold of $100\,e^-$. The mean noise for all HICs was $5.7\,e^-$ with a standard deviation of $0.29\,e^-$. For HICs that were installed onto staves, the mean noise was $5.42\,e^-$ with a standard deviation of $0.26\,e^-$.

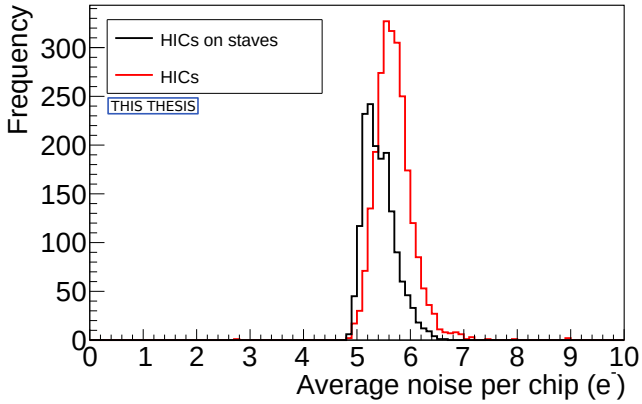

Figure 6.8: Average temporal noise per chip measured for all OB HICs during the threshold scan. HICs selected for staves are shown with the black line.

Given the epitaxial layer thickness for ALPIDE of $d = 25\,\mu$m, the average ionisation energy in silicon of $I_0 = 3.62\,$eV and the mean energy loss per flight path of a Minimum Ionising Particle (MIP) of $\frac{dE}{dx} = 3.87\,\text{MeVcm}^{-1}$, the mean number of charge carriers ionised by a MIP can be calculated as $\frac{dE/dx \times d}{I_0} = 2673\ e^-h^+$ pairs, much larger than the desired threshold of $100\,e^-$. A low threshold is possible because of the low noise.

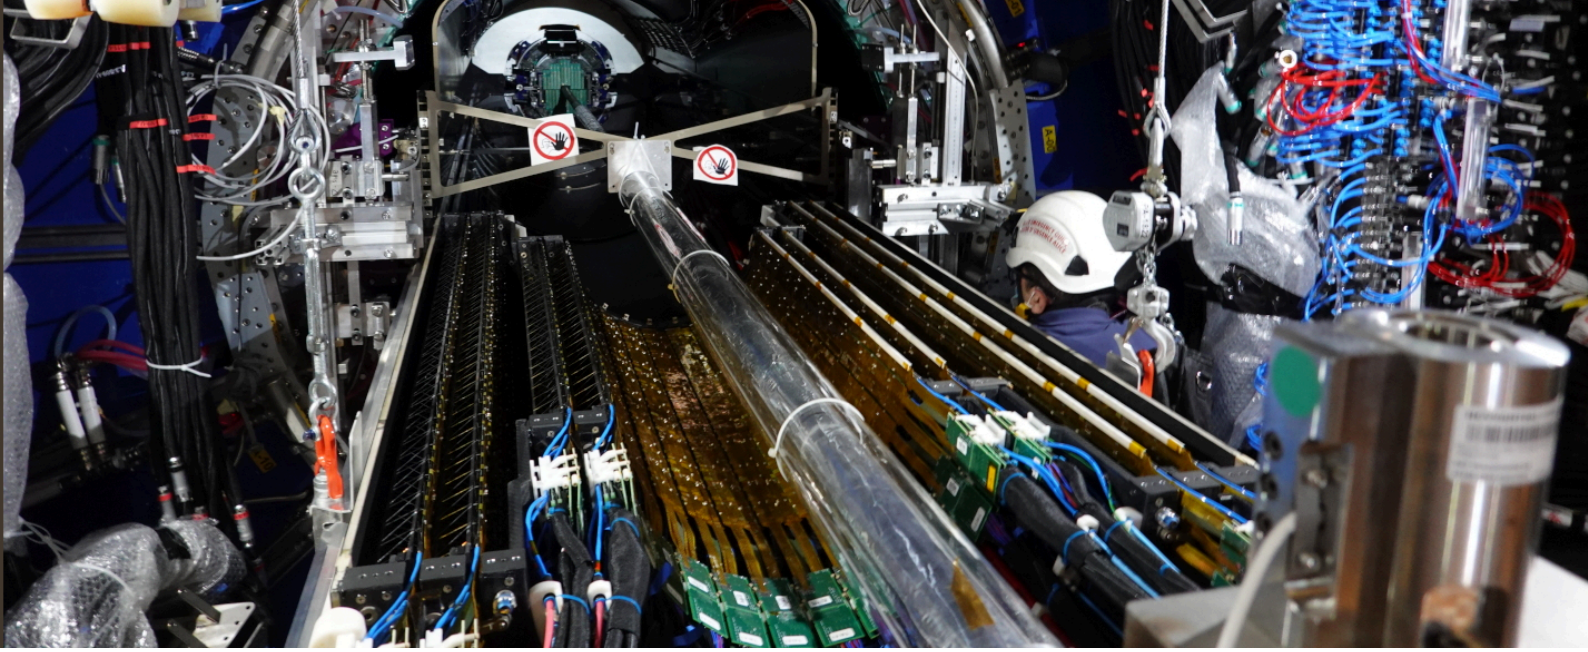

Chapter 7:
ITS readout system

After the reception test conducted at CERN, each stave was integrated into the half-barrel support structure as well as the final services, including the cooling and readout systems.

One FPGA-based Readout Unit (RU) connects to every stave, i.e. 192 RUs in total. Every layer uses identical RUs, where only the configuration changes between the IB and OB. An RU connects to a stave via Twinaxial cables with a total length of up to 7.6 m. The length of data transmission cables to each layer in the ITS2 are shown in Tab. 7.1. The copper cables provide 40 MHz clock and control (DCTRL), as well as unidirectional, 400 Mb/s (OB), 1.2 Gb/s (IB), data lines.

Table 7.1: Cable lengths for each layer in the ITS2. Each cable is split into two parts for detector integration purposes.

Layers	Detector side (cm)	Readout side (cm)	Total (cm)
IB	265	495	760
ML	245	425	670
OL	215	520	735

For the IB, one copper cable connects an RU to a stave, whilst for the OB, four copper cables connect an RU to a stave. For the IB, the copper cable is segmented

Header photograph shows the bottom half of the Outer Barrel before being inserted into the centre of the TPC. © CERN. Taken from [91].

into nine data lines, one clock line and one control line. For the OB, each of the copper cables provides one clock and one control line and four data lines for the ML and seven data lines for the OL. Each master has its own data line.

Each RU conducts the following tasks for its respective stave to via the Main FPGA:

- Provides clock and control to the ALPIDE chips.
- Sends triggers to the ALPIDE chips.
- Receives data from the ALPIDE chips and indicates the quality of those data.
- Controls the powering of the staves HICs via a Power Board (PB).
- Monitors the currents and voltages of the stave, as well as the temperatures of the RU and the PT100 sensors located on the stave.
- Provides access to the GBTx chip registers via Inter-Integrated Circuit (I^2C).

7.1 Readout system architecture

Figure 7.1 shows a block diagram of the Readout system architecture. The RUs, PBs and staves will be installed within the cavern, with the PUs and PBs in the less radiative environment away from the detector ($< 10\,\text{krad}$). The counting room, where there is only atmospheric radiation, houses the Common Readout Unit (CRU), First Level Processor (FLP), Local Trigger Unit (LTU), Event Processing Node (EPN), Central Trigger Processor (CTP), as well as the CAEN RU and stave power supplies. The CTP provides triggers to the LTU which provides triggers to the CRU and RU. The CRU sends control signals to and receives data from the RU via GBTx links. Data received by the CRU from the RU is shipped to the FLP, in which the CRU is hosted. Data is stored directly on the FLP in Offline mode, or transferred to the EPN in online mode. For the commissioning tasks discussed in the next chapters, Offline mode is used exclusively. The CAEN RU and Power Board (PB) power supplies orchestrate the powering of the PB's and RU. The RU controls the PB, specifying the voltage, currents and switch on/off time. Over-currents and over-temperatures in the PB cause the RU to cut power to the affected PB channel.

Two FPGAs are present in the RU. The Static Random-Access Memory (SRAM) based Main FPGA is responsible for communication with the ALPIDE chips and the CRU. The flash-based auxiliary FPGA is responsible for providing the initial

configuration to the Main FPGA on power-up and the repair of the Configuration Random-Access Memory (CRAM) of the main FPGA in the event of radiation upsets [92]. The SRAM based Main FPGA is more prone than the flash-based auxiliary FPGA to Single Event Upsets (SEU) in the CRAM but offers the data throughput required by the ITS, whilst the auxiliary flash-based FPGA can perform in a radiation field without SEU to the CRAM [93].

The RU houses three GBTx chips [94]. The GBTx chips provide reliable communication with the back-end in a high radiation environment. The GBTx0 receives Single Word Transactions (SWT) from the CRU on its downlink. SWT is a slow-control protocol used to control the Main-FPGA and its interfaces on the detector side. For instance, SWT is used to configure the PU and to send DCTRL messages to the detector. The GBTx0 transmits answers to the SWT and detector data to the CRU on its uplink, where the two are multiplexed by the Main-FPGA. The GBTx0 also controls the GBT-Slow Control Adapter (GBT-SCA) chip, which communicates with the auxiliary FPGA via I^2C. The GBTx1 transmits detector data to the CRU on its uplink, whilst its downlink is unused. The GBTx2 receives timing and trigger information form the LTU on its downlink and transmits detector data to the CRU on its uplink. The Versatile TransReceiver (VTRx) and Versatile Twin-Transceiver (VTTx) provide electrical-optical signal translation. One VTRx serves the GBTx0 chip for the uplink and downlink. One VTTx serves the uplink for the GBTx1 and GBTx2 chips. One Versatile Receiver (VRx) receives trigger information for the GBTx2 chip.

7.2 ALPIDE to RU data format

Each ALPIDE chip has two data transmission ports: a serial transmission port and a parallel data port. The data format for both ports is identical. In the OB, slaves transmit data over the parallel port to their respective master. The master samples the parallel port and forwards the data via its serial transmission port at 400 Mb/s. The valid data words are shown in Tab. 7.2, whilst the meaning of each data word is outlined below.

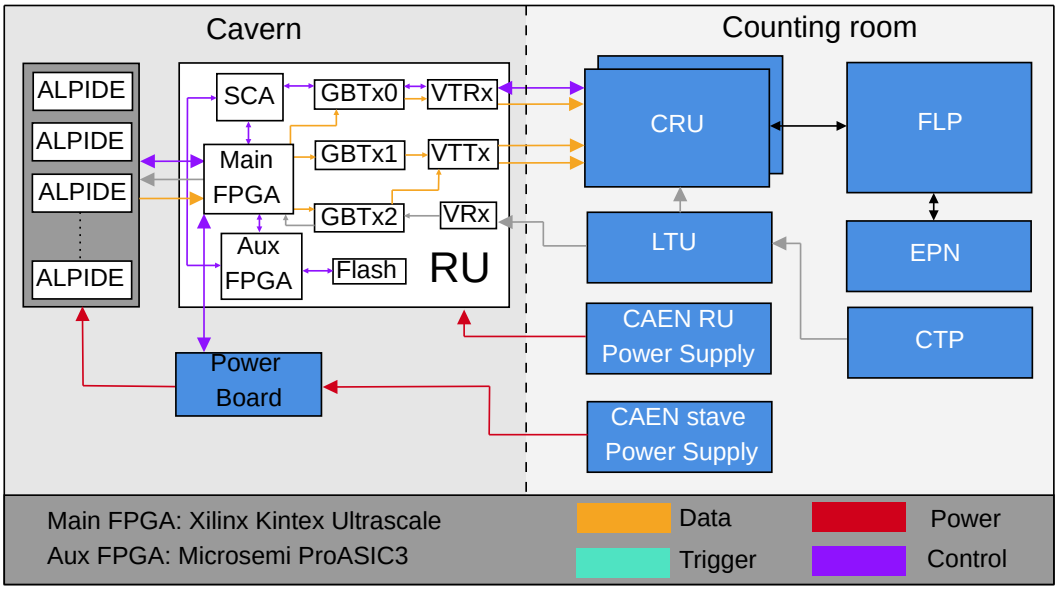

Figure 7.1: Block diagram of the readout system architecture [92], [95]. See text.

Table 7.2: Valid data words transmitted by the ALPIDE chips [96]. Square brackets specify the bits used. For example Chip_ID[3:0] is the first four bits of Chip_ID.

Data word	Length (bits)	Value (binary)
IDLE	8	1111_1111
CHIP HEADER	16	1010<Chip_ID[3:0]><BC_for_frame[10:3]>
CHIP TRAILER	8	1011<Readout_Flags[3:0]>
CHIP EMPTY FRAME	16	1110 <Chip_ID[3:0]><BC_for_frame[10:3]>
REGION HEADER	8	110<Region_ID[4:0]>
DATA SHORT	16	01<Encoder_ID[3:0]><ADDR[9:0]>
DATA LONG	24	00<Encoder_ID[3:0]><ADDR[9:0]>_0_<Hit_Map[6:0]>
BUSY ON	8	1111_0001
BUSY OFF	8	1111_0000

IDLE The IDLE data word is transmitted whenever data is not ready to be transmitted.

CHIP HEADER Each data packet is preceded by a CHIP HEADER word. The bits Chip_ID[3:0] identify the location of the chip on the HIC. The second byte of CHIP HEADER contains BC_for_frame[10:3] which is the value of the internal bunch crossing counter at the moment the trigger was received for the data packet in question.

CHIP TRAILER Each data packet is succeeded by a CHIP TRAILER word. The Readout_Flags[3:0] is comprised of four flags, BUSY_VIOLATION, FLUSHED_INCOMPLETE, STROBE_EXTENDED and BUSY_TRANSACTION. BUSY_VIOLATION is one when the ALPIDE chip transmits an empty data packet due to reaching the limits of its data processing capabilities. FLUSHED_ INCOMPLETE is one when an MEB within the ALPIDE chip was flushed to create memory for further events. This flag is only seen in the continuous mode of operation. STROBE_EXTENDED is one when the framing window of the event was extended after the reception of an external trigger. BUSY_TRANSITION is one when the BUSY_VIOLATION was asserted during the readout of the frame.

CHIP EMPTY FRAME The CHIP EMPTY FRAME data word is present when the readout event is empty. Aside from the first four bits, the CHIP EMPTY FRAME data word is identical to the CHIP TRAILER data word.

REGION HEADER The REGION HEADER data word indicates the start of transmission of data for a given region. Region_ID[4:0] denotes one of 32 regions. The REGION HEADER data word is sent only for regions that have hits. It is sent in ascending order, omitting regions with no hits.

DATA SHORT The DATA SHORT data word provides the location of the hit pixel. Encoder_ID[3:0] is the priority encoder address inside the region, whilst ADDR[9:0] is the pixel address inside the priority encoder.

DATA LONG The DATA LONG data word is transmitted if clustering is enabled. The encoder_ID[3:0] and ADDR[9:0] provide the location of the pixel with the lowest address in a cluster. Hit_Map[6:0] provides the cluster shape.

BUSY ON and BUSY OFF The BUSY ON and BUSY OFF data words are transmitted on the assertion and de-assertion of a BUSY status respectively.

7.3 Decoding of ALPIDE data within the RU

High-speed data from the ALPIDE masters are transmitted via Samtec Twinax cables to the RUs. The Samtec cables consist of 12 differential pairs: two for clock and control, nine for high-speed data transmission and one is unused. For the IB, one Samtec cable connects one IB stave to a single RU, where each of the nine ALPIDE chips has their own high-speed data line. For the MLs, one Samtec cable connects to four master chips on each HS, using four of the high-speed data lines. For the OLs, one Samtec cable connects to seven master chips on each HS, using seven of the high-speed data lines. An illustration of the number of data, clock and control lines for each layer can be seen in Fig. 7.2.

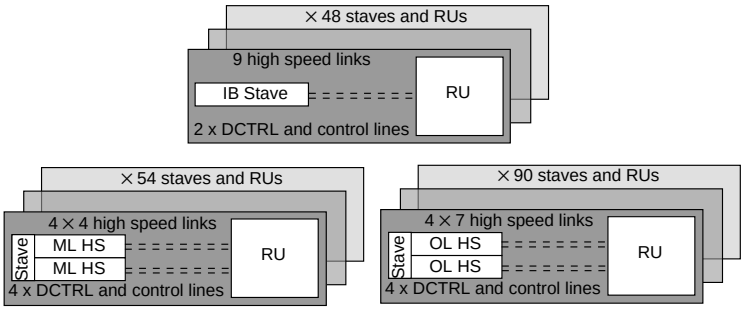

Figure 7.2: Illustration of the number of data, clock and control lines per layer.

Data lanes receive data from the ALPIDE chips, decode that data and transmit them to the GBT packer, which assembles the data into a new format before transmitting it to the CRU. Each trigger sent to the detector generates one data packet to transfer from the data lane to the GBT packer. Errors in the data stream are flagged. The RU will attempt to recover from or ignore errors in the data stream if possible. Errors which cannot be ignored or recovered from the affected data lane result in being unable to receive more data. Provided lost data is flagged, it can be tolerated that a small fraction of data is lost. An overview of the possible errors within the data stream which are provided in real-time is given in Sec. 7.3.1.

7.3.1 Real-time data integrity checks

The RU provides real-time feedback on data integrity. An outline of possible flags corresponding to data stream problems is outlined in this section.

8b10b errors

The data are 8b10b encoded on the serial link. The driver and pre-emphasis driver in the DTU have a configurable output current, customisable via a 4-bit DAC.

For the OB, the DTU provides a line rate of 400 Mb/s over 7.6 m of cable. Due to the time constant, RC, of the cable, a pre-emphasis driver is used to boost high-frequency components which would otherwise be attenuated leading to a slower rise time of the signal. Too much pre-emphasis, however, leads to an overshoot, where the high-frequency components have an amplitude greater than the range of the receiver. Tuning the pre-emphasis is necessary to optimise the rise time whilst forbidding signal loss. The driver amplifies every frequency component of the signal. Tuning the driver is necessary to combat the effects of attenuation along the cable.

8b10b encoding maps 8-bit data words to 10-bit data-words to maintain a DC level close to 0 V. This means there are 2^{10} possible data words, which are mapped to 2^8 data words. Therefore there are $2^{10} - 2^8$ unmapped words. An 8b10b error is raised when an unmapped data word is received. Non-optimal driver and pre-emphasis driver currents lead to errors interpreted by the RU as 8b10b errors. In addition, 8b10b errors can be caused by clock jitter.

Detector timeouts

Each trigger sent to the detector generates one data packet. If a data packet is not received by the GBT packer, a detector timeout flag will be asserted by the affected lane. Detector timeouts can be caused by configuration issues such as a lack of power to the detector, a link not being activated, a problem with the clock propagation or a problem with the transceiver alignment (described in Sec. 8.1.4). They can also be caused by non-optimal driver and pre-emphasis driver currents.

Busies

A BUSY word is transmitted by an ALPIDE master to the RU if it, or one of its slaves, is reaching the limits of its data acquisition capabilities, i.e. its MEBs are becoming full. The BUSY word, and the transceiver which received it, is flagged by the RU. The BUSY state does not prevent data taking, but it indicates there will be some data loss if a further trigger arrives whilst the busy state is ongoing.

The BUSY state often indicates that a particular ALPIDE chip has at least one bad double column. Masking these loud double columns can solve the problem.

Lane FIFO overflows

An excessive input throughput from the ALPIDE chip causes the RU to raise a lane FIFO overflow flag. Table 7.3 shows the maximum throughput capacities for transmission channels in the ITS2. For the OB links to the RU, the maximum throughput of ALPIDE data to the RU is 320 Mb/s.

Lane FIFO overflows are handled by stripping the REGION_HEADER, DATA_SHORT and DATA_LONG data words from the GBT packer, with a code word marking the occurrence of the event in the data stream. When the FIFO is no longer overflowing, these data words are again propagated as normal.

Lane FIFO overflows can be caused by improper masking of bad double columns and pixels or improper threshold tuning.

Table 7.3: Maximum throughput capacities of transmission channels in the ITS2 [97].

Channel	Capacity
IB link	960 Mb/s
OB link	320 Mb/s
RU-CRU GBT link	3200 Mb/s
CRU to FLP RAM	13.750 GB/s
FLP to EPN	8.625 GB/s
Total ITS to EPN	40 GB/s [98]

Illegal triggers

An illegal trigger flag is raised when the RU receives a trigger from the LTU which is out of time with the data packet received by the ALPIDE chip. This means the GBT packer cannot determine which trigger the data packet was received on. The

VTRx transceiver on the RU receives triggers via an optical link from the LTU. The illegal trigger flag is asserted under two conditions. One is when the LTU is sending triggers, as well as the RU. The other is when the VTRx transceiver is contaminated, meaning the optical link is unstable due to a low signal to noise ratio.

7.4 Common Readout Unit

The RUs transmit data to the Common Readout Units (CRUs) via GBT links. For the IB, three GBT links connect to the CRUs, whilst for the OB two GBT links connect to the CRUs. Each CRU has compatibility for 24 input links. A total of 22 CRUs are used for the readout of all 192 RUs in ITS2 over 432 GBT links. The CRUs control the RUs via optical links. Table 7.4 shows the number of RUs and CRUs per layer. The CRUs are installed in the First Level Processor (FLP) units.

Table 7.4: Number of RUs per CRU for each layer

Layer	RUs	GBTx per RU	CRUs	RUs per CRU
0	12	3	2	6
1	16	3	2	8
2	20	3	4	5
3	24	2	2	12
4	30	2	4	8, 7, 8, 7
5	42	2	4	11, 10, 11, 10
6	48	2	4	12

7.5 RU to CRU data format

This section will describe the format of data packaged within the RU and transferred to the CRU during the detector verification and cosmic data taking. It's important to note that the data format is evolving and the format described here is likely to be deprecated after commissioning.

The GBT packer is responsible for organising data received from the ALPIDE chips by the data lanes into a format to be transmitted to the CRUs. The GBT packer takes the data packet from the ALPIDE chips and adds header, trailer and status packets. For the IB, one GBT packer serves three ALPIDE chips. Three GBT packers are needed to serve an IB stave, with each GBT packer transmitting data to

the CRU via three GBT links. For the OB, one GBT packer serves one Half-Stave (HS). Therefore, two GBT packers are needed for each stave, with each GBT packer transmitting data to the CRU via two GBT links. The RU is identical for the IB and the OB, but in the OB, the third GBT packer (and GBT link) is unused.

A round-robin multiplexer collects the data from each lane and arranges it into a GBT word once data from one event are collected. Each GBT word is arranged in the RU firmware and is 80 bits long, consisting of nine 8-bit data words and one 8-bit identifier, plus one single bit, 'DATA_VALID'.

Data are transferred from the RU to the CRU upon receipt of a heartbeat trigger from the CTP. The heartbeat triggers are sent at the same frequency as the orbit frequency, i.e. 10kHz. One heartbeat frame contains at least an SOP, Raw Data Header (RDH) and EOP. It can also contain user data after the RDH and before the EOP. In the case that the user data exceeds the maximum packet size (512 GBT words), the page counter in the RDH is augmented, whilst the stop bit remains equal to 0. Upon receipt of the next heartbeat trigger, the current CRU packet is closed with the EOP control data word. Then, a heartbeat closing packet is sent containing an RDH with the stop bit equal to 1.

The data transmitted between the SOP and EOP is named the CRU data packet. Each CRU data packet is preceded by a 4 GBT word, the RDH, the purpose of which is to identify the CRU data packet. The structure of an RDH can be seen in Fig. 7.3.

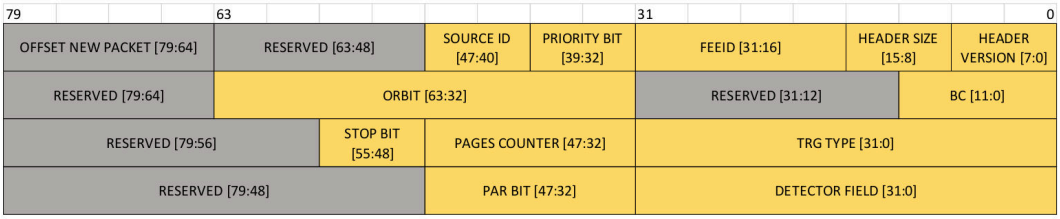

Figure 7.3: The structure of RDH version 6. This is the RDH version used in the verification and cosmic data collection presented in this thesis.

The meaning of the entries in the RDH is as follows:

- Header version (8 bits): the RDH version.
- FEE ID (16 bits): the unique ID of each RU link, i.e. one per OB half stave. Each RU has a 10-bit dipswitch with a unique identifier specifying the layer and stave number. This is mapped to the 16-bit FEE ID to specify the layer, stave and link.

- Priority bit (8 bits): if this is set to 0x1, the data packet is given a higher priority
- Source ID (8 bits): indicates the detector ID, i.e. the detector ID of ITS2 is 32.
- Orbit counter (32 bits): specifies the orbit from which the data is from.
- BC (12 bits): the bunch crossing counter.
- TRG type (32 bits): trigger type set by the CTP on the reception of the HB.
- Detector field (32 bits): detector field.
- PAR (16 bits): Pause And Recover.
- Stop bit (8 bits): equal to 8'b11111111 to identify last data page, 8'b00000000 otherwise.
- Page counter (16 bits): counter to denote different CRU data packets belonging to the same heartbeat trigger.
- Reserved (64 bits total): filled with 0 by the RU. The CRU uses these fields to add extra padding to the data packets. The CRU adds the following data, which are irrelevant to the RU firmware:
  - Link ID (8 bits): the GBT channel from which the data are received. Can have a value from 0 - 11.
  - Memory size (16 bits): size of the RDH and payload in bytes.
  - Offset next packet (16 bits): offset in bytes to the next RDH in memory. Dynamic offset allocation based on the payload size reduces the size of the data packets.
  - Packet counter (8 bits): counter which increases for every packet received on a particular link.
  - CRU ID (12 bits): the unique identifier for the CRU from which the data is received.
  - Datapath wrapper ID (4 bits): used to identify one of two PCI endpoints within the CRU.

The data received by the CRU is then stored in FLP memory.

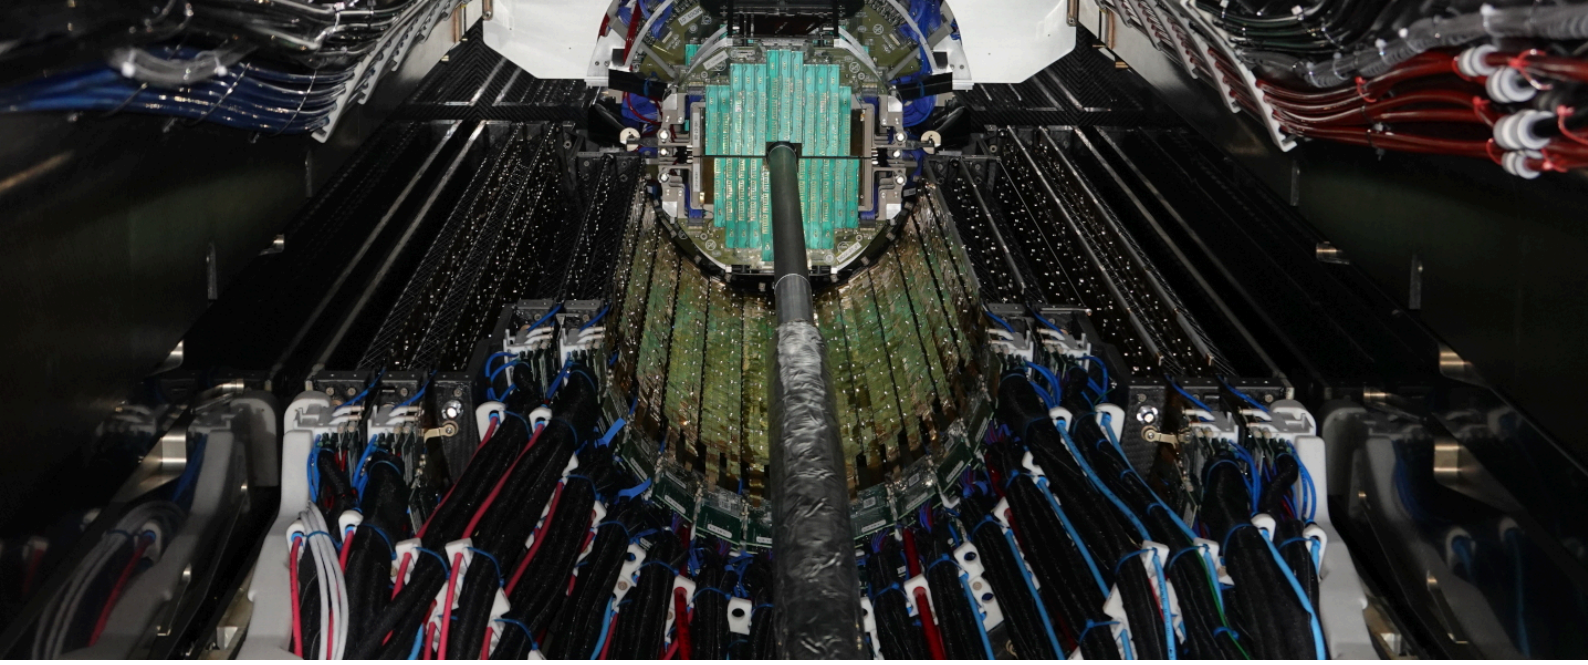

Chapter 8:
Outer Barrel verification

After installation of the staves into the half barrels, each stave was verified using the final services including the readout and power system with a sequence of staves tests, which are outlined in this section. Each stave was verified independently. The first steps in the verification process involved powering the stave, configuring the ALPIDE chips and aligning the RU GPIO inputs on the RU. Two scans are then carried out, one to check the DCTRL communication with the stave and one to verify data transmission from the stave to the RU via the high-speed lines. Next, the chip threshold is tuned and noisy double columns and pixels are determined. The final two scans determine the FHR and chip threshold respectively. An overview of the verification procedure is shown in Fig 8.1.

8.1 Powering and configuration

The optimum operating voltage of the ALPIDE chip is 1.8 V for both digital and analogue supplies. To achieve this, the filterboard and cable resistance was taken into account as well as the tolerances of the power-unit channel potentiometers and the common ground of the power bus. These considerations are explained below.

Header photograph shows the bottom half of the Outer Barrel fully installed around the beam-pipe. © CERN. Taken from [91].

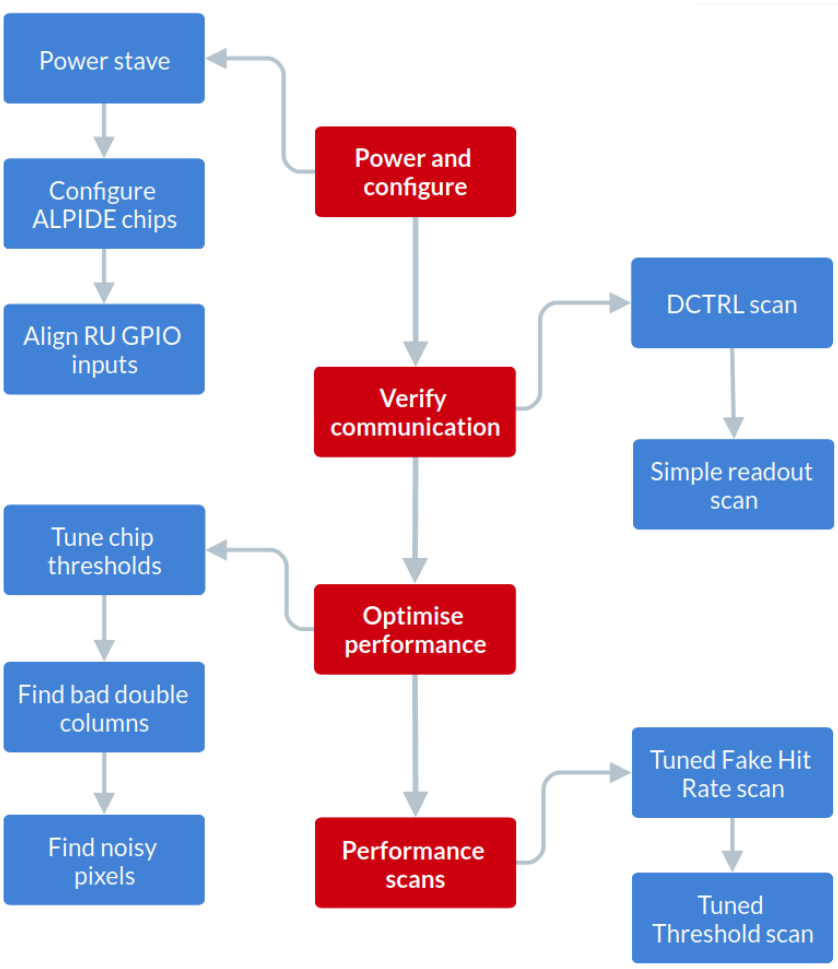

Figure 8.1: Flowchart showing the sequence of the OB verification procedures. The procedures are grouped into four categories, shown in red. Blue boxes denote the procedures of each category.

8.1.1 Potentiometer offset

Each power unit channel has an 8-bit digital potentiometer which sets the output voltage according to the value supplied via I^2C transaction from the RU. Tolerances in the 256-step wiper resistance and end-to-end resistance of each digital potentiometer affect the voltage it actually supplies. Calibration is needed to ensure each supplies the intended voltage. Once calibrated, the selected value can be stored in its non-volatile memory such that it is recalled on power-up of the PU. The potentiometers can be calibrated using the on-board 12-bit, 8-channel ADCs. The calibration procedure is as follows:

- Set digital potentiometers to supply 1.82 V
- Measure the actual supplied voltage with the on-board ADCs
- Calculate the offset

Figure 8.2 shows the distribution of both analogue and digital PU offsets calculated for the entire OB. The conversion from ADC units to mV is as follows:

$$V = \frac{3.072}{4096} ADC. \tag{8.1}$$

The ADC has a resolution of 0.75 mV. The mean PU offset was 30 ADC units and 22 mV, with a standard deviation of 3 ADC units and 2 mV respectively. No significant difference between analogue and digital channels is expected and, indeed, the mean PU offset of the analogue channels and digital channels is equal.

8.1.2 Common ground of the power bus

The ground line on the power bus is shared by all HICs. The voltage supplied to each HIC needs to be increased to account for losses due to the resistance of the power supply cable and the power bus. The method used to achieve this, accounting for the common ground line, will be outlined here.

Firstly, all HICs are powered without voltage drop compensation or chip configuration. The digital and analogue currents, I^D and I^A respectively, are measured by the on-board ADCs. Each HIC is iterated through, starting from the HIC closest to the power input, measuring I^D and I^A. The voltage drop via ground ignoring the common ground, α, for each HIC is calculated as:

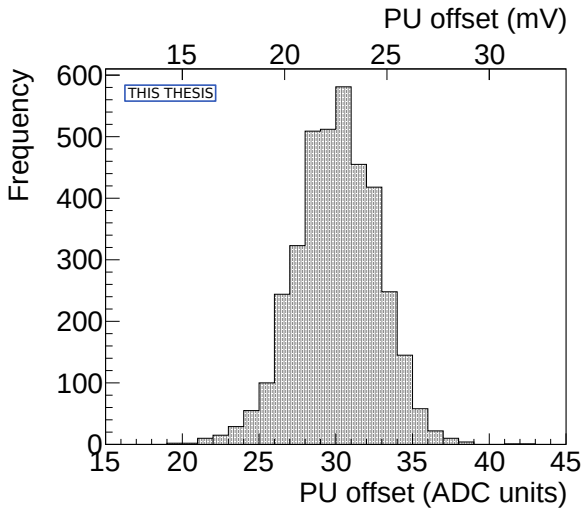

Figure 8.2: Distribution of PU offsets shown in ADC units and mV.

$$\alpha = (I^D + I^A)R, \tag{8.2}$$

where R is the resistance of the ground line on the powerbus for the HIC in question. Since the ground line is shared, the HIC closest to the power input has a higher current passing through the power bus ground than the HIC furthest away from the power bus input. Each HIC is iterated through, starting from the HIC furthest away from the power input. The contribution to the voltage through the shared line from a particular HIC, β with index i, where i starts at the HIC furthest from the power source is given by:

$$\beta_i = I^T(R_i - R_{i-1}), \tag{8.3}$$

where I^T is the sum of I^D and I^A for the HIC of index i and all HICs with a index smaller than i.

The voltage drop for a given HIC, γ, with index j, where j starts at the HIC closest to the power source, is given by:

$$\gamma_j = \alpha_j + \sum_{0}^{j-1} \beta_j. \tag{8.4}$$

The difference between supplied voltage and on-HIC voltage for a HIC with index j, is given by:

$$V_j^{D/A} = I_j^{D/A}(R_{cable} + R_j) + \gamma_j, \tag{8.5}$$

where R_{cable} is the resistance of the power supply cable. To compensate for the voltage drop, the voltage is increased from the supplied value by $V_j^{D/A}$.

8.1.3 ALPIDE configuration

A global broadcast to the ALPIDE chips is asserted stave by stave. The Control Management Unit (CMU) and Data Management Unit (DMU) modules are configured first, setting the previous chip ID and the initial token. The initial token is given to the master chip, whilst the previous chip IDs denote the order in which the slaves are read by the master. Manchester encoding and Double Data Rate (DDR) are configured within the register too. For the verification scans, both Manchester encoding and DDR were enabled.

The high-speed line driver, pre-emphasis driver and PLL charge pump, are configured next. For the verification scans, the standard settings used were 0x3, 0x8 and 0x8 DACs for the driver, pre-emphasis and charge pump respectively. Each of the three DACs is four bits wide, giving a maximum setting of 0xF DACs. For all staves, the standard pre-emphasis and charge pump were optimal. For two chips on two staves, outlined in Table 8.1, a non-standard driver setting was used.

Table 8.1: Abnormal driver settings in the ITS2 OB.

Stave	Chip ID	Driver DAC setting
L5_00	72	0x6
L5_23	96	0x2

34 slave chips are known to be dead from the MOSAIC scans. For these slaves, the previous chip ID is set to be 0 and the DDR on the parallel port is disabled, so they are never read out by the master.

The sensor matrix is configured, according to the scan type and the characteristics of each stave. Bad pixels and bad double columns are masked in the ALPIDE chip, disabling their readout.

The ALPIDE chips draw a higher current after configuration. The voltage drop is recalculated and applied.

8.1.4 Transceiver alignment

As mentioned in section 4.4.3, the ALPIDE chip has a PRBS-7 generator in the DTU. The GPIO transceivers in the RU have an in-built PBRS pattern checker, capable of checking one of five industry standard PRBS patterns: PRBS-7, PRBS-9, PRBS-15, PRBS-23 and PRBS-31.

The PRBS pattern is generated in the ALPIDE master. The GPIO transceiver locks onto the pattern and calculates the expected next data word. The idelay is the analogue delay of the signal set on the FPGA pin and is a fraction of the bit period. The known PRBS period is divided into 512 idelay units. The number of PRBS errors, where the pattern received by the GPIO transceiver does not match the pattern calculated by the FPGA, is recorded as a function of idelay in steps of 10 idelay units. The most suitable idelay is the centre of an open window. Figure 8.3 shows the size of the open window for all instances of transceiver alignment conducted throughout the OB verification. Every stave in the OB has at least one entry.

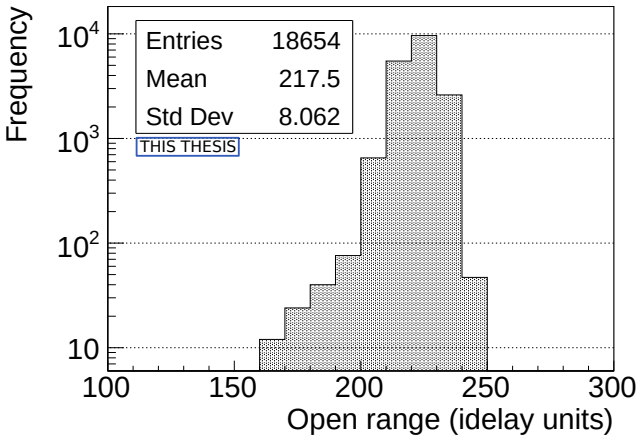

Figure 8.3: Size of open window in idelay units for all staves in the OB across every case where transceivers were aligned.

8.2 Communication tests

Two scans are run to verify the communication between the stave and RU. The control test verifies the functionality of the DCTRL communication, whilst the simple readout test verifies the functionality of the high-speed line communication.

8.2.1 Control test

For each ALPIDE chip, the 16-bit wide DTU test register is written to and read from 1000 times. For each iteration, the pattern written to the register is varied. If the read-back of the register does not match the test pattern written to the register, then the chip fails the control test.

The control test will be passed if the ALPIDE chip is functioning correctly, if the data cables are correctly connected and if the transceiver delay is optimally set. A control test which fails for an entire HS indicates a problem with the powering or data cable connection for that HS.

The control test will only indicate that the read and write is working or is not working. It cannot determine if the read is not working or the write is not working. In order to find this out, the IBIAS ALPIDE register can be written to. Writing 0x0 to the IBIAS register results in a decrease in analogue current consumption of the ALPIDE chips, if the register is correctly written to.

After ensuring power cables were correctly attached, no chip was found to have control issues aside from the 34 known bad chips from the MOSAIC scans.

8.2.2 Simple readout test

The chip thresholds are not tuned and no masking of pixels or double columns is applied. The ALPIDE chips are put into continuous trigger mode with a trigger frequency of 11.2 kHz with no internal stimuli. The scan is a simple readout test, checking that reasonable data is received by the RU from each link.

The RU could potentially flag 8b10b errors, detector timeouts and busies here. 8b10b errors and detector timeouts would motivate a look at the driver and pre-emphasis driver currents, whilst detector timeouts would also motivate an investigation into the transceiver alignment and detector powering.

8.3 Tuning and masking

Three scans are performed to ensure reasonable data is sent from the ALPIDE chips. The threshold tuning scan is used to ensure the response to a given input charge is uniform across the detector. The double column and tuned fake hit scan without masking are both used to determine which areas of the pixel matrix should be masked due to being noisy. The double column scan finds double columns which

have a faulty priority encoder, preventing their successful readout. The tuned fake hit scan without masking finds pixels whose hit rate is above the expected threshold, characterising certain pixels as noisy.

8.3.1 Threshold tuning

A threshold tuning scan is performed to tune the chip thresholds to $100\,e^-$. The threshold scan works as follows. Five rows are scanned (1, 2, 254, 255, 509, 510) with 21 injections of $100\,e^-$ at varying VCASN and ITHR register values. First, VCASN is augmented, from 30 to 70 in steps of 1 DAC, whilst ITHR is held at the reset value. Next ITHR is augmented from 20 to 130 in steps of 1 DAC, whilst VCASN is set to be the most suitable value determined from the VCASN scan.

In each case, the number of hits per scanned pixel is recorded with the corresponding register setting appearing in the data stream as the orbit counter. The register setting at which the number of hits per stave first reaches 50% of the maximum number of pixels per stave is recorded as the optimum register value.

Figure 8.4 shows the optimum VCASN and ITHR values found for each chip in the OB. Figure 8.5 shows the distribution of optimum VCASN and ITHR values. Both follow a Gaussian distribution. For VCASN, the mean value is 56.6 ± 0.02 DACs. For ITHR, the mean value is 45.0 ± 0.01 DACs.

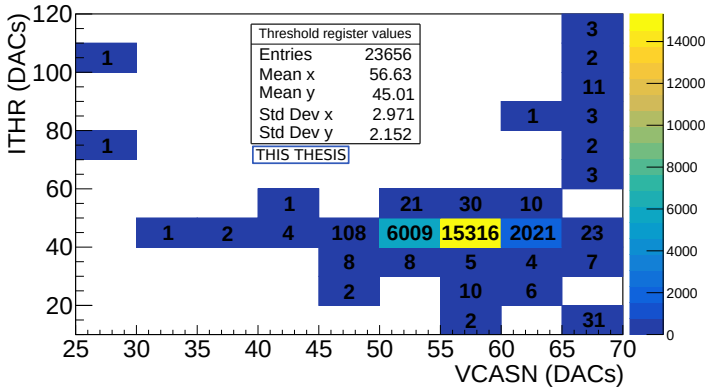

Figure 8.4: VCASN and ITHR values obtained from the threshold tuning scan.

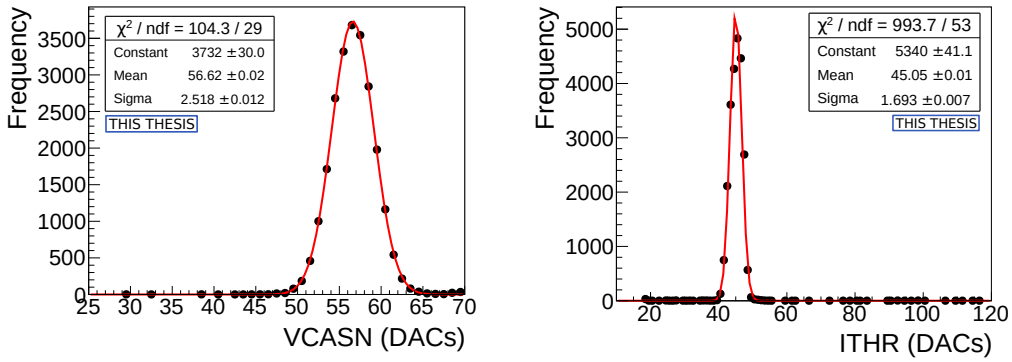

Figure 8.5: VCASN and ITHR values obtained from the threshold tuning scan.

8.3.2 Noise mask generation

Two scans were responsible for identifying noise sources within the pixel matrix. Two sources of noise within the pixel matrix were expected. One is stuck pixels, which read out a hit more often than is reasonable. The other source is broken double columns. Faulty priority encoder logic leads to the inability of some double columns to mask a noisy pixel, causing all pixels within the double column to be unusable.

Double column scan

Each double column is served by a single priority encoder. The logic of some priority encoders is broken and they are unable to mask a stuck pixel. In this case, the priority encoder continuously reads out the stuck pixel hit, ignoring the other hit pixels in its domain. The double column scan finds the double columns where there is an excess of hits caused by the dysfunctional priority encoder.

Every pixel in each ALPIDE chip was pulsed 21 times with a charge of 500 e^- from C_{inj}, pulsing single rows at a time. Pixels that did not respond to the charge injection were noted as being unresponsive. Double columns which had a number of recorded hits higher than a threshold were noted as being stuck. The unresponsive pixels and stuck double columns are written to file to be used in later tests for masking.

Figure 8.6 shows the percentage of double columns masked for all staves. 13 staves had no double columns masked. The highest number of double columns

masked was $\approx 0.25\%$ of all double columns on a stave, achieved by four staves.

Tuned fake hit scan without masking

The chip thresholds are tuned according to the VCASN and ITHR values obtained in the threshold tuning scans. The masks calculated in the double column scan are not applied. The ALPIDE chips are set to run in continuous trigger mode with no internal stimuli for 30 seconds. Pixels that record a hit are either noisy or responding correctly to a particle event, such as a cosmic muon. Pixels that record a hit more than a given cut value are classified as noisy and written to file for future masking. The cut used was 1×10^{-6}/pixel/event, corresponding to a single pixel firing more than three times in a 30 s run.

The distribution of noisy pixels is shown in Fig. 8.6. The proportion of pixels masked ranges from 0% to 0.004%.

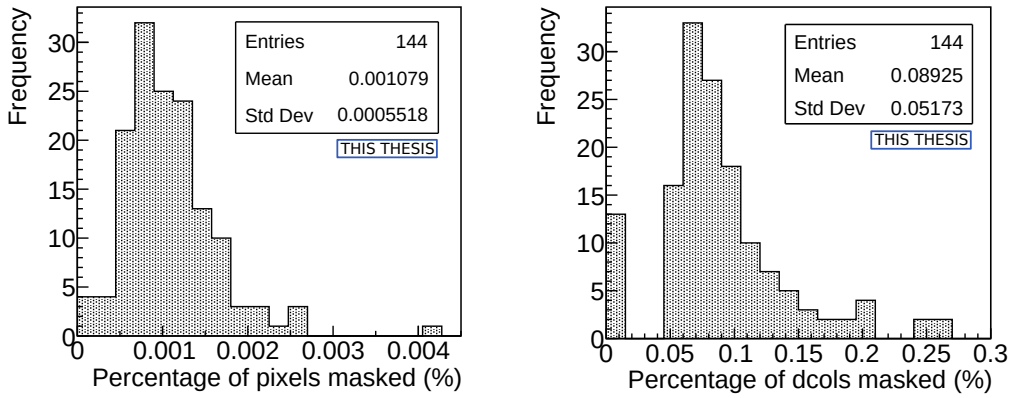

Figure 8.6: Percentage of pixels (left) and double columns (right) masked per stave. 144 is the total number of OB staves.

8.4 Performance scans

Two scans were used to gather two key performance statistics: the Fake-Hit Rate (FHR) and the threshold uniformity.

8.4.1 Tuned fake hit scan with masking

With the chip thresholds tuned and the noisy pixels and double columns masked, a FHR scan is launched to determine the FHR of the stave. The chips are run in continuous triggering mode with a trigger rate of 11.2 kHz and no internal stimuli. The FHR rate is the number of hits recorded per event per pixel.

Figure 8.7 shows the FHR as a function of the number of masked pixels for staves in the OB, for both the unmasked and masked tuned FHR scans. For the unmasked data, the pixels were masked offline. For the masked data, the pixels were masked offline where the number of masked pixels is greater than zero. For zero masked pixels, the pixels were hardware masked according to the masks generated by the tuned FHR scan without masking, as well as the double column scan. In the case of the masked FHR data, the FHR is stable at 10^{-10}/pixel/event for all staves after a further eight pixels have been masked offline. Approximately 60 pixels per HIC are known to be unmaskable from the MOSAIC scans due to faulty masking logic.

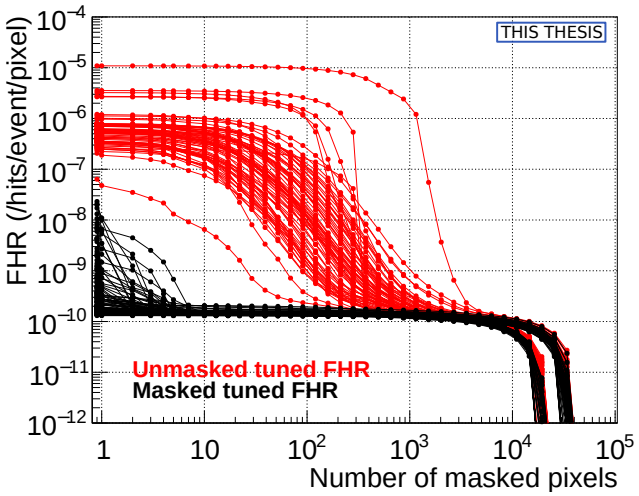

Figure 8.7: Fake hit rate as a function of the number of masked pixels for staves in the OB, for both masked and unmasked tuned FHR scans.

Figure 8.8 shows the FHR of all staves. Pixels are masked in order of increasing noise, i.e. the loudest pixel is masked first. Three cases are presented: no masking, hardware (HW) masking and HW masking plus 10 pixels masked offline. The no masking case refers to the tuned FHR scan without HW masking. The HW masking case refers to the tuned FHR scan with masking, where the masks are calculated from the double column scan and the tuned FHR with-

out masking. The HW masking plus 10 pixels masked offline case, uses the data from the tuned FHR with HW masking and masks the ten loudest pixels offline. The FHR improves from $(9.6 \pm 16) \times 10^{-7}$/pixel/event in the no masking case, to $(5.0 \pm 11) \times 10^{-9}$/pixel/event in the HW masking case. Masking a further 10 pixels offline improves the FHR to $(1.4 \pm 0.3) \times 10^{-10}$/pixel/event.

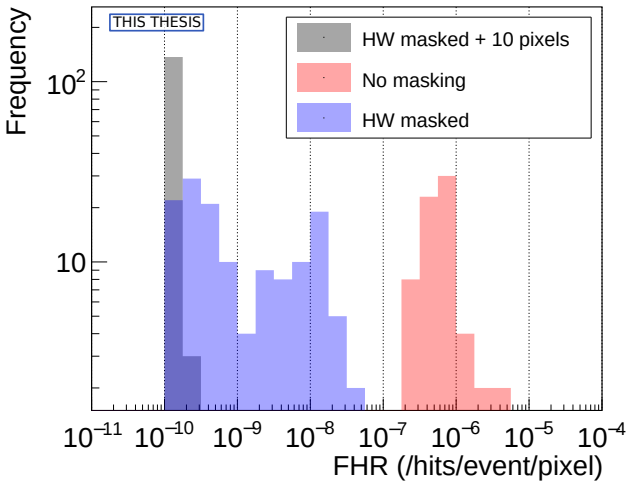

Figure 8.8: Fake hit rate per stave calculated during the verification for various masking scenarios.

8.4.2 Tuned threshold scan

Similarly to the tuned threshold scan performed with the MOSAIC in Sec. 6.2.3, the tuned threshold scan determines the operating threshold of each chip after tuning. C_{inj} has a nominal capacitance of 230 aF. The injected charge, Q_{inj} is given by:

$$Q_{inj} = C_{inj}(VPULSEH - VPULSEL) \tag{8.6}$$

VPULSEH and VPULSEL are set by 8-bit DACs. The maximum value of (VPULSEH - VPULSEL) is 1.8 V. Therefore, the minimum voltage step is 7 mV, equivalent to one DAC unit. The minimum charge step is $10\,e^-$. For the threshold scan, (VPULSEH - VPULSEL) was varied from 0 to 30 DACs. Thus, the injected charge was varied from $0\,e^-$ to $3000\,e^-$ in steps of $10\,e^-$. For each value of Q_{inj}, 21 injections were made.

The value of the charge once the firing probability of all pixels within a single chip reaches 50% is the threshold of the chip. For each Q_{inj}, N_{inj} injections are

made, where N_{hit} pixels detect a hit. The probability of a pixel firing, p_{hit} is given by an error function:

$$p_{hit}(Q_{inj}) = \frac{1}{2}\left[1 + erf\left(\frac{Q_{inj} - \mu}{\sigma\sqrt{2}}\right)\right], \tag{8.7}$$

where μ is the threshold of the chip and σ is indicative of the threshold uniformity across the stave. The shape of the error function gives the measurement its name, an S-curve measurement. If the same measurement is done on a pixel scale, σ is the temporal noise of the pixel. The differential of the fitted error function produces a Gaussian centred on the threshold value with a standard deviation indicative of the threshold uniformity across the stave. Figure 8.9 shows N_{hits}/N_{inj} as a function of Q_{inj} for each stave in the OB, as well as the fitted error function and resulting Gaussian.

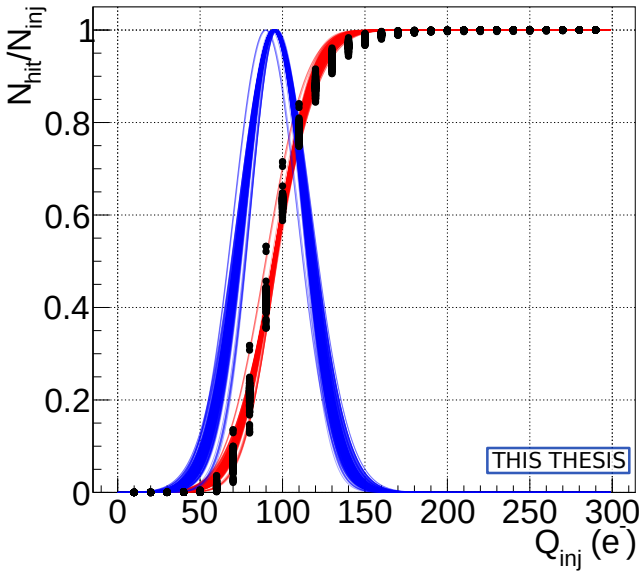

Figure 8.9: S-curve measurement performed for each stave in the OB at least once. The black points represent the raw data, whilst the red lines represent the fitted error function. The blue lines are the derivative of the error functions.

Figure 8.10 shows the threshold and standard deviation distributions per stave obtained from each fitted Gaussian. The mean threshold per stave in the OB, $\bar{\mu} = (94 \pm 5)\,\mathrm{e}^-$. The mean standard deviation per stave in the OB, $\bar{\sigma} = (19 \pm 2)\,\mathrm{e}^-$. The threshold tuning scan aimed to tune the chip thresholds to $100\,\mathrm{e}^-$, not $94\,\mathrm{e}^-$. The reason for this discrepancy is the difference between methods of determining the threshold used by the threshold tuning scan and the tuned threshold scan. In the

case of the threshold tuning scan, Q_{inj} is set to $100\,e^-$ and the VCASN or ITHR registers are adjusted. The value of the VCASN or ITHR register at the point where 50% of total pixels fire is recorded as being the optimum register setting. Unfortunately, this method fails to accurately measure the chip threshold, unlike the S-curve measurement. The method used in the threshold tuning scan was discovered to be incorrect after all data had been taken and unfortunately could not be corrected within the timeframe of this thesis. However, a threshold of $(94 \pm 5)\,e^-$ is sufficiently high for a low fake-hit rate and low enough for good detection efficiency [99].

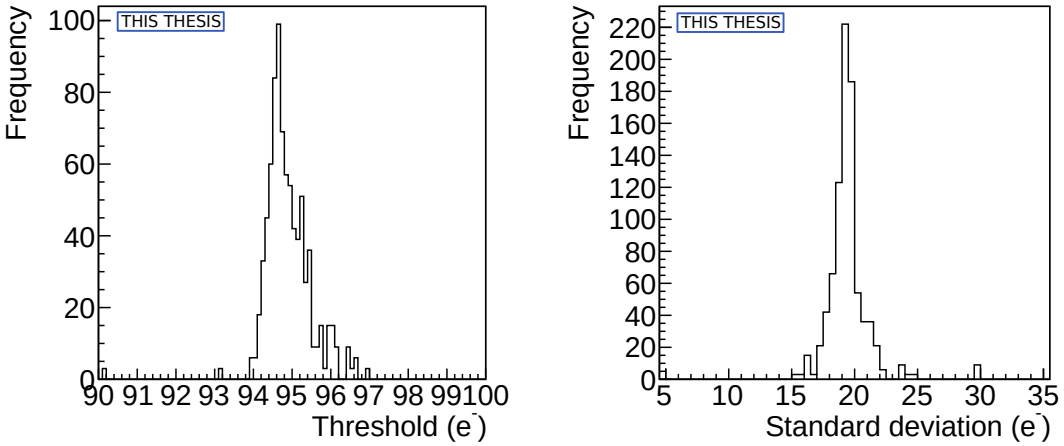

Figure 8.10: (Left) Threshold distribution per stave in the OB with $\bar{\mu} = (94 \pm 5)$ e^-. (Right) Standard deviation per stave in the OB with $\bar{\sigma} = (19 \pm 2)$ e^-. For each stave, at least one measurement was made.

Figure 8.11 shows the threshold and standard deviation distributions per chip in the OB. The mean threshold per chip in the OB was $\bar{\mu} = (94.00 \pm 0.02)\,e^-$, with a standard deviation of $8.36\,e^-$. The mean standard deviation per chip in the OB was $\bar{\sigma} = (19.236 \pm 0.008)\,e^-$, with a standard deviation of $2.74\,e^-$.

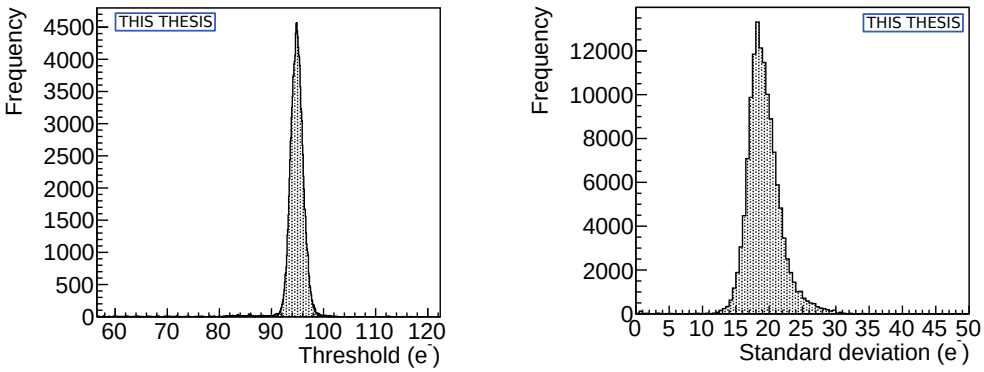

Figure 8.11: (Left) Threshold distribution per chip in the OB with $\bar{\mu} = (94 \pm 8)$ e$^-$. (Right) Standard deviation per chip in the OB with $\bar{\sigma} = (19 \pm 3)$ e$^-$. For each chip, at least one measurement was made.

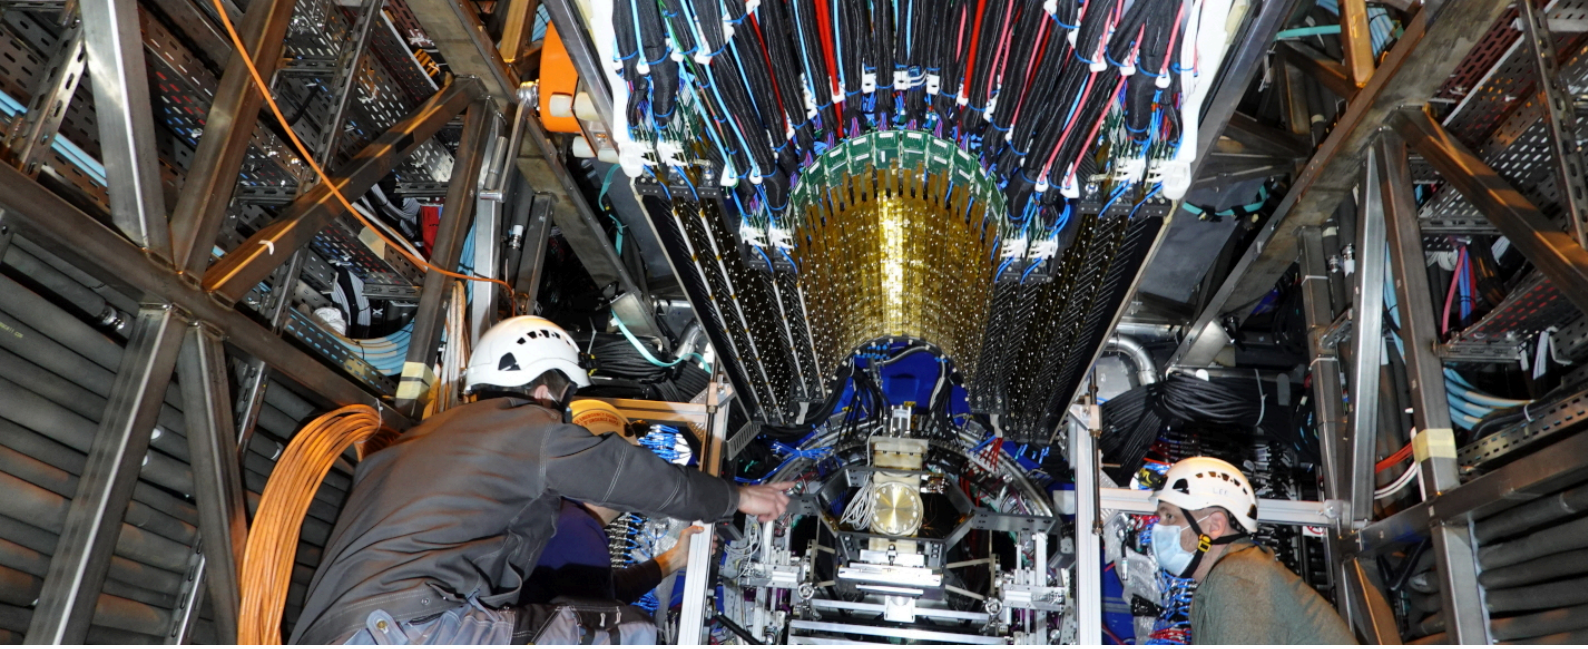

Chapter 9:

Measurement of cosmic muons

A cosmic muon data collection campaign was carried out at the end of 2020. The following chapter describes the cosmic muon acquisition scan, the fitting of muon tracks and the use of these muon tracks to measure the detection efficiency of the ITS2 Outer Barrel (OB).

9.1 Cosmic acquisition scan

The Cosmic acquisition (CAQ) scan is very similar to the masked FHR scan used during the stave verification. The ALPIDE chips are put into continuous mode with a trigger frequency of 11.2 kHz and the chip thresholds are tuned to 94 e^- according to the optimum VCASN and ITHR values found during the verification. The main difference between the CAQ scan and the tuned FHR scan during the stave verification is the use of triggers synchronised across staves provided by the LTU.

Synchronised triggers used in the CAQ scan were implemented with a Run Control (RC) system. A schematic of the RC can be seen in Fig. 9.1. One RC server is started for each CRU on the FLPs in the OB. Each RC server communicates with a RC client running on another machine. The RC client sends a command to each

Header photograph shows the top half of the Outer Barrel ready for installation after rotation. © CERN. Taken from [91].

RC server to prepare the chips for the upcoming run. Each RC server then communicates with its respective CRU which communicates with each RU via SWT. Each RU then configures the ALPIDE chips via DCTRL. Once all chips are configured, the ALPIDE and RU configuration is dumped to a log file and the RC server sends a 'READY' signal to the RC client, which waits for all RC servers to become 'READY'. Each CRU and RU receives its clock from the LTU.

Once all RC servers are 'READY', the RC client sends a signal to start the run. One CRU controlled by an RC server is chosen to be the 'LTU master', which means only this CRU communicates with the LTU server. The LTU master sends a signal to the LTU server, which sends a signal to the LTU to start sending synchronised triggers via GBTx2 to each RU. The CAQ scan is then started.

The RU counters for every RU are printed to the terminal and logged, as well as the power consumption of each stave. The run continues for a set amount of time (chosen to be 500 seconds). Once the run is finished, the ALPIDE and RU configuration is again dumped to a log file and each RC server sends a 'DONE' signal to the RC client.

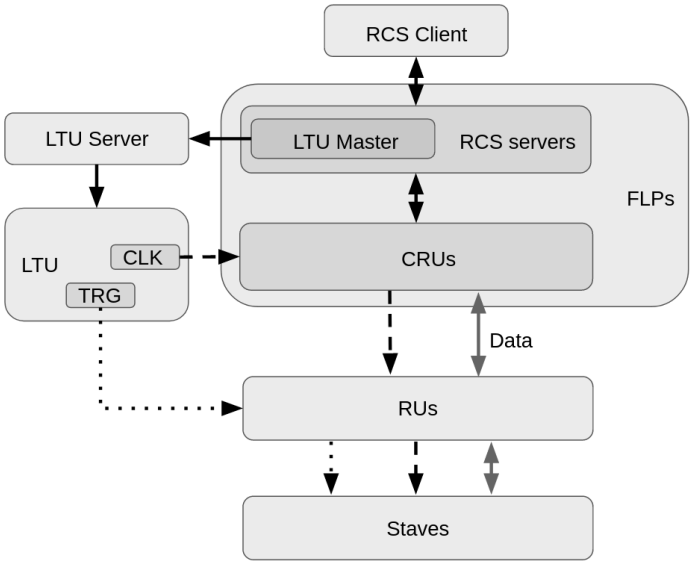

Figure 9.1: Schematic of the Run Control system. The LTU provides synchronised triggers to each RU.

9.2 Cosmic muon data collection campaign

In December 2020, approximately 5×10^6 cosmic muon events were collected with the Outer Barrel of the ALICE ITS2. The raw cluster distribution for each layer can be found in Appendix B. The Outer Barrel was situated in the on-ground commissioning laboratory at CERN with each half-barrel separated. The top and bottom halves of the Outer Barrel will be referred to as OBTOP and OBBOT respectively. Each half-barrel was orientated in the same way spatially, such that both were in the final orientation of OBBOT. The data was acquired independently for each half-barrel over a number of 500 s runs with a trigger frequency of 11.2 kHz and a pixel threshold tuned to $94\,e^-$. The duty cycle was 100% because of the long strobe length and overlap of the analogue response of the ALPIDE chips. Synchronised triggers were achieved with the RC protocol outlined in Sec. 9.1. The RU firmware used was v1.5.0.

98% of pixels were included in the runs, where pixels not included either belonged to dead chips or were hardware masked beforehand due to being noisy. Two staves missed roughly 30% of runs. L6_26 missed the first 33% of runs due to DCTRL communication issues caused by a non-optimal phase offset, which was subsequently noticed and optimised, such that the stave was included in the rest of the runs. L6_41 started to show DCTRL issues in the final 28% of runs, also caused by a non-optimal phase offset. In both cases, the non-optimal phase offsets were caused by an RU firmware bug later rectified. L3_03 missed approximately 38% of runs due to a faulty power-board.

A number of staves were affected by timeouts, as shown in Tab. 9.1. Only the cluster distribution for L6_11 was affected however, as discussed in Sec. 9.4.1.

Table 9.1: Proportion of runs affected by timeouts for affected staves.

Stave	Runs affected
L3_16	$\approx 7\%$
L4_23	$\approx 78\%$
L5_05	$\approx 56\%$
L5_14	100%
L4_36	$\approx 9\%$
L6_11	$\approx 18\%$

In total, 151692 pixels were masked (0.001% of total OB pixels) and 11337

double columns were masked (0.1%). These masks were obtained during the stave verification described in Sec. 8.4.2.

Over the course of the cosmic muon data acquisition runs, the readout units reported no increase in the number of data corruption issues, aside from a known issue with the strength of the optical trigger signal supplied to the VTRx transceiver, which has since been understood and fixed.

Figure 9.2 shows the number of RU counter issues over time. Two RUs were affected by illegal triggers, caused by VTRx issues. L5_10 showed illegal triggers suddenly in all of the last 10 runs. L4_14 showed illegal triggers intermittently from halfway through the scans, before reliably appearing in the last 10 runs.

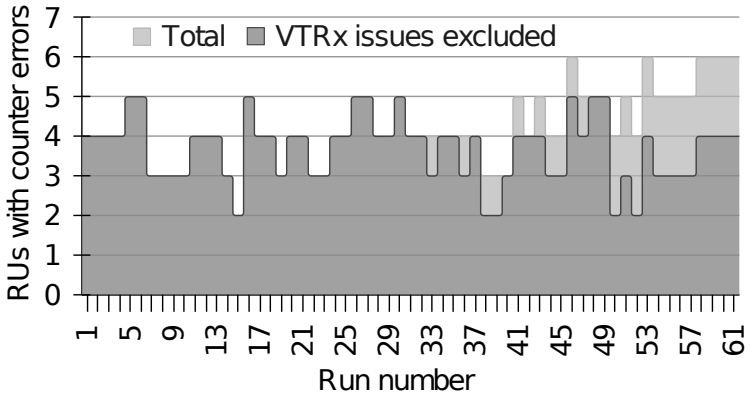

Figure 9.2: Number of RUs with counter errors as a function of run number.

Figure 9.3 shows the breakdown in types of RU counter issues, excluding illegal triggers caused by the VTRx. 12 staves were responsible for all RU counter issues seen in the cosmic campaign. The busies and timeouts are due to non-optimal pixel masks, whist the 8b10b errors are due to problematic high-speed link drivers and transmission lines which were addressed after the cosmics data-taking campaign. The non-optimal pixel masks can be accounted for during the offline noise suppression.

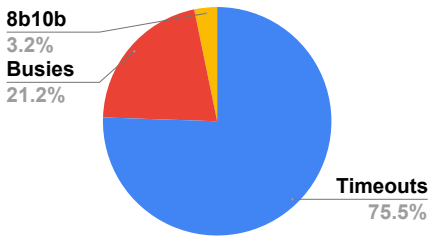

Figure 9.3: Breakdown of RU counter issues by type.

9.3 Cosmic track reconstruction

The cosmic muon data was first noise suppressed offline, before muon tracks were fit. The procedure for noise suppression, track fitting, parameter range and track fitting minimisation variable are described in this section. Multi-track events were ignored.

9.3.1 Noise reduction

Cosmic flux at sea level is about 1 muon per square cm per minute. Each cosmic muon data acquisition run lasted 500 s with a trigger frequency of 11.2 kHz and a 100% duty cycle. Therefore roughly 7×10^{-10} muons are expected per pixel per event, or approximately 4×10^{-3} muons per pixel per run.

Two noise suppression passes of the data were made. The first pass, used to reduce the raw data volume, excluded any cluster which had a hit rate of more than 3×10^{-6} /event/pixel on a run by run basis, corresponding to a single cluster firing no more than 16 times in a single 500 s run. This noise cut reduced the raw data volume by 95%. Figure 9.4 shows the FHR per stave after the first noise suppression pass. The mean FHR per HS was $(1.094000 \pm 0.000002) \times 10^{-9}$ /event/pixel, whilst the median FHR was 7.5×10^{-11} /event/pixel.

The first noise pass, whilst successful at reducing the raw data volume, has a cut too lenient for a full suppression of noise. The second noise pass consisted of cutting any cluster which fired more than twice in a single run. A single pixel is unlikely to be hit more than once by a cosmic muon during a single run. However, since the CAQ scan was undertaken using the continuous operating mode of the ALPIDE chip, an event on the boundary of acquisition windows can be recorded twice, as mentioned in Sec. 4.4.3.

Each run was looped through and the coordinates of each cluster are determined. If a cluster appears more than twice, the cluster coordinates are written to a file. For the cosmic track selection, discussed in Sec. 9.3.3, these noisy clusters are ignored.

9.3.2 Displacement from cluster to line fit

All cluster combinations are fitted with a parametric line in 3D using $\langle d^2 \rangle$ minimisation with MIGRAD [89] using MINOS [89] to generate errors on the free parameters. $\langle d^2 \rangle$ is the average squared distance from the cluster to the fitted track.

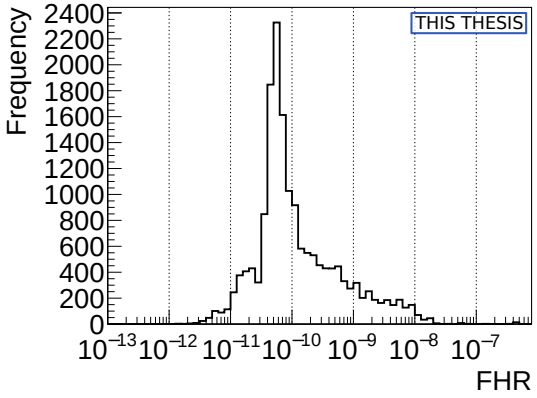

Figure 9.4: Fake Hit Rate per Half Stave after the first noise suppression pass, where pixels whose hit rate exceeded 3×10^{-6}/pixel/event were cut from the data set. Every run in the conducted in the cosmic muon data collection campaign is included.

A line in three dimensions can be specified by two points, $\vec{x_0} = (x_0, y_0, z_0)$ and $\vec{x_1} = (x_1, y_1, z_1)$. A vector, $\vec{v}$, along this line is given by [100]:

$$\vec{v} = \begin{pmatrix} x_0 + (x_1 - x_0)t \\ y_0 + (y_1 - y_0)t \\ z_0 + (z_1 - z_0)t \end{pmatrix}. \tag{9.1}$$

The squared distance between a point on the line described by the parameter t and another point, not necessarily on the line, $\vec{x_p} = (x_p, y_p, z_p)$ is therefore:

$$\begin{aligned} d^2 = &[(x_0 - x_p) + (x_1 - x_0)t]^2 \\ &+ [(y_0 - y_p) + (y_1 - y_0)t]^2 \\ &+ [(z_0 - z_p) + (z_1 - z_0)t]^2. \end{aligned} \tag{9.2}$$

The most appropriate value of t is where d^2 is minimum. Therefore, setting $\frac{d(d^2)}{dt} = 0$ and solving for t, results in the following appropriate value for t:

$$t = -\frac{(\vec{x_0} - \vec{x_p}) \cdot (\vec{x_1} - \vec{x_0})}{|\vec{x_1} - \vec{x_0}|^2}. \tag{9.3}$$

Substituting Eq. 9.3 into the Eq. 9.2 results in:

$$d^2 = \frac{|(\vec{x_1} - \vec{x_0}) \times (\vec{x_0} - \vec{x_p})|^2}{|\vec{x_1} - \vec{x_0}|^2}. \tag{9.4}$$

The full derivation of d^2 can be found in Appendix C. $\langle d^2 \rangle$ was minimised using MIGRAD for each track where $\langle d^2 \rangle$ is defined as:

$$\langle d^2 \rangle = \frac{\sum d_i^2}{N}, \tag{9.5}$$

where N is the number of clusters belonging to a track and d_i^2 is the distance from each cluster to the track squared, where the index i refers to the ITS2 layer. d_i is illustrated in the xy plane in Fig. 9.5.

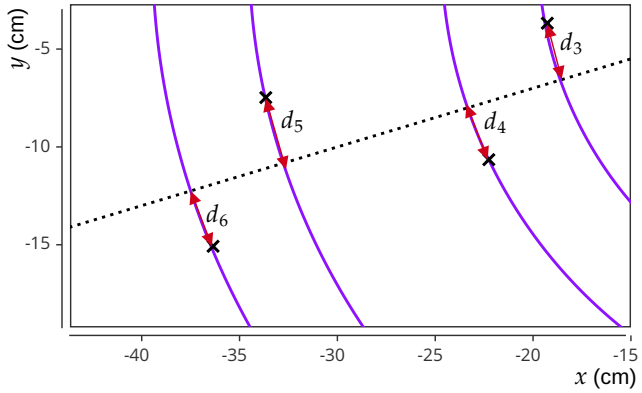

Figure 9.5: Illustration of the distance from each cluster, d_i, belonging to a track shown in x and y. The subscript denotes the layer which the cluster belongs to. The dotted line represents a muon track and the purple concentric lines illustrate a fraction of the ITS2 layers.

9.3.3 Fitting cosmic tracks

An example muon track is shown in red in Fig. 9.6. The vectors $\vec{r}_0$ and $\vec{r}_t$ are the position vectors to points A and B respectively, where point A is the point on the muon track where $y = 0$ and point B is the point on the muon track where $y = t$. A vector, $\vec{a}$, describes the muon track from points A to B. The vector $\vec{v}$ is a unit vector parallel to $\vec{a}$. $\vec{r}_t$ is therefore:

$$\vec{r}_t = \vec{r}_0 + t\vec{v}, \tag{9.6}$$

where t is some scalar. Defining point A as (x_0, y_0, z_0), the unit vector $\vec{v}$ as (v_x, v_y, v_z), and any point on the muon track as (x, y, z), results in the following parametric form of the line describing the muon track:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} x_0 \\ y_0 \\ z_0 \end{pmatrix} + t \begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix}. \tag{9.7}$$

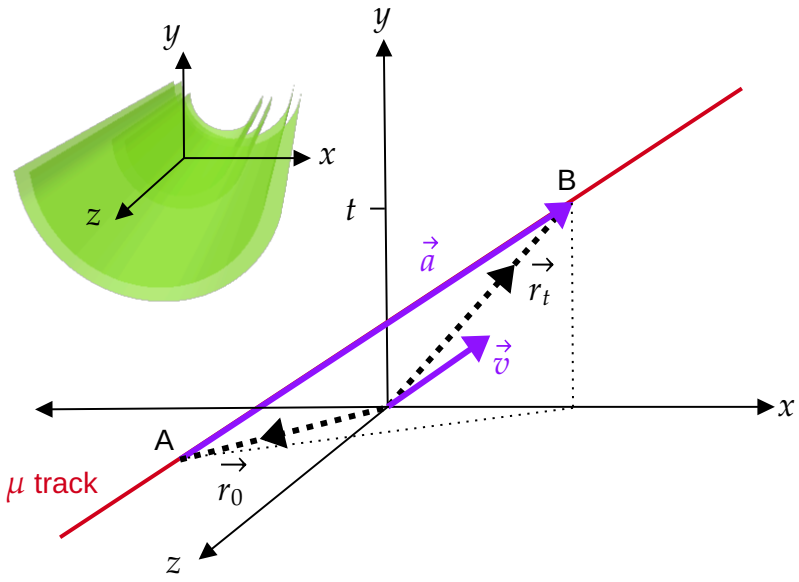

Figure 9.6: An example muon track (red) and its parametric line fit, $\vec{a}$. $\vec{r}_0$ and $\vec{r}_t$ define the position vectors to points A and B respectively, which are two points on $\vec{a}$. The direction vector of $\vec{a}$ is $\vec{v}$. The coordinate system relative to the ITS2 is shown on the left. See text.

Six parameters, x_0, y_0, z_0, v_x, v_y and v_z, define the parametric line. A 3D line can be defined with four parameters if the line is not parallel to a single plane. An over parametrised line will have no unique solution and the minimisation process will not converge. Since cosmic muon flux is greatest from the zenith, parallel to the y axis in Fig. 9.6, a muon track is unlikely to be parallel to the xz plane. If it is assumed that the line is not parallel to the xz plane, y_0 can be forced to 0, where x_0 and z_0 are the coordinates on the xy plane where the line penetrates the xz plane. Similarly, since we require the line to penetrate the xz plane, v_y must be non-zero. Thus, we can scale the direction vector by $1/v_y$. With this manipulation, the parametric equation for the 3D line becomes:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} x_0 \\ 0 \\ z_0 \end{pmatrix} + t \begin{pmatrix} v_x/v_y \\ 1 \\ v_z/v_y \end{pmatrix}. \tag{9.8}$$

The line equation can then be parametrised as follows:

$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} p_0 \\ 0 \\ p_2 \end{pmatrix} + t \begin{pmatrix} p_1 \\ 1 \\ p_3 \end{pmatrix}, \tag{9.9}$$

where $p_0 = x_0$, $p_2 = z_0$, and the parameters p_1 and p_3 are the x and z components of the direction vector respectively.

9.3.4 Parameter range

MINUIT is able to produce more reliable error estimates without specifying the range on parameters [89]. If the minimisation function (in this case d^2) has a clear minimum, MINUIT will be able to find it without parameter limits. However, it is necessary to understand the range in parameter values which are expected, so that the parameters determined by MINUIT can be verified.

The parameters p_0 and p_2 are the x and z coordinates of the point on the fitted line where $y = 0$. The spatial limits of the detector provide a physical range for these parameters, x_{min} and z_{min}. Parameters p_1 and p_3 describe the x and z components of the direction vector, $\vec{v}$, respectively. The components of the direction vector $\vec{v}$ are shown in spherical coordinates in Fig. 9.7.

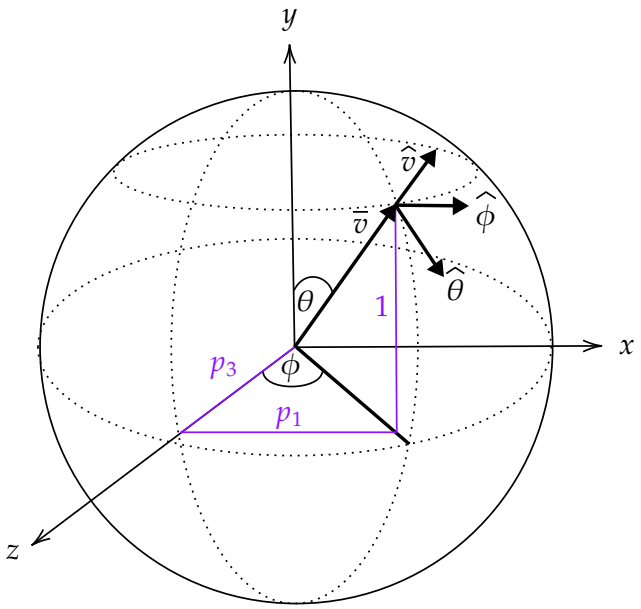

Figure 9.7: The direction vector, $\vec{v}$, displayed in spherical coordinates.

The direction vector, $\vec{v}$, can be written in spherical coordinates as follows:

$$\begin{pmatrix} p_1 \\ 1 \\ p_3 \end{pmatrix} = \begin{pmatrix} |\vec{v}|sin(\phi)sin(\theta) \\ |\vec{v}|cos(\theta) \\ |\vec{v}|cos(\phi)sin(\theta) \end{pmatrix}, \tag{9.10}$$

where $|\vec{v}| \in [1,\infty)$, $\phi \in [-\pi,\pi]$ and $\theta \in [0,\pi/2]$. In addition $sin(\theta)sin(\phi) \in [-1,1]$, $cos(\theta)sin(\phi) \in [-1,1]$ and $cos(\theta) \in [0,1]$. Most cosmic muon tracks will be best described by values of p_1 and p_3 close to 0, since the cosmic muon flux is greatest from the zenith (when $\theta = 0$). The polar angle, θ, and the azimuthal angle, ϕ are related to the components of the direction vector as follows:

$$\theta = cos^{-1}\left(\frac{1}{|\vec{v}|}\right), \tag{9.11}$$

$$\phi = tan^{-1}\left(\frac{p_1}{p_3}\right). \tag{9.12}$$

The effect of the parameters p_1 and p_3 on θ are shown in Fig. 9.8. Figure 9.8a shows the effect on θ when p_1 and p_3 are varied within a range of $[-1,1]$. As expected, the minimum value of θ is indeed 0. θ becomes asymptotic with 90° as $|p_1|$ and $|p_3|$ increase, as shown in 9.8b.

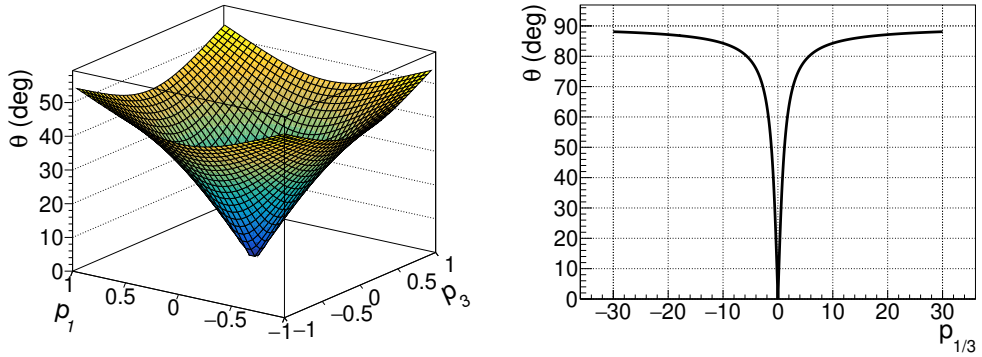

(a) Values of θ in the range of $p_1 \in [-1,1]$ and $p_3 \in [-1,1]$.

(b) Values of θ in the range of $p_1 \in [-30,30]$ whilst $p_3 = 0$ and $p_3 \in [-30,30]$ whilst $p_1 = 0$.

Figure 9.8: Values of θ with varying values of the p_1 and p_3 parameters.

Errors on the free parameters are gathered using MINOS. MINOS defines the error on a parameter as the minimum change in that parameter required to change the minimisation variable (in this case $\langle d^2 \rangle$) by some variable amount. This variable was set to be 2×10^{-4} cm^2, which is the mean absolute displacement from the chip marker position, measured by the CMM, squared, as shown in Fig. 6.3.

9.4 Data analysis of cosmic muon tracks

This section describes features of the cosmic muon data set after track fitting.

9.4.1 Cluster distribution

Figure 9.9 shows the fitted cluster distribution of all the cosmic muon tracks detected in each half barrel. θ is the polar angle to each cluster in the track fit, where OBBOT is in the range $-\pi < \theta < 0$ rad and OBTOP is in the range $0 < \theta < \pi$rad. Masked chips are visible as small rectangles with no clusters. Vertical stripes with a larger number of clusters denote stave overlaps. It can be seen that OBBOT detects more clusters than OBTOP, due to the larger exposure time OBBOT received.

Three staves can be noticed which detected fewer hits than their neighbours, one at ≈ 0.8 rad on L3 (L3_03), one at ≈ -1 rad on L6 (L6_41) and one at ≈ 1.3 rad on L6 (L6_11). Of these three staves, L3_03 missed $\approx 38\%$ of runs due to a faulty power-board, L6 _41 missed $\approx 30\%$ of runs to a non-optimal phase offset and L6 _11 missed $\approx 20\%$ of runs due to timeouts and busies, as described in Sec. 7.3.1, caused by non-optimal pixel masks and driver settings respectively. In addition, one half-stave can be seen which detected fewer clusters than its neighbours, at ≈ -1.4 rad for L4 (L4_23L). L4_23 was affected by timeouts for $\approx 77\%$ of the runs.

L6_26 and L6_41 showed DCTRL issues for $\approx 30\%$ of runs. The decrease in run time for L6 _41 can be seen in the cluster distribution at ≈ -1 rad. Since L6_26 covers the range $-2.9 < \theta < -2.8$ rad, the runs missed due to DCTRL issues are not noticeably visible in the cluster distribution due to the low cosmic flux in this region of θ.

9.4.2 Angle of incidence

Figure 9.10 shows the polar-azimuthal distribution of all reconstructed cosmic muons tracks. Tracks are included if they hit all four half-layers of the detector in their respective half-barrel. The polar angle, θ, and azimuthal angle, ϕ, are described by Fig. 9.7, where z is parallel to the length of each stave.

Three regions of lower acceptance in the polar angle are visible around the ϕ values of -180°, 0° and 180°. The polar angle in these regions of ϕ is limited by the geometry of layer 3 with respect to layer 6. The maximum polar angle in these

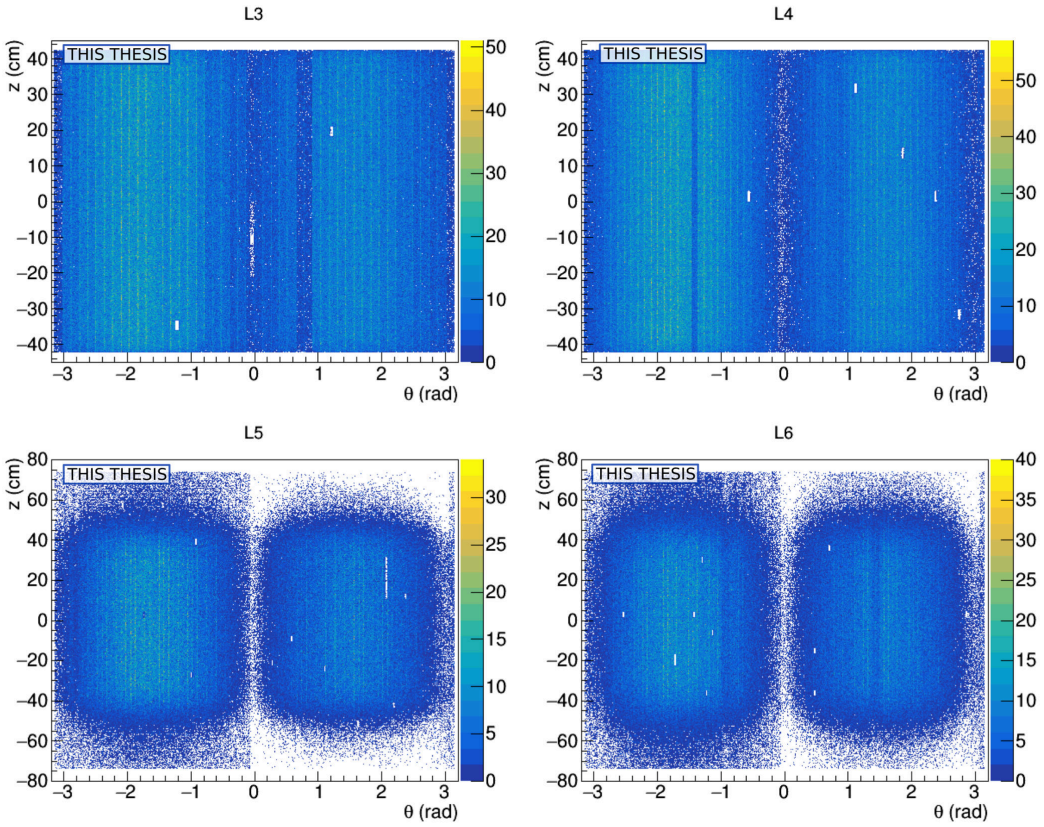

Figure 9.9: Fitted cluster distribution. Layer 3 is shown in the top left, layer 4 in the top right, layer 5 in the bottom left and layer 6 in the bottom right. In each case, both half-barrels are shown. White areas indicate regions where no hits are detected. See text.

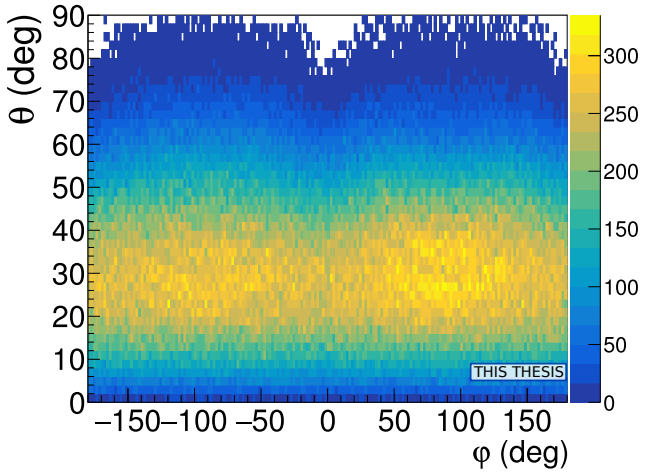

Figure 9.10: Angle of incidence of tracks for the OB. The polar and azimuthal angle for each cosmic muon track which hits all four half-layers of the detector in its respective half-barrel are shown. See text.

regions, illustrated in Fig. 9.11, is approximately 80°, which is as expected in the polar-azimuthal distribution of reconstructed tracks shown in Fig. 9.10.

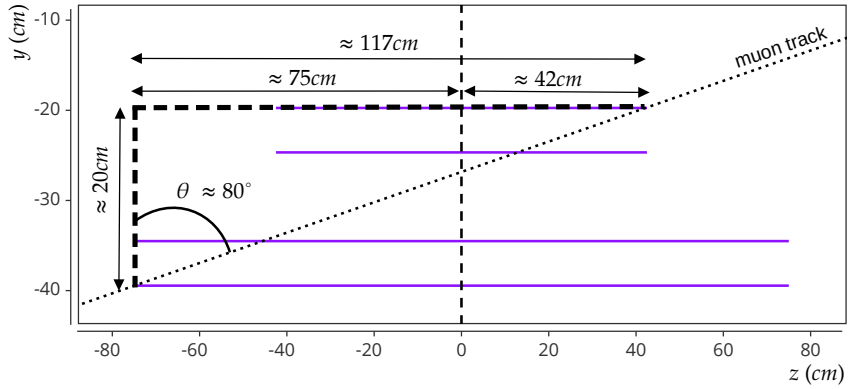

Figure 9.11: Schematic representing the maximum polar angle at $\phi = -180°, 0°, 180°$. Each half-layer is shown in blue, with layer 3 at the top and layer 6 at the bottom. For a cosmic muon track with an azimuthal angle of $\phi = -180°, 0°, 180°$, the maximum polar angle is approximately 80°.

9.5 Measurement of the detection efficiency

The detection efficiency is the ratio of measured signals to total expected signals. For example, a muon track traversing four layers of a perfectly efficient detector would result in a signal in each of the four layers. The detection efficiency of the active area of silicon sensors is close to 100% [25], meaning that a detector which had no dead areas should have a detection efficiency close to 100%. In practice, dead areas are present on most silicon detectors, in the form of non-responsive chips, or physical gaps between chips arising from the construction procedure. This analysis aims to determine the detection efficiency of the OB of the ITS2 in two cases: the detection efficiency of the active area of the OB, by making appropriate cuts to reject the known dead regions, as well as the intrinsic detection efficiency of the OB, where known dead regions are included. The MC physics simulations conducted in the ITS2 Technical Design Report use a baseline intrinsic detection efficiency of 95% [25] per layer.

Detection efficiency is defined as the number of tracks detected within a search window within a Region Of Interest (ROI) divided by the total number of tracks expected within that ROI:

$$e = \frac{n}{t}, \tag{9.13}$$

where e is the efficiency, n is the number of tracks detected within a search window within a ROI, and t is the total number of tracks expected within the ROI. An illustration of the criteria needed for a track to be counted as efficiently detected is shown in Fig. 9.12.

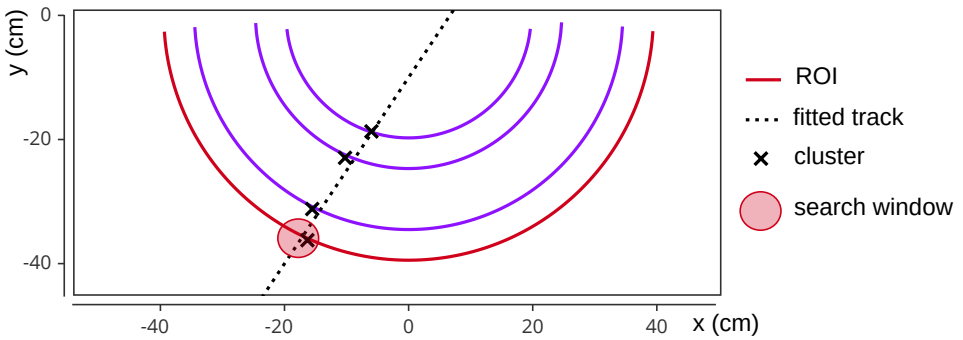

Figure 9.12: Illustration depicting the criteria for a track to be counted as efficiently detected. For a fitted track which is expected to hit all four half-layers, the expected cluster position of the track on the ROI is determined from the track fit. A search window (red circle) is expanded around this expected cluster position. A track is counted as being efficiently detected in the ROI if a cluster lies within this search window.

In this analysis, the ROI was chosen to be each half-layer. A single cosmic muon track must have clusters within the search window on at least three half-layers to be included in the efficiency measurement. The expected cluster position on each half-layer is determined from the track fit. If a cluster is present on each half-layer within the search window, then the track is classified as efficiently detected and both n and t are incremented by one for all ROIs. If the track has clusters within the search window for only three half-layers, the track is counted as being inefficiently detected for the ROI where no cluster is present. It is counted as being efficiently detected for the other three ROIs. In this case n and t are incremented by one for three ROIs, whilst only t is incremented by one for the ROI which did not have a cluster within the search window.

9.5.1 Choosing appropriate cosmic muon tracks

A number of cuts need to be made on the cosmic muon data set before making a measurement of the detection efficiency. The cuts used are described in this section.

An overview of the cuts used is displayed in Tab. 9.2.

Table 9.2: Overview of the purpose of the cuts made on the cosmic muon data set for the efficiency measurement.

Cut	Purpose
Convex-hull cut	Exclude tracks which are not expected to hit the ROI.
Cylinder cut	Exclude tracks which pass through the gaps between HSs.
Dead regions cut	Exclude tracks which are expected to hit dead chips.
Gaps between chips cut	Exclude tracks which pass through the gaps between ALPIDE chips.
χ^2 cut	Exclude tracks based on the goodness of the track fit.

Of the cuts described in Tab. 9.2, three are used in every analysis scenario presented in Sec. 9.6. All analysis scenarios require the χ^2 of a given track to be less than four. In addition, the cylinder cut and convex-hull cut are applied globally to ensure the a given track is expected to hit the ROI and to reduce the effect that the gaps between half-staves have on the efficiency measurement. Since the ITS2 is geometrically designed to accept tracks from the IP, the cosmic muon tracks which pass through these physical gaps between half-staves are not relevant for the efficiency measurement.

Convex-hull cut

A given track must be expected to hit the ROI. Examples of tracks which are not expected to hit the ROI are shown in Fig. 9.13. Given the geometry of the OB, some tracks could for instance intersect only the OLs and not the MLs, as shown in Fig. 9.13a. In addition, the curvature of the ITS2 leaves gaps between HSs through which cosmic muons can pass without detection, as illustrated in Fig. 9.13b.

To check that a given track hits the ROI in question, a geometric model of the OB was produced with the cosmic muon data. A convex hull was fitted around cluster positions for each HS. In addition a line of best fit was fitted to all cluster positions in each HS. A track that intersects the line of best fit will satisfy the following equations:

$$y = \frac{p_0 m + c}{1 - p_1 m}, \tag{9.14}$$

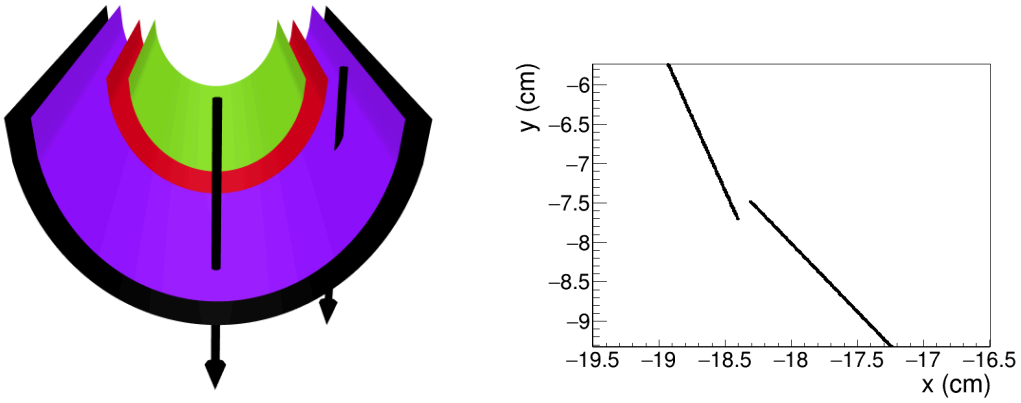

(a) **(b)**

Figure 9.13: Examples of regions of the OB where a track is not expected to be detected. (a) shows an illustration of two example muon tracks, illustrated by black arrows, which do not intersect all four half-layers. (b) shows the gap between two overlapping staves, where a muon could pass without detection.

$$x = p_0 + p_1 y, \tag{9.15}$$

where m and c are the gradient and y-intercept respectively for the line of best fit describing the clusters in a given HS. The coordinates of intersection were then checked to determine if they were within the convex hull fitted around all cluster positions. In addition, the z coordinate of intersection was determined and checked that it was within the z spatial limits of the layer in question. This process was repeated for all HSs in each layer. If a track intersected at least one HS in each layer it was considered as a valid track and included in the efficiency measurement.

Cylinder cut

The data set obtained is of course comprised of cosmic muons, and not from colliding beams at the interaction point inside the beam pipe. However, the ITS2 was never designed to be a cosmic muon detector and so its cosmic muon detection efficiency is not a performance property with much value. Its curvature leads to gaps between HSs that are not optimal for cosmic muon detection. It is known from pilot beam studies that particles originate from the interaction point extended by a sphere of radius $O(\mu\text{m})$. Requiring tracks to originate from the IP extended by a sphere of radius $O(\mu\text{m})$, however, results in too few tracks for a meaningful measurement of detection efficiency to be made. Extending the sphere by some radius to improve

statistics has the negative effect of biasing the angular distribution of tracks such that the tracks no longer pass through the interaction point.

The following cut was implemented to roughly replicate beam data in the xy plane and reduce the effect of gaps between staves on the efficiency measurement. Each track had to pass through a cylinder of radius 10 cm and length 40 cm in the z axis, centred about the origin. For muons whose tracks satisfy this requirement, the gaps between staves are not 'visible' due to the overlap of HSs. Figure 9.14 shows 1000 randomly selected tracks in OBBOT that pass the 'cylinder cut'.

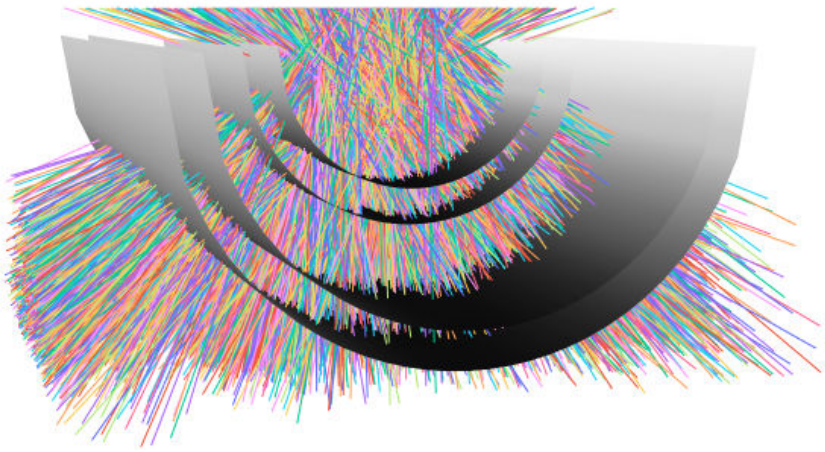

Figure 9.14: 1000 randomly selected tracks that pass the 'cylinder cut' described in the text.

Dead regions cut

The efficiency measurement must consider the dead regions of the detector. Dead chips, as well as staves that detect fewer hits than expected, can be seen in Fig. 9.9. Tracks for which extrapolated hits are located on a dead chip or a less responsive stave are discarded from the efficiency measurement. This is done by finding the coordinates of the regions in the cluster distribution shown in Fig. 9.9 and rejecting any tracks which are expected to hit them.

Gaps between chips cut

Small physical gaps between ALPIDE chips originate during the HIC and stave construction. These gaps appear in the cluster distribution. The gaps between chips are illustrated for layer three in Fig. 9.15, where each horizontal and vertical line

corresponds to a region where no hits are detected in the cosmic muon cluster distribution. The number of grid lines in z is 27. A HS in layer three has four HICs, so a total of three gaps between HICs visible in z. Each of the four HICs has six gaps between chips visible in z, therefore 27 gaps in total. There are 48 grid lines in ϕ. There are 24 staves in layer three, where each stave has two gaps between chips visible in ϕ, therefore 48 gaps.

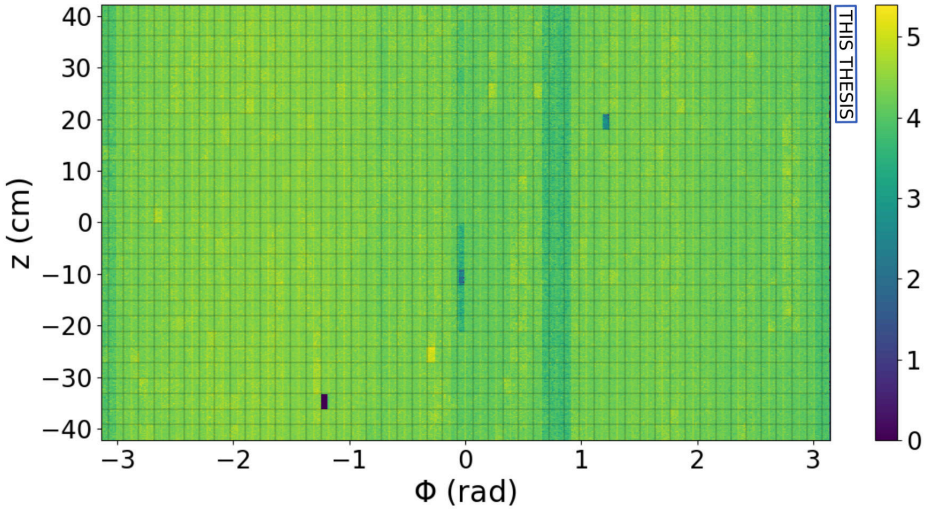

Figure 9.15: Raw cluster distribution for Layer 3. Areas of no clusters form a grid pattern, illustrated by the grey lines.

χ^2 cut

χ^2 is defined as follows:

$$\chi^2 = \sum_i \left[\frac{\Delta d_i}{\sigma_i}\right]^2, \tag{9.16}$$

where Δd_i is the distance from the fitted track to the cluster and σ_i is the expected error on the cluster position. The best estimate for the expected error on the cluster position is given by the residual marker distribution measured by the CMM and shown in Fig. 6.3. Therefore, σ_i is assumed to be equal for all clusters. The expression for χ^2 becomes:

$$\chi^2 = \frac{1}{\sigma^2}\sum_i (\Delta d_i)^2, \tag{9.17}$$

$$\chi^2 = \left(\frac{d_T}{\sigma}\right)^2, \tag{9.18}$$

where d_T is the sum of Δd_i. Figure 9.16 shows the average efficiency per half layer as a function of χ^2 for both half-barrels, with all previously mentioned cuts, i.e. the cylinder cut, the dead regions cut and the gaps between chips cut. As χ^2 increases, the efficiency decreases. A track with a higher χ^2 is less straight, therefore, above some χ^2 threshold, the expected cluster coordinates in the ROI are no longer accurate and fall outside of the search window. For $\chi^2 > 4$, the number of tracks is too low for a meaningful efficiency measurement.

Figure 9.17 shows the efficiency per half layer as a function of χ^2 up to a maximum χ^2 of four.

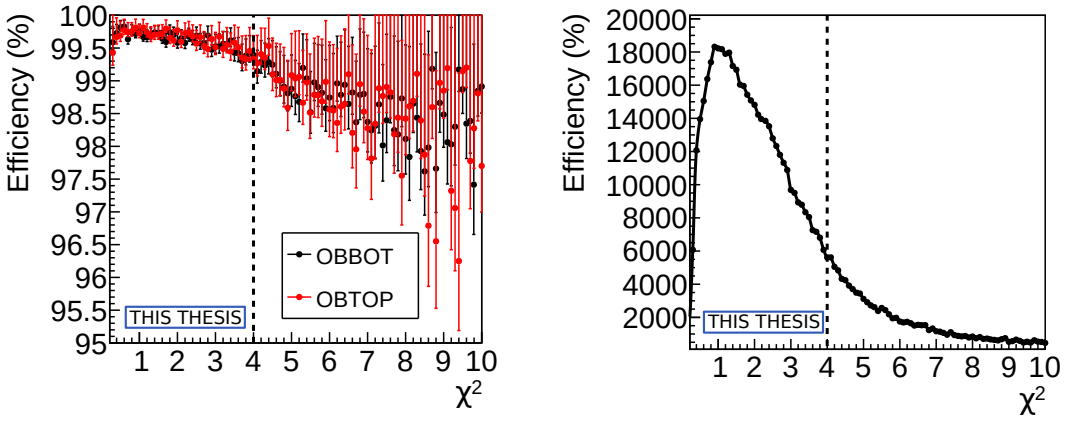

Figure 9.16: Average efficiency per half layer as a function of χ^2 for both half-barrels (left). The χ^2 distribution is shown on the right. The following cuts were used: the cylinder cut, the dead-regions cut and the gaps-between-chips cut.

9.6 Efficiency measurement results

The detection efficiency of the active area of the OB of the ITS2 is presented in Sec. 9.6.1. The detection efficiency of the OB of the ITS2 including the dead regions caused by dead chips and the physical gaps between chips is presented in Sec. 9.6.2.

The following cuts were applied to all analysis scenarios: the convex-hull cut, cylinder cut, and the χ^2 cut, where the χ^2 of a given track must be less than four. A cluster was considered to be inside the search window if the cluster had an absolute

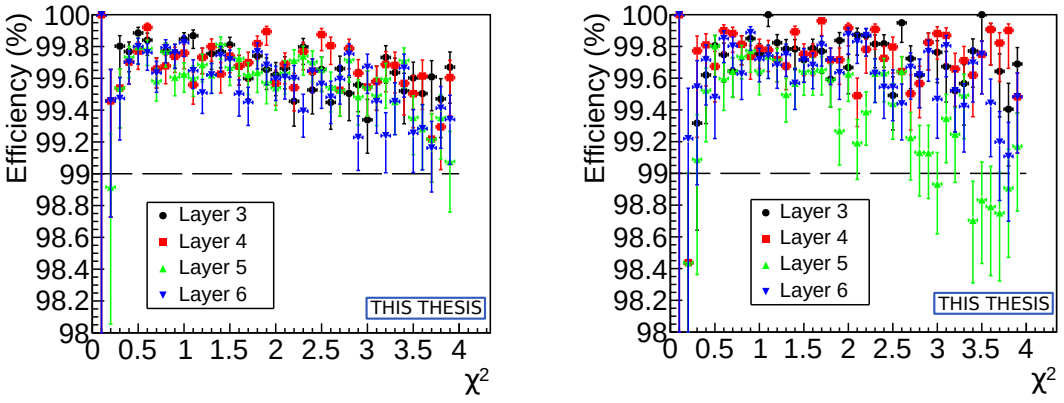

Figure 9.17: Efficiency as a function of χ^2 for both half-barrels for $\chi^2 \in [0,4]$. The following cuts were used: the cylinder cut, the dead-regions cut and the gaps-between-chips cut.

displacement, $|d|$, of less than 2 mm from the expected cluster position. Figure 9.18 shows the average efficiency per half layer as a function of $|d|$. The efficiency starts to drop as $|d|$ falls below 2 mm. The search window needs to be large enough to take into account the misalignment of the chips on the space frames, the misalignment of the staves in the barrel structure and the error on the extrapolated track. Below a $|d|$ of 2 mm, the efficiency starts to steadily decrease, suggesting real hits no longer fall within the search window since the search window is too small.

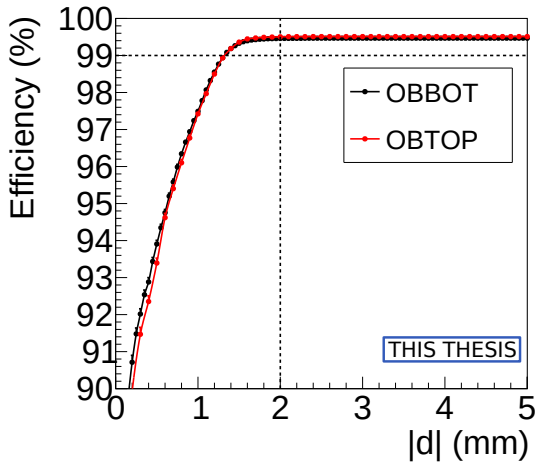

Figure 9.18: Average efficiency per half layer as a function of the $|d|$, the absolute displacement from the cluster to the extrapolated track. The search window is defined as a 2 mm sphere around the expected cluster position, as determined from the extrapolation of the track fit.

9.6.1 Efficiency of the active area

To gain a measure of the detection efficiency of the active area of the ITS2 OB the following cuts were used in addition to the global cuts: the dead-regions cut and the gaps-between-chips cut.

Table 9.3 lists the efficiency of each half-layer with the cuts described. The efficiency for each layer is above 99%, consistent with the detection efficiency measurement made of the ALPIDE chips with test-beam experiments [101].

Table 9.3: Efficiency the active area of each half-layer. OBBOT had 120171 total tracks after the cuts, whilst OBTOP had 73647 total tracks after the cuts.

Half-layer	Efficiency (%)
L3_BOT	$99.67 + 0.02 - 0.02$
L4_BOT	$99.68 + 0.02 - 0.02$
L5_BOT	$99.62 + 0.02 - 0.02$
L6_BOT	$99.60 + 0.02 - 0.03$
L3_TOP	$99.74 + 0.02 - 0.02$
L4_TOP	$99.75 + 0.02 - 0.02$
L5_TOP	$99.44 + 0.03 - 0.03$
L6_TOP	$99.67 + 0.02 - 0.02$

9.6.2 Efficiency including dead regions

To gain a measure of the intrinsic detection efficiency of the ITS2 OB, only the global cuts were applied. The dead-regions cut and the gaps-between-chips cut were not applied. Since the dead-regions represent intrinsically inefficient areas of the OB of the ITS2, this measurement represents the intrinsic detection efficiency of the OB of the ITS. Table 9.4 shows the intrinsic detection efficiency per half-layer. A comparison between the detection efficiency of the active area and the intrinsic detection efficiency for each half-layer is shown in Fig. 9.19.

9.6.3 Conclusion

The detection efficiency of the active area of the OB of the ITS2 was measured to be greater than 99% using a cosmic muon data set, consistent with the measured detection efficiency of the ALPIDE chip obtained from test-beam measurements [101]. The detection efficiency of the OB of the ITS2, including the intrinsically

Table 9.4: Efficiency of each half-layer. OBBOT had 176773 total tracks after the cuts, whilst OBTOP had 114499 total tracks after the cuts.

Half-layer	Efficiency (%)
L3_BOT	$99.27 + 0.02 - 0.02$
L4_BOT	$98.23 + 0.03 - 0.03$
L5_BOT	$99.00 + 0.02 - 0.02$
L6_BOT	$98.05 + 0.03 - 0.03$
L3_TOP	$98.46 + 0.04 - 0.04$
L4_TOP	$99.60 + 0.02 - 0.02$
L5_TOP	$98.64 + 0.03 - 0.04$
L6_TOP	$97.98 + 0.04 - 0.04$

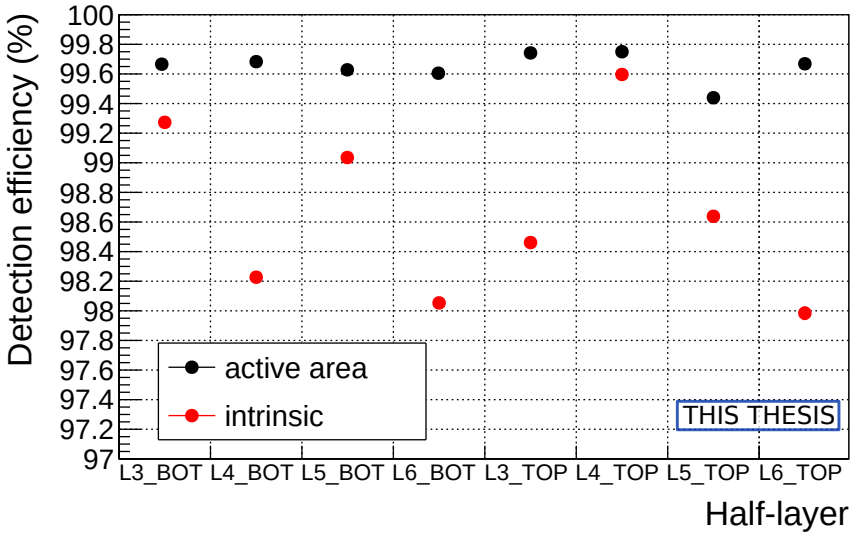

Figure 9.19: Detection efficiency per half-layer. The detection efficiency of the active area is shown with black points. The intrinsic detection efficiency for the OB of the ITS2 is shown in red. The intrinsic detection efficiency includes dead areas such as non-responsive chips and the physical gaps between chips.

inefficient regions of non-responsive chips and the physical gaps between chips was found to be greater than 97.94%.

Five cuts were used in the analysis. Of these, three were used in all analysis scenarios. These three cuts were the convex-hull cut, the cylinder cut and the χ^2 cut. The latter was used to ensure only tracks with an appropriate goodness of fit were included in the efficiency measurement. The two former cuts guaranteed a track was expected to hit the ROI. Two further cuts were used to determine if a given track hits the active area of the detector. These two cuts were the dead-regions cut and the gaps-between-chips cut.

The measured intrinsic detection efficiency of the OB of the ITS2 of greater than 97.94% exceeds the baseline detection efficiency of 95% foreseen and used in the physics simulations described in Sec. 4.4.2.

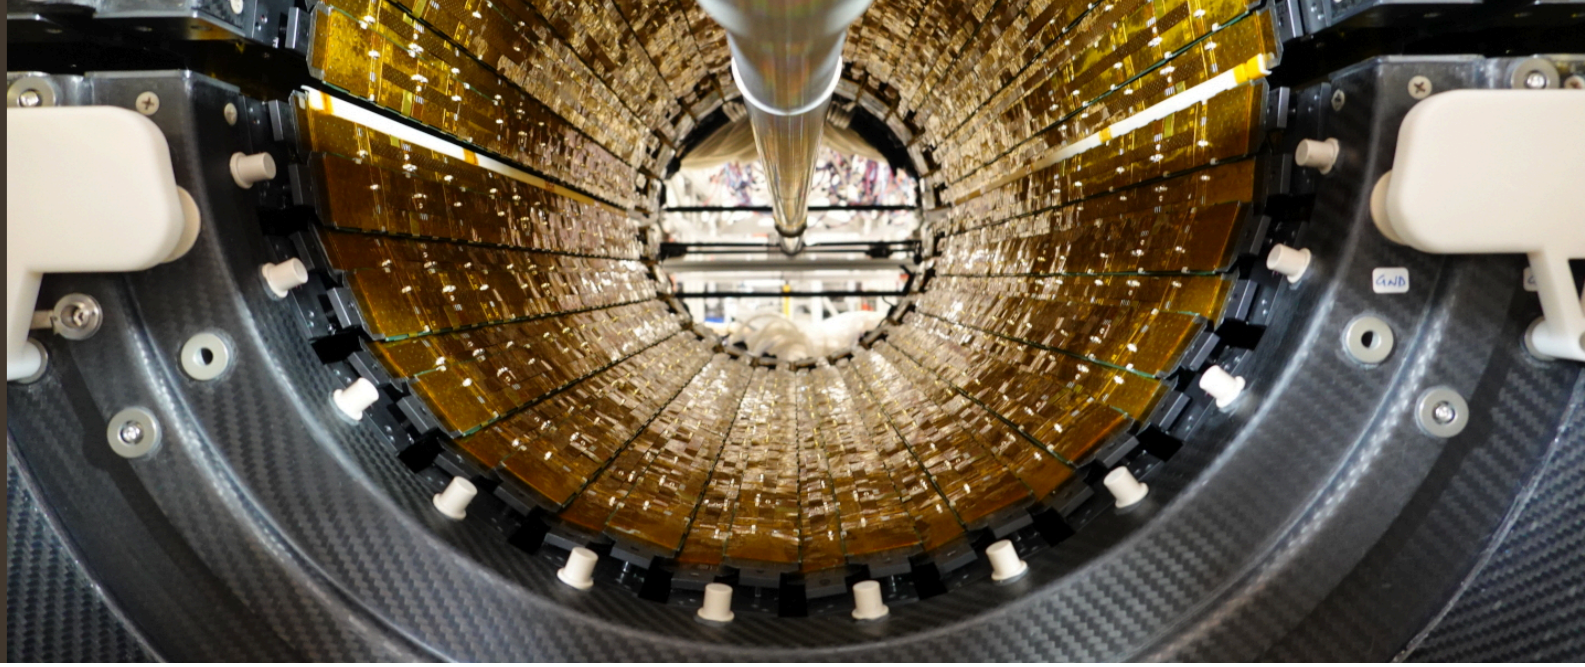

Chapter 10:
Conclusions

A Large Ion Collider Experiment (ALICE) aims to characterise the Quark-Gluon Plasma (QGP) by studying relativistic heavy-ion collisions produced by the Large Hadron Collider (LHC). Several observables of the QGP exist. Some of these observables, such as heavy-flavour hadrons, lose energy to the medium, meaning they traverse the ALICE detector with low transverse momentum (p_T). For this reason, the Inner Tracking System (ITS) of ALICE needs to have a minimum material budget, to minimise the energy loss of low p_T particles such that they can traverse multiple tracking layers before being lost due to momentum cut-off. In addition, an ITS with a high impact parameter resolution and high vertexing capabilities is needed to measure short-lived heavy-flavour hadrons such as the charmed baryon, Λ_c, the measurement of which will help to understand the nature of hadronisation and its dependence on quark coalescence.

High vertexing capabilities can be achieved with an ITS whose spatial resolution is high and whose first layer is close to the interaction point. The LHC will deliver Pb-Pb collisions with an interaction rate of 50 kHz after Long Shutdown 2 (LS2). Due to the large combinatorial background of relativistic heavy-ion collisions, and the increase in luminosity expected after LS2, triggering is inefficient, necessitating an ITS with a readout rate high enough to cope with recording all events.

Header photograph shows the fully installed Outer Barrel of ITS2. © CERN. Taken from [2].

The ALICE ITS has been successfully upgraded over LS2. The entirely new tracker, ITS2, uses a unique Monolithic Active Pixel Sensors (MAPS) design known as ALPIDE throughout. The ALPIDE chip has a readout rate twice the expected Pb-Pb interaction rate after LS2, a spatial resolution of 5 µm and a detection efficiency greater than 99%. The ALPIDE chips are arranged into Hybrid Integrated Circuits (HICs) and staves, which form the seven layers in the concentric barrel structure of the tracker. The material budget is $0.35\%\,X_0$ for the innermost three layers, the Inner Barrel (IB), and $1.1\%\,X_0$ for the outermost four layers, the Outer Barrel (OB). ITS2 is expected to improve the resolution of the track impact parameter by a factor of three for p_T below $1\,\mathrm{GeV}/c$ and a factor of six for p_T above $10\,\mathrm{GeV}/c$ compared to the ITS used for LHC Runs 1 and 2. A measurement of the Λ_c was unachievable at the limit of the capabilities of the previous ITS, but detailed new measurements will become possible with good precision down to a p_T of $2\,\mathrm{GeV}/c$ with ITS2.

OB HICs were successfully constructed in Bari, Liverpool, Strasbourg and Wuhan, whilst OB staves were successfully constructed in Amsterdam, Berkeley, Daresbury, Frascatti and Turin. The construction procedure was outlined in this thesis. A characterisation of OB staves was presented in this thesis, where it was found that all constructed staves fulfilled the minimum metrological requirements, sufficient for a successful installation within the barrel structure of the ITS2.

The constructed half-barrels of the ITS2 were connected to the readout and cooling system in an above-ground laboratory before undergoing an extensive period of commissioning. The Fake-Hit Rate (FHR) per stave was measured to be $(1.4 \pm 0.3) \times 10^{-10}$ /pixel/event for the most optimum masking scenario.

A cosmic muon acquisition campaign was conducted for the OB of the ITS2, from which approximately 5×10^6 cosmic muon tracks were reconstructed. From these data, a measurement of the detection efficiency of the active area of the OB was made, where it was found to be greater than 99% for the active area of all half-layers, and greater than 97.94% for the intrinsic detection efficiency of all half-layers of the OB of the ITS2, exceeding the foreseen baseline detection efficiency of 95% [25].

The ITS2 is successfully installed around the LHC beam-pipe, a process illustrated chronologically by the photographs displayed in each chapter header within this thesis. The ITS2 has $10\,\mathrm{m}^2$ of active area segmented into 12.5 billion pixels, making it the largest silicon pixel detector ever constructed. The ITS2 is the first detector based solely on MAPS technology. The ITS2 has undergone a commissioning

campaign in-situ around the beam-pipe, including data taking with 450 GeV LHC so-called pilot beams providing low-intensity collision events.

A Letter of Intent has recently been submitted for ITS3, which will replace the ITS2 IB during the LHC Long Shutdown 3 (LS3) for data taking in LHC Run 4.

The silicon sensor in the ITS2 contributes only 15% to the total material budget. As shown in Tab. 4.2 the FPC contributes 50%, the cooling circuit contributes 20% and the carbon fibre space frame contributes 15%. The goal of the next iteration of the ALICE ITS, ITS3, is to cut the material budget significantly by removing the FPC, cooling and carbon fibre space frame, resulting in a detector constructed out of silicon sensors only. The envisioned detector would be made of an assembly of a small number of large scale silicon sensors, cooled via air-flow with minimal mechanical support structure.

Silicon, despite being brittle at thicknesses of approximately 100 µm becomes bendy at thicknesses of around 50 µm. ALPIDE chips, with dimensions of 1.5cm $\times$ 3cm, have been successfully thinned to 50 µm and bent to a radii of 22mm without signs of mechanical or electrical damage [102]. The bent ALPIDE chips had charge thresholds unaffected by the bending procedure and a detection efficiency of greater than 99.9% [102].

In conventional CMOS manufacturing, the maximum chip size is limited by the field of view of the photolithographic process, typically a few cm^2. CMOS stitching allows sensors to be fabricated with dimensions limited only by the size of the wafer. Power and electrical signal distribution can be done inside the silicon chip, eliminating the need for an FPC.

The ALICE ITS3 will consist of three layers of curved wafer-scale sensors, with a material budget of $0.05\% X_0$ per layer and a first layer radius of 18 mm, resulting in a significant improvement to the tracking precision and efficiency at low p_T .

Lists

List of acronyms

AERD Address Encoder Reset Decoder

ALICE A Large Ion Collider Experiment

ALPIDE ALice PIxel DEtector

CMM Coordinate Measuring Machine

CMOS Complementary Metal Oxide Semiconductor

CMU Control Management Unit

CP Cold Plate

CPV Charged Particle Veto detector

CRAM Configuration Random-Access Memory

CRU Common Readout Unit

CTP Central Trigger Processor

DCal Di-jet Calorimeter

DDR Double Data Rate

DMA Direct Memory Access

DMU Data Management Unit

DTU Data Transmission Unit

EMCal ElectroMagnetic Calorimeter

EPN Event Processing Node

ESD Electro-Static Discharge

FHR Fake-Hit Rate

FIFO First In First Out

FLP First Level Processor

FMC FPGA Mezzanine Card

FPC Flexible Printed Circuit

FPN Fixed Pattern Noise

GEM Gas Electron Multiplier

HF Heavy-Flavour

HIC Hybrid Integrated Circuit

HMPID High-Momentum Particle Identification Detector

HS Half-Stave

HS_L Half-Stave Lower

HS_U Half-Stave Upper

IB Inner Barrel

ITS Inner Tracking System

LH Left-Handed

LHC Large Hadron Collider

LILO Last In Last Out

LS2 Long Shutdown 2

LSDC Liverpool Semiconductor Detector Centre

LTU Local Trigger Unit

LVDS Low Voltage Differential Signal

MAPS Monolithic Active Pixel Sensors

MCS Multiple Coulomb Scattering

MEB Multi-Event Buffer

MFT Muon Forward Tracker

MIP Minimum Ionising Particle

ML Middle Layers

M-LVDS Multipoint Low Voltage Differential Signal

MOSAIC MOdular System for Acquisition Interface and Control

MWPC Multi-Wire Proportional Chamber

NIEL Non-Ionising Energy Loss

NMOS n-channel Metal Oxide Semiconductor

OB Outer Barrel

OBBOT Bottom half of the Outer Barrel

OBTOP Top half of the Outer Barrel

OL Outer Layers

PB Power Board

PHOS PHOton Spectrometer

PID Particle IDentification

PLL Phase-Locked Loop

PMOS p-channel Metal Oxide Semiconductor

PRBS Pseudo-Random Bit Sequence

QC Quality Control

QCD Quantum Chromodynamics

QED Quantum Electro-Dynamics

QGP Quark-Gluon Plasma

RC Time constant

RDH Raw Data Header

RH Right-Handed

RRU Region Readout Unit

GBT-SCA GBT-Slow Control Adapter

SDD Silicon Drift Detector

SEU Single Event Upsets

SF Space Frame

SM Standard Model

SPD Silicon Pixel Detector

SRAM Static Random-Access Memory

SSD Silicon Strip Detector

SWT Single Word Transactions

TID Total Ionising Dose

TN Temporal Noise

TOF Time Of Flight detector

TPC Time Projection Chamber

TRD Transition Radiation Detector

UDP User Datagram Protocol

VRx Versatile Receiver

VTRx Versatile TransReceiver

VTTx Versatile Twin-Transceiver

List of Figures

List of Tables

Appendices

Appendix A

Derivation of transverse momentum resolution

Figure A.1 shows the geometry of a particle traversing a region with a magnetic field, B. The particle of charge q moves along an arc of radius R, given by:

$$R = \frac{p_T}{qB}. \tag{A.1}$$

Geometrically, R, can be expressed as:

$$R^2 = \left(\frac{L}{2}\right)^2 + (R-s)^2, \tag{A.2}$$

$$R = \frac{s}{2} + \frac{L^2}{8s}. \tag{A.3}$$

For $s \ll R$:

$$s \approx \frac{L^2}{8R}. \tag{A.4}$$

The transverse momentum resolution, σ_{p_T}, is given by:

$$\sigma_{p_T} = \frac{\partial p_T}{\partial s}\sigma_s, \tag{A.5}$$

$$\sigma_{p_T} = \frac{\partial}{\partial s}\left(\frac{qBL^2}{8s}\right)\sigma_s. \tag{A.6}$$

Therefore, the normalised transverse momentum resolution depends on $\mathrm{p_T}$, L,

B and the resolution of the sagitta, σ_s:

$$\frac{\sigma_{p_T}}{p_T} \propto \frac{p_T}{BL^2}\sigma_s. \tag{A.7}$$

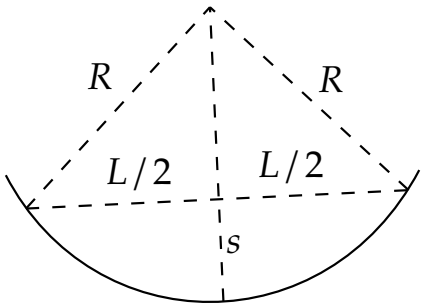

Figure A.1: Geometry of a particle moving in a magnetic field.

Appendix B

Raw cluster distributions from cosmic muon data set

The raw cluster distribution for each layer in the Outer Barrel are shown here. This distribution is shown for layer three in Fig. B.1, layer four in Fig. B.2, layer five in Fig. B.3 and layer six in Fig. B.4. The total number of masked chips visible in the cluster distributions is 40, corresponding to the 40 intentionally masked chips in the cosmic muon aquisition scans.

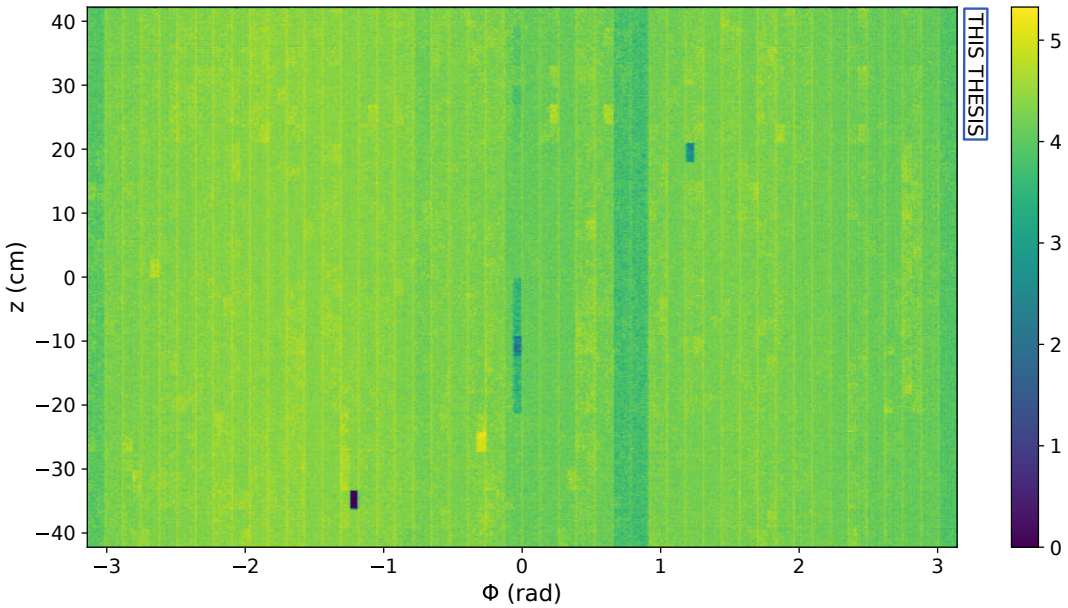

Figure B.1: Raw cluster distribution of all cosmic events for layer three.

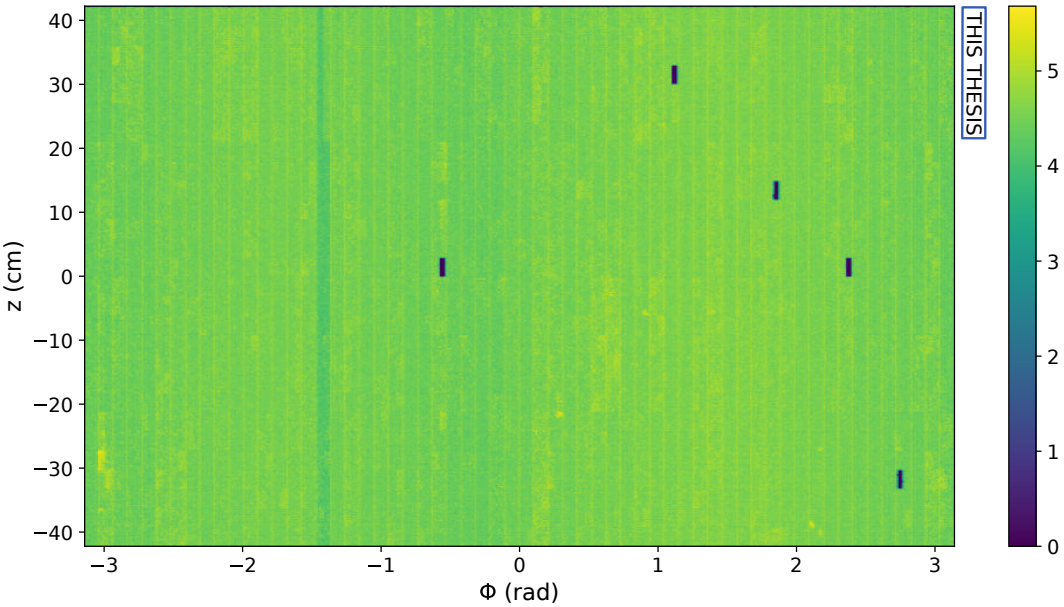

Figure B.2: Raw cluster distribution of all cosmic events for layer four.

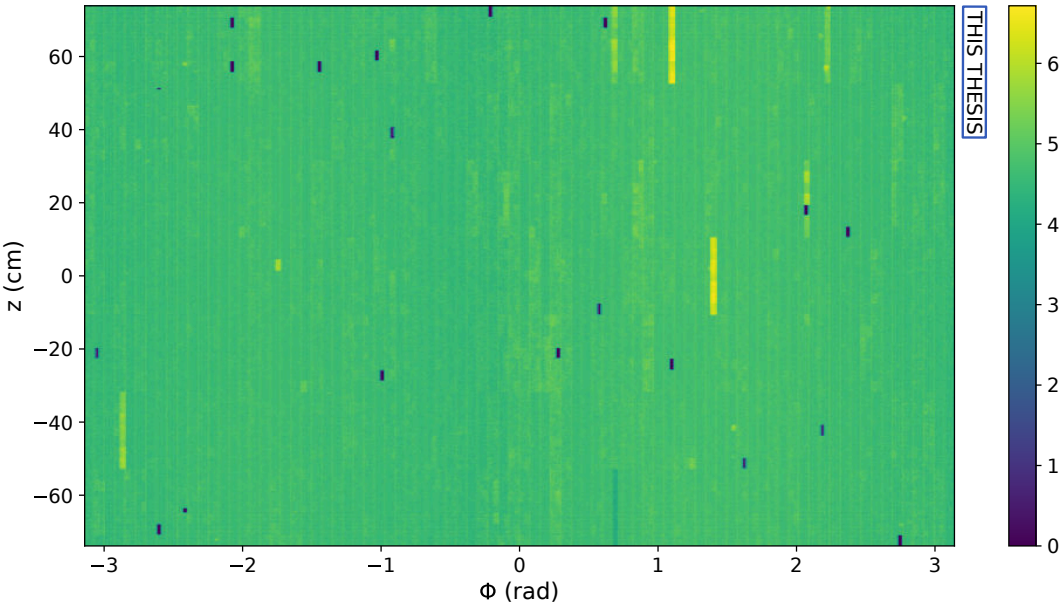

Figure B.3: Raw cluster distribution of all cosmic events for layer five.

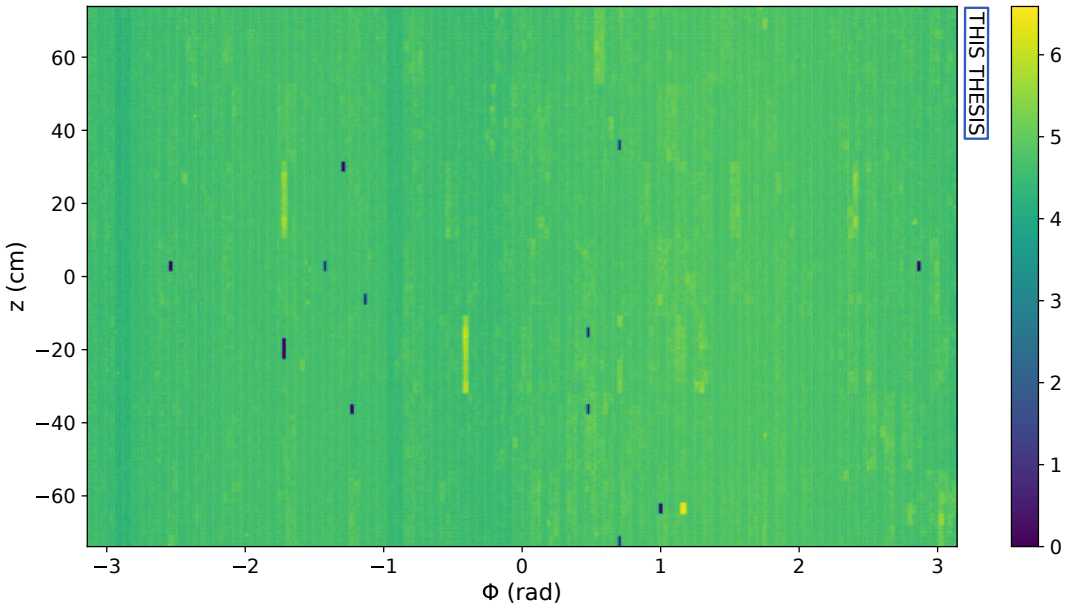

Figure B.4: Raw cluster distribution of all cosmic events for layer six.

Appendix C

Derivation of distance from cluster to line fit

A line in three dimensions can be specified by two points, $\vec{x_0} = (x_0, y_0, z_0)$ and $\vec{x_1} = (x_1, y_1, z_1)$. A vector, $\vec{v}$, along this line is given by [100]:

$$\vec{v} = \begin{pmatrix} x_0 + (x_1 - x_0)t \\ y_0 + (y_1 - y_0)t \\ z_0 + (z_1 - z_0)t \end{pmatrix}. \tag{C.1}$$

The squared distance between a point on the line described by the parameter t and another point, not necessarily on the line, $\vec{x_p} = (x_p, y_p, z_p)$ is therefore:

$$\begin{aligned} d^2 = & [(x_0 - x_p) + (x_1 - x_0)t]^2 \\ & + [(y_0 - y_p) + (y_1 - y_0)t]^2 \\ & + [(z_0 - z_p) + (z_1 - z_0)t]^2 . \end{aligned} \tag{C.2}$$

The most appropriate value of t is where d^2 is minimum, i.e. when $\frac{d(d^2)}{dt} = 0$. Equation C.2 is written in a more easily differentiable form below:

$$\begin{aligned} d^2 = & \left[(x_0 - x_p)^2 + (y_0 - y_p)^2 + (z_0 - z_p^2)\right] \\ & + 2t\left[(x_1 - x_0)(x_0 - x_p) + (y_1 - y_0)(y_0 - y_p) + (z_1 - z_0)(z_0 - z_p)\right] \\ & + t^2\left[(x_1 - x_0)^2 + (y_1 - y_0)^2 + (z_1 - z_0)^2\right] . \end{aligned} \tag{C.3}$$

Differentiating Eq. C.3 with respect to t results in the following equation:

$$\frac{d(d^2)}{dt} = 2\left[(x_1-x_0)(x_0-x_p)+(y_1-y_0)(y_0-y_p)+(z_1-z_0)(z_0-z_p)\right] + 2t\left[(x_1-x_0)^2+(y_1-y_0)^2+(z_1-z_0)^2\right]. \tag{C.4}$$

The most appropriate value of t is where d^2 is minimum. Therefore, setting $\frac{d(d^2)}{dt} = 0$ and solving for t, results in the following appropriate value for t:

$$t = \frac{-\left[(x_1-x_0)(x_0-x_p)+(y_1-y_0)(y_0-y_p)+(z_1-z_0)(z_0-z_p)\right]}{\left[(x_1-x_0)^2+(y_1-y_0)^2+(z_1-z_0)^2\right]}. \tag{C.5}$$

The following identities are useful for rearranging Eq. C.5 into a neater form:

$$\vec{x}\cdot\vec{y} = \sum x_i y_i, \tag{C.6}$$

$$\vec{x}\cdot\vec{y} = \vec{y}\cdot\vec{x}, \tag{C.7}$$

$$|\vec{x_1}-\vec{x_0}| = \sqrt{(x_1-x_0)^2+(y_1-y_0)^2+(z_1-z_0)^2}. \tag{C.8}$$

Using Eq. C.6 and Eq. C.7, the numerator of Eq. C.5, α, can be written as:

$$\alpha = -(\vec{x_0}-\vec{x_p})\cdot(\vec{x_1}-\vec{x_0}). \tag{C.9}$$

Using Eq. C.8, the denominator of Eq. C.5, β, can be written as:

$$\beta = |\vec{x_1}-\vec{x_0}|^2. \tag{C.10}$$

Dividing Eq. C.9 by Eq. C.10, i.e. $t = \frac{\alpha}{\beta}$, results in the following form of Eq. C.5:

$$t = \frac{-(\vec{x_0}-\vec{x_p})\cdot(\vec{x_1}-\vec{x_0})}{|\vec{x_1}-\vec{x_0}|^2}. \tag{C.11}$$

Using Eq. C.6, Eq. C.7 and Eq. C.8, Eq. C.3 can be rewritten as:

$$d^2 = |\vec{x_0}-\vec{x_p}|^2 + 2t\left[(\vec{x_1}-\vec{x_0})\cdot(\vec{x_0}-\vec{x_p})\right] + t^2\left[|\vec{x_1}-\vec{x_0}|^2\right] \tag{C.12}$$

Substituting Eq. C.11 into Eq. C.12 results in:

$$
\begin{aligned}
d^2 = {} & |\vec{x_0} - \vec{x_p}|^2 \\
& - 2\frac{(\vec{x_0} - \vec{x_p}) \cdot (\vec{x_1} - \vec{x_0})}{|\vec{x_1} - \vec{x_0}|^2} \left[(\vec{x_0} - \vec{x_p}) \cdot (\vec{x_1} - \vec{x_0})\right] \\
& + \left(\frac{(\vec{x_0} - \vec{x_p}) \cdot (\vec{x_1} - \vec{x_0})}{|\vec{x_1} - \vec{x_0}|^2}\right)^2 |\vec{x_1} - \vec{x_0}|^2,
\end{aligned}
\tag{C.13}
$$

$$
\begin{aligned}
d^2 = {} & |\vec{x_0} - \vec{x_p}|^2 \\
& - 2\frac{\left[(\vec{x_0} - \vec{x_p}) \cdot (\vec{x_1} - \vec{x_0})\right]^2}{|\vec{x_1} - \vec{x_0}|^2} \\
& + \frac{\left[(\vec{x_0} - \vec{x_p}) \cdot (\vec{x_1} - \vec{x_0})\right]^2}{|\vec{x_1} - \vec{x_0}|^2},
\end{aligned}
\tag{C.14}
$$

$$
d^2 = |\vec{x_0} - \vec{x_p}|^2 - \frac{\left[(\vec{x_0} - \vec{x_p}) \cdot (\vec{x_1} - \vec{x_0})\right]^2}{|\vec{x_1} - \vec{x_0}|^2}, \tag{C.15}
$$

$$
d^2 = \frac{|\vec{x_0} - \vec{x_p}|^2 |\vec{x_1} - \vec{x_0}|^2 - \left[(\vec{x_0} - \vec{x_p}) \cdot (\vec{x_1} - \vec{x_0})\right]^2}{|\vec{x_1} - \vec{x_0}|^2}. \tag{C.16}
$$

The vector quadruple product is defined as:

$$
(\vec{A} \times \vec{B})^2 = \vec{A}^2 \vec{B}^2 - (\vec{A} \cdot \vec{B})^2 \tag{C.17}
$$

Using Eq. C.17, Eq. C.16 can be written in its final form as:

$$
d^2 = \frac{|(\vec{x_1} - \vec{x_0}) \times (\vec{x_0} - \vec{x_p})|^2}{|\vec{x_1} - \vec{x_0}|^2} \tag{C.18}
$$

Bibliography

[1] M. Brice and J. M. Ordan. *Photograph: ALICE Inner Tracking System (ITS) detectors removal*. 2019. URL: http://cds.cern.ch/record/2663202?ln=en.

[2] F. Reidt, J. Klein, and M. A. Imhoff. *Photograph: ALICE ITS Outer Barrel Installation*. 2021. URL: http://cds.cern.ch/record/2758788.

[3] M. Brice. *Photograph: ALICE ITS upgrade detector and subsystems*. 2020. URL: http://cds.cern.ch/record/2746491.

[4] M. Wilde. "Measurement of Direct Photons in pp and Pb-Pb Collisions with ALICE". In: *Nucl. Phys. A* 904-905 (2013), pp. 573c–576c. DOI: 10.1016/j.nuclphysa.2013.02.079. arXiv: 1210.5958.

[5] R. Snellings. "Elliptic flow: A brief review". In: *New J. Phys.* 13.June (2011). DOI: 10.1088/1367-2630/13/5/055008. arXiv: 1102.3010.

[6] CERN. *CERN Accelerating science*. URL: https://press.cern/press-releases/2000/02/new-state-matter-created-cern.

[7] M. Gyulassy. "The QGP discovered at RHIC". In: *Struct. Dyn. Elem. matter. Proceedings, NATO Adv. Study Institute, Camyuva-Kemer, Turkey, Sept. 22-October 2, 2003*. 2004, pp. 159–182. arXiv: nucl-th/0403032 [nucl-th].

[8] The ALICE Collaboration. "Upgrade of the ALICE Experiment: Letter Of Intent". In: *J. Phys.* G41 (2014), p. 87001. DOI: 10.1088/0954-3899/41/8/087001.

[9] The ALICE Collaboration. "Technical Design Report for the Upgrade of the ALICE Inner Tracking System". In: *J. Phys. G Nucl. Part. Phys.* 41.8 (2014), p. 87002. URL: http://stacks.iop.org/0954-3899/41/i=8/a=087002.

[10] L. Musa for The ALICE Collaboration. *Conceptual Design Report for the Upgrade of the ALICE ITS*. Tech. rep. CERN-LHCC-2012-005. LHCC-G-159. Geneva: CERN, 2012. URL: https://cds.cern.ch/record/1431539.

[11] L. Greiner et al. "A MAPS based vertex detector for the STAR experiment at RHIC". In: *Nucl. Instrum. Meth.* A650 (2011), pp. 68–72. DOI: 10.1016/j.nima.2010.12.006.

[12] P. D. Group et al. "Review of Particle Physics". In: *Prog. Theor. Exp. Phys.* 2020.8 (2020). DOI: 10.1093/ptep/ptaa104. URL: https://doi.org/10.1093/ptep/ptaa104.

[13] H. Bichsel. "Straggling in thin silicon detectors". In: *Rev. Mod. Phys.* 60.3 (1988), pp. 663–699. DOI: 10.1103/RevModPhys.60.663.

[14] K. Nakamura et al. "Review of particle physics". In: *J. Phys. G Nucl. Part. Phys.* 37.7 A (2010), pp. 3–5. DOI: 10.1088/0954-3899/37/7A/075021.

[15] A. H. Compton. "A Quantum Theory of the Scattering of X-rays by Light Elements". In: *Phys. Rev.* 21.5 (1923), pp. 483–502. DOI: 10.1103/PhysRev.21.483.

[16] C. Kittel. *Introduction to solid state physics*. New York: John Wiley and Sons, 2004.

[17] H. Spieler. *Introduction to Radiation Detectors and Electronics*. UC Berkeley, 1998.

[18] K. K. N. Simon M. Sze. "Physics and Properties of Semiconductors—A Review". In: *Phys. Semicond. Devices*. John Wiley and Sons, 2006. DOI: https://doi.org/10.1002/9780470068328.ch1. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/9780470068328.ch1.

[19] L. Rossi. *Pixel Detectors: From Fundamentals to Applications*. Particle Acceleration and Detection. Springer, 2006. URL: https://books.google.co.uk/books?id=Jbp73yTz-LYC.

[20] G. Lutz. *Semiconductor radiation detectors*. 2007. DOI: 10.1007/978-3-540-71679-2.

[21] W. Shockley. “The Theory of p-n Junctions in Semiconductors and p-n Junction Transisters”. In: *Bell Syst. Tech. J.* 28.3 (1949), pp. 435–489. DOI: 10.1002/j.1538-7305.1949.tb03645.x..

[22] C. Grupen and B. Shwartz. “Characteristic properties of detectors”. In: *Part. Detect.* (2009), pp. 56–70. DOI: 10.1017/cbo9780511534966.005.

[23] E. Gatti and P. F. Manfredi. “Processing the signals from solid-state detectors in elementary-particle physics”. In: *La Riv. del Nuovo Cim.* 9.1 (1986), pp. 1–146. DOI: 10.1007/BF02822156.

[24] M. Turala. “Silicon tracking detectors - Historical overview”. In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 541.1-2 (2005), pp. 1–14. DOI: 10.1016/j.nima.2005.01.032.

[25] The ALICE Collaboration. “Technical Design Report of the ITS Upgrade”. In: *J. Phys. G Nucl. Part. Phys.* 41 (2014), pp. 1–181. URL: http://cds.cern.ch/record/1625842.

[26] J. M. Ordan. *Photograph: ALICE ITS Outer Barrel installation*. 2021. URL: http://cds.cern.ch/record/2756341.

[27] M. Gell-Mann. “A Schematic Model of Baryons and Mesons”. In: *Phys. Lett.* 8.3 (1964), pp. 214–215. DOI: 10.1016/S0031-9163(64)92001-3.

[28] G. Zweig. “An SU_3 model for strong interaction symmetry and its breaking; Version 2”. In: *CERN-TH-412* February (1964), 80 p. URL: https://cds.cern.ch/record/570209.

[29] E. D. Bloom et al. “High-Energy Inelastic e-p Scattering at 6° and 10°”. In: *Phys. Rev. Lett.* 23.16 (1969), pp. 930–934. DOI: 10.1103/PhysRevLett.23.930. URL: https://link.aps.org/doi/10.1103/PhysRevLett.23.930.

[30] M. Breidenbach et al. “Observed Behavior of Highly Inelastic Electron-Proton Scattering”. In: *Phys. Rev. Lett.* 23.16 (1969), pp. 935–939. DOI: 10.1103/PhysRevLett.23.935. URL: https://link.aps.org/doi/10.1103/PhysRevLett.23.935.

[31] S. L. Glashow, J. Iliopoulos, and L. Maiani. “Weak interactions with lepton-hadron symmetry”. In: *Phys. Rev. D* 2.7 (1970), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285.

[32] J. J. Aubert et al. “Experimental observation of a heavy particle J”. In: *Phys. Rev. Lett.* 33.23 (1974), pp. 1404–1406. DOI: 10.1103/PhysRevLett.33.1404.

[33] J. E. Augustin et al. “Discovery of a narrow resonance in e+e- annihilation”. In: *Phys. Rev. Lett.* 33.23 (1974), pp. 1406–1408. DOI: 10.1103/PhysRevLett.33.1406.

[34] S. W. Herb et al. “Observation of a dimuon resonance at 9.5 GeV in 400-GeV proton-nucleus collisions”. In: *Phys. Rev. Lett.* 39.5 (1977), pp. 252–255. DOI: 10.1103/PhysRevLett.39.252. URL: https://link.aps.org/doi/10.1103/PhysRevLett.39.252.

[35] The CDF Collaboration. “Observation of Top Quark Production in $p\bar{p}$ Collisions with the Collider Detector at Fermilab”. In: *Phys. Rev. Lett.* 74.14 (1995), pp. 2626–2631. DOI: 10.1103/PhysRevLett.74.2626. URL: https://link.aps.org/doi/10.1103/PhysRevLett.74.2626.

[36] O. W. Greenberg. “Spin and unitary-spin independence in a paraquark model of baryons and mesons”. In: *Phys. Rev. Lett.* 13.20 (1964), pp. 598–602. DOI: 10.1103/PhysRevLett.13.598.

[37] D. J. Gross and F. Wilczek. “Asymptotically free gauge theories. I”. In: *Phys. Rev. D* 8.10 (1973), pp. 3633–3652. DOI: 10.1103/PhysRevD.8.3633.

[38] F. Wilczek. “Asymptotically free gauge theories. II*”. In: *Phys. Rev. D* 9.4 (1974).

[39] D. J. Gross and F. Wilczek. “Ultraviolet behavior of non-abelian gauge theories”. In: *Phys. Rev. Lett.* 30.26 (1973), pp. 1343–1346. DOI: 10.1103/PhysRevLett.30.1343.

[40] H. D. Politzer. “Reliable perturbative results for strong interactions?” In: *Phys. Rev. Lett.* 30.26 (1973), pp. 1346–1349. DOI: 10.1103/PhysRevLett.30.1346.

[41] Nobel Nobel Prize Outreach. *The Nobel Prize in Physics 2004*. 2022. URL: https://www.nobelprize.org/prizes/physics/2004/summary/ (visited on 03/10/2022).

[42] H. Satz. *Ultimate Horizons: Probing the Limits of the Universe*. The Frontiers Collection. Springer Berlin Heidelberg, 2014. URL: https://books.google.co.uk/books?id=g1S-BAAAQBAJ.

[43] F. Karsch, E. Laermann, and A. Peikert. "Quark mass and flavour dependence of the QCD phase transition". In: *Nucl. Phys. B* 605.1-3 (2001), pp. 579–599. DOI: 10.1016/S0550-3213(01)00200-0. arXiv: 0012023 [hep-lat].

[44] G. Shaw and B. Martin. *Particle Physics*. Manchester. Wiley, 2008.

[45] H. Satz. *Lecture Notes in Physics 945 Extreme States of Matter in Strong Interaction Physics*.

[46] G. Martinez. "Advances in Quark Gluon Plasma". In: (2013). arXiv: 1304.1452. URL: http://arxiv.org/abs/1304.1452.

[47] S. Shusu. "Event anisotropy v2 at STAR". PhD thesis. Huazhong Normal University, 2010.

[48] J. Adam et al. "Centrality Dependence of the Charged-Particle Multiplicity Density at Midrapidity in Pb-Pb Collisions at sNN =5.02 TeV". In: *Phys. Rev. Lett.* 116.22 (2016), pp. 1–17. DOI: 10.1103/PhysRevLett.116.222302. arXiv: 1512.06104.

[49] M. . Thoma. "Damping rate of a hard photon in a relativistic plasma". In: *Phys. Rev. D* 51.2 (1995), pp. 2–5. URL: https://doi.org/10.1103/PhysRevD.51.862.

[50] J. Adam et al. "Direct photon production in Pb–Pb collisions at $\sqrt{s_{NN}}$=2.76 TeV". In: *Phys. Lett. Sect. B Nucl. Elem. Part. High-Energy Phys.* 754 (2016), pp. 235–248. DOI: 10.1016/j.physletb.2016.01.020.

[51] The ALICE Collaboration. "Direct photon production in Pb – Pb collisions at $\sqrt{s_{NN}}$ = 2.76 TeV". In: 754 (2016), pp. 235–248. DOI: 10.1016/j.physletb.2016.01.020.

[52] K. Ackermann, N. Adams, and C. Adler. "Elliptic flow in Au+Au collisions at $\sqrt{s_{NN}}$=130 GeV". In: *Phys. Rev. Lett.* 86.3 (2000), pp. 1–5. DOI: 10.1103/PhysRevLett.86.402. arXiv: 0009011 [nucl-ex].

[53] J. Adam et al. “Anisotropic Flow of Charged Particles in Pb-Pb Collisions at $\sqrt{s_{NN}}$ =5.02 TeV”. In: *Phys. Rev. Lett.* 116.13 (2016). DOI: 10.1103/PhysRevLett.116.132302. arXiv: arXiv:1602.01119v2.

[54] B. Abelev et al. “Transverse momentum distribution and nuclear modification factor of charged particles in p+Pb collisions at $\sqrt{s_{NN}}$=5.02 TeV”. In: *Phys. Rev. Lett.* 110.8 (2013). DOI: 10.1103/PhysRevLett.110.082302. arXiv: 1210.4520.

[55] C. Manuel. “Viewpoint The Stopping Power of Hot Nuclear Matter”. In: 97 (2014). DOI: 10.1103/Physics.7.97.

[56] F. M. Liu and S. X. Liu. “Quark-gluon plasma formation time and direct photons from heavy ion collisions”. In: *Phys. Rev. C - Nucl. Phys.* 89.3 (2014), pp. 1–5. DOI: 10.1103/PhysRevC.89.034906.

[57] T. Matsui and H. Satz. “J/ψ suppression by quark-gluon plasma formation”. In: *Phys. Lett. B* 178.4 (1986), pp. 416–422. DOI: 10.1016/0370-2693(86)91404-8.

[58] The ALICE Collaboration. “Differential studies of inclusive J/Ψ and Ψ(2S) production at forward rapidity in Pb–Pb collisions at $\sqrt{s_{NN}}$ = 2.76 TeV”. In: June (2015). arXiv: arXiv:1506.08804v2.

[59] H. Satz. “Quarkonium Binding and Dissociation : The Spectral Analysis of the QGP *”. In: (2006). arXiv: 0609197v1 [arXiv:hep-ph].

[60] K. S. Sateesh. “Experimental signal for diquarks in quark-gluon plasma”. In: 45.3 (1992), pp. 866–868.

[61] ALICE Collaboration. “Constraining hadronization mechanisms with Λ_c^+/D^0 production ratios in Pb-Pb collisions at $\sqrt{s_{\mathrm{NN}}} = 5.02$ TeV”. In: December (2021). arXiv: 2112.08156. URL: http://arxiv.org/abs/2112.08156.

[62] ALICE Collaboration. “Prompt D^0, D^+, and D^{*+} production in Pb-Pb collisions at $\sqrt{s_{\mathrm{NN}}}$ = 5.02 TeV”. In: October (2021). DOI: 10.1007/JHEP01(2022)174. arXiv: 2110.09420. URL: http://arxiv.org/abs/2110.09420%0Ahttp://dx.doi.org/10.1007/JHEP01(2022)174.

[63] P. Kuijer. “The alice experiment at the CERN LHC”. In: *Nucl. Phys. B - Proc. Suppl.* 117.SUPPL. 1 (2003), pp. 62–64. DOI: 10.1016/S0920-5632(03)90488-9.

[64] The ALICE Collaboration. *Letter of Intent for A Large Ion Collider Experiment*. Tech. rep. Geneva: CERN, 1993, CERN–LHCC–93–016, LHCC–I–4. URL: https://cds.cern.ch/record/290825?ln=en.

[65] J. Alme et al. "The ALICE TPC, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 622.1 (2010), pp. 316–367. DOI: 10.1016/j.nima.2010.04.042.

[66] M. Alviggi et al. "The ATLAS experiment at the CERN Large Hadron Collider". In: *J. Instrum.* 3.S (2008), pp. 8003–8008.

[67] G. Petrucciani. "The CMS experiment at the CERN LHC". In: *Search Higgs Boson C.* (2013), pp. 15–58. DOI: 10.1007/978-88-7642-482-3_2.

[68] The ALICE Collaboration. "The alice experiment at the CERN LHC". In: *J. Instrum.* 3.S08002 (2008). DOI: 10.1088/1748-0221/3/08/s08002.

[69] M. Suljic. "Study of Monolithic Active Pixel Sensors for the Upgrade of the ALICE Inner Tracking System". PhD thesis. Universita Degli Studi Di Trieste, 2017.

[70] G. Contin. "Performance of the present ALICE Inner Tracking System and studies for the upgrade". In: *J. Instrum.* 7.6 (2012). DOI: 10.1088/1748-0221/7/06/C06007.

[71] V. Altini. "Trigger performance of the ALICE Silicon Pixel Detector". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 650.1 (2011), pp. 30–32. DOI: 10.1016/j.nima.2010.11.167. URL: http://dx.doi.org/10.1016/j.nima.2010.11.167.

[72] E. Botta on behalf of the ALICE Collaboration. "ALICE ITS: Operational experience, performance and lessons learned". In: *Proc. Sci.* 373 (2019). DOI: 10.22323/1.373.0002.

[73] H. Appelshauser et al. "Space-charge distortion measurements and their calibration in the ALICE TPC". In: *Alice Note 2016* (2016), pp. 1–73. URL: https://indico.cern.ch/event/587756/.../TPC_distortions_summary_Nov_11.pdf.

[74] The ALICE Collaboration. "Technical Design Report for the Upgrade of the ALICE Time Projection Chamber". In: (2014). URL: https://cds.cern.ch/record/1622286/files/ALICE-TDR-016.pdf.

[75] The ALICE Collaboration. "ALICE MFT TDR". In: *Cern. ALICE-TDR-018* (2015). URL: https://cds.cern.ch/record/1981898.

[76] M. Mager on behalf of the ALICE Collaboration. "ALPIDE, the Monolithic Active Pixel Sensor for the ALICE ITS upgrade". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 824.2016 (2016), pp. 434–438. DOI: 10.1016/j.nima.2015.09.057. URL: http://dx.doi.org/10.1016/j.nima.2015.09.057.

[77] X. N. Wang and M. Gyulassy. "Hijing: A Monte Carlo model for multiple jet production in pp, pA, and AA collisions". In: *Phys. Rev. D* 44.11 (1991), pp. 3501–3516. DOI: 10.1103/PhysRevD.44.3501.

[78] T. Sjöstrand, S. Mrenna, and P. Skands. "PYTHIA 6.4 physics and manual". In: *J. High Energy Phys.* 2006.5 (2006). DOI: 10.1088/1126-6708/2006/05/026. arXiv: 0603175 [hep-ph].

[79] Y. Oh et al. "Ratios of heavy baryons to heavy mesons in relativistic nucleus-nucleus collisions". In: *Phys. Rev. C - Nucl. Phys.* 79.4 (2009), pp. 1–10. DOI: 10.1103/PhysRevC.79.044905.

[80] M. He, R. J. Fries, and R. Rapp. "Ds meson as a quantitative probe of diffusion and hadronization in nuclear collisions". In: *Phys. Rev. Lett.* 110.11 (2013), pp. 1–5. DOI: 10.1103/PhysRevLett.110.112301. arXiv: 1204.4442.

[81] H. Hillemanns et al. "Radiation hardness and detector performance of new 180nm CMOS MAPS prototype test structures developed for the upgrade of the ALICE Inner Tracking System". In: *IEEE Nucl. Sci. Symp. Conf. Rec.* (2013). DOI: 10.1109/NSSMIC.2013.6829475.

[82] F. Reidt. "Studies for the ALICE Inner Tracking System Upgrade". In: *Nucl. Phys. A* 956 (2016), pp. 866–869. arXiv: 1211.4494.

[83] A. Szczepankiewicz on behalf of the ALICE Collaboration. "Readout of the upgraded ALICE-ITS". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 824 (2016), pp. 465–469. DOI: 10.1016/j.nima.2015.10.056. URL: http://dx.doi.org/10.1016/j.nima.2015.10.056.

[84] IBS Precision Engineering. *https://www.ibspe.com/*. URL: https://www.ibspe.com/ (visited on 03/14/2022).

[85] IBS Precision Engineering. *ALICIA User manual*. 2016.

[86] A. Junique. "The engineering, production and quality assurance of the inner barrel staves for the upgrade of the ALICE inner tracking system". In: *Proc. Sci.* 370 (2019), pp. 2–6. DOI: 10.22323/1.370.0081.

[87] L. de Cilladi. *Characterization of Monolithic Pixel Detectors for the ALICE Inner Tracking System*. 2019. URL: http://www.infn.it/thesis/PDF/getfile.php?filename=14402-De%20Cilladi-magistrale.pdf.

[88] Glenn Roberts Jr. *How to Get a Particle Detector on a Plane*. 2019. URL: https://newscenter.lbl.gov/2019/09/19/how-to-get-a-particle-detector-on-a-plane/ (visited on 04/14/2022).

[89] F. James and M. Roos. "Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations". In: *Comput. Phys. Commun.* 10 (1975), pp. 343–367. DOI: 10.1016/0010-4655(75)90039-9.

[90] G. De Robertis et al. "A modular system for acquisition, interface and control (MOSAIC) of detectors and their related electronics for high energy physics experiment". In: *EPJ Web Conf.* 174 (2018), pp. 2016–2019. DOI: 10.1051/epjconf/201817407002.

[91] F. Reidt, J. Klein, and M. A. Imhoff. *Photograph: ALICE ITS Outer Barrel (OB) installation*. 2021. URL: http://cds.cern.ch/record/2759554.

[92] M. Lupi. "Design , development , and experimental assessment of a highly-reliable , radiation-tolerant readout system for the upgrade of the ALICE inner tracker". PhD thesis. 2020. URL: https://cds.cern.ch/record/2742653/files/CERN-THESIS-2020-150.pdf.

[93] J. J. Wang et al. "Total ionizing dose effects on flash-based field programmable gate array". In: *IEEE Trans. Nucl. Sci.* 51.6 II (2004), pp. 3759–3766. DOI: 10.1109/TNS.2004.839255.

[94] P. Moreira et al. "The GBT Project". In: (2009). DOI: 10.5170/CERN-2009-006.342. URL: https://cds.cern.ch/record/1235836.

[95] A. Velure. "Integration, commissioning and first experience of ALICE ITS control and readout electronics". In: *Proc. Sci.* 370.September (2019). DOI: 10.22323/1.370.0113.

[96] ALICE ITS ALPIDE Development Team. *ALPIDE Operations Manual*. 2016.

[97] ALICE ITS Upgrade. *ITS Readout Fact Sheet*. 2020.

[98] The ALICE Collaboration. "Technical Design Report for the Upgrade of the Online-Offline Computing System". In: CERN-LHCC-.ALICE-TDR-019 (2015). URL: https://cds.cern.ch/record/2011297/?ln=fr.

[99] P. Martinengo. "The new Inner Tracking System of the ALICE experiment". In: *Nucl. Phys. A* 967 (2017), pp. 900–903. DOI: 10.1016/j.nuclphysa.2017.05.069. URL: http://dx.doi.org/10.1016/j.nuclphysa.2017.05.069.

[100] E. W. Weisstein. *Point-Line Distance–3-Dimensional*. URL: https://mathworld.wolfram.com/Point-LineDistance3-Dimensional.html (visited on 03/30/2022).

[101] M. Šuljić. "ALPIDE: The Monolithic Active Pixel Sensor for the ALICE ITS upgrade". In: *J. Instrum.* 11.11 (2016). DOI: 10.1088/1748-0221/11/11/C11025.

[102] ALICE ITS Collaboration. "First demonstration of in-beam performance of bent Monolithic Active Pixel Sensors". In: *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.* 1028.August 2021 (2022), p. 166280. DOI: 10.1016/j.nima.2021.166280. arXiv: 2105.13000. URL: https://doi.org/10.1016/j.nima.2021.166280.