


ALICE Forward Multiplicity Detector

From Design to Installation



Christian Holm Christensen

December 2007

Niels Bohr Institute
University of Copenhagen

Supervisor:

Jens Jørgen Gaardhøje
Niels Bohr Institute
Blegdamsvej 17
DK-2100 København Ø

Contact information:

Home: Sankt Hansgade 23, 1th
DK-2200 København N

(+45) 35 35 96 91

(+45) 24 61 85 91

Email: cholm@nbi.dk

Institute: Niels Bohr Institute
Blegdamsvej 17
DK-2100 København Ø

(+45) 35 32 54 04

Homepage: <http://www.nbi.dk/~cholm>

The cover shows the ALICE detector, with a cut-away to show the inside of the central barrel detectors and the forward detectors.

This document is set by L^AT_EX, using Computer Modern Fonts.

Figures are created in XFig <http://www.xfig.org> and ROOT <http://root.cern.ch>.

Preface

In the ALICE Technical Proposal [1] a forward detector based on Micro-Channel Plate (MCP) sensors was proposed to provide charged particle multiplicity measurements and trigger information. However, quoting the ALICE Forward Detectors Technical Design Report (TDR) [2]

The MCP option would have been a novel and elegant way to build a forward detector system, but it was realized that it would have required substantial funds and a major R&D effort and that, the desired functionality could be achieved with existing and proved technologies based on Cherenkov radiators, scintillators, and Si-strip detectors. This led to the division into the T0, the V0 and the FMD, for convenience named the Forward Detectors (FWD). These systems provide different functionalities, but also to some extent functional overlap and complementarity, which is considered an advantage for ALICE.

This thesis summarises my work over the past 4 years on the ALICE Forward Multiplicity Detector, starting as a technical editor of the FWD TDR, and ending with the installation of 4/5 of the FMD detector in the ALICE experimental hall. During these past 4 years, I have been in contact with virtually all parts of the FMD, ranging from hardware design and testing to implementation of analysis software and simulated data analysis.

From January 2005 up to the summer of that same year, the premises, requirements, and design of the FMD was formulated as part of the preparation of the FWD Technical Design Report. For me, it provided an opportunity to acquaint my self with the ALICE experiment and to provide input into the design process through simulations of the detector and software expertise.

The physical implementation of the detector started shortly after the acceptance of the TDR. My next task was to understand the basic front-end electronic components adopted by the FMD from the ALICE Time Projection Chamber (TPC) project, with the ultimate goal of customising that front-end electronics to the needs of the FMD. In that context I collaborated closely with electronics engineers and technicians to formulate the requirements of the FMD front-end electronics and in particular to implement the firmware for the FMD digitiser cards board controller chip.

Once we had a more or less complete system, the focus shifted towards testing and control of the system. To that end, I developed a large base of control, monitoring, and analysis software to perform the necessary tests. Much of this software has since then evolved into the crucial parts of the FMD control, monitoring, and analysis software used in the final installation, including among other things the on-line calibration programs, configuration database, and the front-end server.

The first sub-detector, FMD3 was assembled in early 2007, and install April of that same year. Before, during, and after the installation Kristján Gulbrandsen, Gaël Renault, and I did extensive tests of the sub-detector at CERN, and took an active part in the final installation. This work culminated in the successful commissioning of the FMD.

Parallel to the hardware related parts, I have also developed off-line simulation, geometric description, reconstruction, and basic analysis tools for the FMD.

Finally, a large portion of my time has been spent on guiding and helping the two master students Carsten Søgaaard and Hans Dalsgaard in their work on test-beam results and analysis

of simulated data respectively. Both students handed in their master thesis in the summer 2007 and received high marks for their work.

Acknowledgements

All mankind is of one author, and is one volume; when one man dies, one chapter is not torn out of the book, but translated into a better language; and every chapter must be so translated. . . . As therefore the bell that rings to a sermon, calls not upon the preacher only, but upon the congregation to come: so this bell calls us all: but how much more me, who am brought so near the door by this sickness. . . . No man is an island, entire of itself. . . any man's death diminishes me, because I am involved in mankind; and therefore never send to know for whom the bell tolls; it tolls for thee.

John Donne (1572-1631) in "Devotions Upon Emergent Occasions", Meditation XVII.

First off all many thanks to the High Energy Heavy Ion (HEHI) group at the Niels Bohr Institute (NBI), and in particular my supervisor Jens Jørgen Gaardhøje, for giving me the opportunity to do this Ph.D. project. Also from the HEHI, I would like to thank Hans Bøggild for many inspiring discussions, Ian Gardner Bearden for the proper use of the 'motivator', and Kristján Herlache Gulbrandsen for being my brother-in-arms in what sometimes seemed like a war, to get the FMD to work. Special thanks goes to master students Carsten Søgaaard and Hans Hjersing Dalsgaard — it was a pleasure to help and guide you as best I could.

On the technical side, I would like to thank the entire Technical Infrastructure group at NBI under the fearless leadership of Børge Svane Nielsen. Special thanks for countless afternoons of brain-storming over problems, large or small, Henrik Bertelsen, for superb support and craftsmanship Jimmy Cali Hansen, Thomas Dyhre, and Bjørn Rasmussen. Last but not least, Jørn Westergaard for his often ingenious and simple solutions to often complicated problems.

Many thanks to the various groups at CERN and other places. Especially thanks to Luciano Musa and his group for help and guidance in understanding the TPC front-end electronics.

Two persons deserve some very special thanks: Tiago Perez and Daniel Kirschner, both of University of Gießen — they taught me everything I know about firmware programming, and without their help I sincerely doubt the FMD project would have progressed without significant delays.

Finally, I would to thank my family, especially Christina Back, for all their love and support.

Christian Holm Christensen
Copenhagen, December 2007

Contents

I	Introductory Material	xv
1	Introduction	1
1.1	Quarks and Gluons — The Strong Interaction	2
1.2	The QCD Phase Transition and Chiral Symmetry	3
1.3	Relativistic Heavy-ion Collisions	5
1.4	Collider Physics	7
1.4.1	Impact Parameter, Participants, and Spectators	7
1.4.2	Transparency and stopping	8
1.4.3	Chemical and Thermal Freeze-out	11
1.5	Signatures for QGP	11
1.5.1	Kinematics	11
1.5.2	Dynamics	12
1.5.3	Chiral Symmetry Restoration — Disoriented Chiral Condensates	13
2	Observables Relevant to the ALICE FMD	15
2.1	Impact Parameter and Centrality	15
2.2	Charged Particle Density ($dM_{ch}/d\eta$)	16
2.2.1	Overview of previous $dM_{ch}/d\eta$ measurements	17
2.2.2	Lessons to be learned from the $dM_{ch}/d\eta$ measurements	21
2.3	Azimuthal Anisotropy (v_2)	22
2.3.1	Fourier decomposition of the azimuthal particle distribution	22
2.3.2	Overview of previous v_2 measurements	24
2.3.3	Lessons to be learned from the v_2 measurements	27
3	The LHC and the ALICE Experiment	29
3.1	The Large Hadron Collider	29
3.2	Overview of ALICE	30
3.2.1	Barrel Detectors	31
3.2.2	Forward Detectors	33
3.2.3	Muon Spectrometer	34
II	The ALICE Forward Multiplicity Detector	37
4	Overview and Design Considerations	39
4.1	Overview	39
4.2	Design Considerations	39
5	Sensors and Electronics	41
5.1	The FMD Sensors	41
5.1.1	Semi-conductor detectors	42
5.1.2	Energy loss in Silicon	46

5.2	The Hybrid cards and VA1 pre-amplifier	50
5.2.1	The VA1 ₃ chip	52
5.3	The Digitiser Cards	56
5.3.1	The ALTRO Chip	61
5.3.2	The Board Controller	65
5.3.3	ALTRO bus protocol	70
5.4	The Read-out Controller Unit	72
5.4.1	DCS Daughter Card	74
6	Data Acquisition and Control	77
6.1	Data Acquisition	77
6.1.1	Event Format	79
6.2	Detector Control System	80
6.2.1	Architecture	80
6.2.2	Low and bias voltage system	81
6.2.3	Front-end Electronics Chain	82
6.2.4	Configuration of the Front-end	88
7	Test beam results	91
7.1	Initial Tests of the Modules	91
7.2	Full Test of the Modules and Electronics, and Efficiency	93
7.2.1	Setup	93
7.2.2	Overall performance	94
7.2.3	Signal-to-Noise	95
7.2.4	Efficiency	95
8	An Analysis Chain for the FMD Data	101
8.1	Geometry	102
8.2	Simulations of events	103
8.3	Digitisation of Simulated Data	104
8.4	Reading Raw Data	106
8.5	Reconstructing Bare Multiplicity	107
8.6	Hit Sharing Corrections and Particle Cuts	110
8.7	Alternative for Low Multiplicity Events	112
8.8	Acceptance and Secondary Particle Corrections	114
8.9	Charged Particle Multiplicity Density	115
8.10	Event Plane Angle and Azimuthal Anisotropy	115
8.10.1	Event Plane	116
8.10.2	Azimuthal Anisotropy — v_2	119
9	Conclusion	121
A	Variables	131
A.1	Physical Variables	131
B	Images of the ALICE Detectors	133
C	Additional Figures for Chapter 2	135
D	Additional Figures for Chapter 5	137
E	Additional Figures for Chapter 8	141

F Miscellaneous Photos

List of Figures

1.1	ALICE	1
1.2	The quark–quark potential	3
1.3	The strong coupling constant α_s as a function of the momentum transfer Q	4
1.4	Lattice QCD calculations of $\langle\psi\bar{\psi}\rangle$, ε and P as a function of T	4
1.5	$\langle\bar{\psi}\psi\rangle$ as a function of density ρ and temperature T	5
1.6	The phase diagram of nuclear matter	6
1.7	(a) Impact parameter in heavy ion collisions. (b) Spectators and participants.	7
1.8	Number of participants and produced charged particles versus b	8
1.9	Stopping in heavy ion collisions.	9
1.10	Transparency in heavy ion collisions.	9
1.11	Ratios of like-particles	10
1.12	Net protons and average rapidity loss	10
1.13	The disappearance of the away–side jet in central Au–Au	12
1.14	The potential V between particles in a linear σ model	14
2.1	Centrality cuts	16
2.2	Charged particle densities from RHIC	17
2.3	$dM_{\text{ch}}/d\eta$ as a function of $\sqrt{s_{NN}}$	18
2.4	Gluon density, and saturation in CGC	19
2.5	$dM_{\text{ch}}/d\eta$ at different beam energies and centralities.	20
2.6	The charged particle rapidity as a function of the scaled psuedo–rapidity $\eta' = \eta - y_{\text{beam}}$	21
2.7	CGC prediction on $dM_{\text{ch}}/d\eta$ at different centralities.	22
2.8	Event P_E plane spanned by the colliding nuclei.	23
2.9	v_2/ε as a function overlap normalised $dM_{\text{ch}}/d\eta$	25
2.10	Pseudo–rapidity dependent v_2	25
2.11	Scaling of v_2	26
3.1	The CERN accelerator complex	29
3.2	ALICE	30
3.3	Section through the magnet and barrel detectors.	31
3.4	The forward detectors and the ITS.	34
3.5	The muon spectrometer.	35
4.1	The forward detectors and the ITS	39
4.2	Total charged particle multiplicity coverage of ALICE	40
5.1	The ALICE FMD electronics chain.	41
5.2	Inner and outer sensor types cut from 6” silicon wafers.	42
5.3	Arrangement of inner (left) and outer (right) sensor into rings.	42
5.4	Valence and Conduction band gap in insulators, semi–conductors, and metals.	43
5.5	Bindings in the silicon crystal	44
5.6	Cross–section through a silicon sensor.	45

5.7	Capacitance and leakage as a function of bias voltage	46
5.8	Sensor thickness and stability	47
5.9	Bethe–Bloch equation evaluated for μ on copper	48
5.10	Energy loss straggling in thin absorbers	49
5.11	Mean and most probable energy loss in Si	51
5.12	Schematic drawing of an inner–type hybrid	51
5.13	Architecture of the VA1 ₃ chip.	52
5.14	Noise (in electrons) of the VA1 ₃ chip	53
5.15	Read–out of all 128 channel in a VA1 ₃ chip.	55
5.16	Read–out of all 128 channel in a VA1 ₃ chip.	56
5.17	Setting up for gain calibration of the 3 rd channel of a VA1 ₃	56
5.18	Testing of a single Sensor–Hybrid compound.	57
5.19	The two types of FMDD	58
5.20	Schematic diagram of FMDD BC cards. Note, that both inner and outer cards have, in total, 20 VA1 ₃ ’s connected.	59
5.21	Schematic architecture of the ALTRO chip	62
5.22	One ALTRO channel sampling data from 128 strips	63
5.23	ALTRO channel event format.	64
5.24	Timing of ALTRO channel readout.	65
5.25	Generation of \overline{dstb}	65
5.26	Schematic overview of the FMDD BC architecture	66
5.27	Timing of triggers	68
5.28	State changes in the trigger handling module	68
5.29	Write instruction sequence on the ALTRO bus	71
5.30	Read instruction sequence on the ALTRO bus	72
5.31	Schematic architecture of the RCU firmware with peripheral devices indicated. . .	73
5.32	Example of a program for the RCU instruction interpreter. On the left is the semi–human–readable code, and on the right, the output of the compiler that can be interpreted by the RCU PU.	75
6.1	The three tier DAQ architecture.	78
6.2	The global event structure.	79
6.3	The FMD DCS architecture.	81
6.4	The FMD DCS software architecture.	82
6.5	The FMD top–level DCS finite state machine.	86
6.6	Steps executed in going from the <i>Off</i> state to the <i>Ready</i> state.	88
6.7	Schema of the DCS front–end electronics configuration database.	89
7.1	Extraction of parasitic beam from ASTRID.	91
7.2	Spectrum of calibrated single strip signal	92
7.3	Hit sharing from particle traversing more than one strip	92
7.4	Sharing of signal between neighbouring strips	93
7.5	Arrangment of the silicon modules in specially manufactured box.	94
7.6	Measured signal–to–noise ratio of inner (left) and outer (right) silicon sensors using 630 MeV e^- from ASTRID	95
7.7	Count statistics, signal, noise, and signal–to–noise ratio of all strips in a single inner module	96
7.8	Count statistics, signal, noise, and signal–to–noise ratio of all strips in a single outer module	97
7.9	Efficiency of strips in a single module.	99
7.10	Efficiency of strips in a single module.	99

8.1	Flow of data through ALIROOT	102
8.2	Energy loss in the FMD silicon sensor from a GEANT 3 simulation.	104
8.3	Conversion of energy deposition into ADC counts	105
8.4	Illustration of ADC value generation in case of over sampling rate 1, 2, 3, and 4.	106
8.5	The process of reconstruction the bare multiplicity in a strip from the raw ADC counts.	107
8.6	Particles traverse the silicon sensors at some angle, meaning that it will see more or less material depending on strip. Hence the signal Δ' need to be path-length corrected by a factor of $\cos \vartheta$	108
8.7	The distribution of strip multiplicity before and after the sharing correction, as well as the merged signals and the signals removed.	110
8.8	Analysing the single strip multiplicity for single and double particle hits	112
8.9	Reconstruction of the charged particle density.	113
8.10	Comparison of methods to reproduce the bare charged particle multiplicity	114
8.11	Relative strength sources of secondary particles seen by the FMD	115
8.12	The primary charged particle density as reconstructed by the FMD, and compared to the input HIJING model data	116
8.13	Difference of reconstructed event plane Ψ_2 to true event plane Ψ_R	118
8.14	Reconstructed v_2 as a function of η	120
B.1	ALICE detectors.	133
B.2	The ALICE event display, showing a central Pb–Pb. Top panel shows the hits generated in the simulation, and the bottom panel shows the particle tracks through the detector simulation.	134
C.1	Fragmentation region in Au–Au at $\sqrt{s_{NN}} = 19.4$ GeV, $\sqrt{s_{NN}} = 130$ GeV, and $\sqrt{s_{NN}} = 200$ GeV at various centralities, as measured by the PHOBOS experiment. Adapted from [3, 4].	135
D.1	Characteristics of all inner sensors, as measured by Hamamatsu. Capacitance and leakage current as a function of bias voltage for the inner sensors. In the 2 nd panel from the top, on the right, one can clearly see the break–down of the semi–conductor at a bias voltage of around 300 V.	137
D.2	Characteristics of all outer sensors, as measured by Hamamatsu. Capacitance and leakage current as a function of bias voltage for the outer sensors. In the 2 nd panel from the top, on the right, one can clearly see the break–down of the semi–conductor at a bias voltage of around 300 V.	138
E.1	Simulated energy loss in the FMD with no δ ray production.	141
F.1	An inner type sensor	143
F.2	The FMD front–end electronics	144
F.3	Test–beam setup at ASTRID	144
F.4	FMD2 and 3 installed in ALICE	145
F.5	Some members of the FMD underneath the installed FMD3 sub–detector	145

List of Tables

1.1	Relativistic heavy-ion accelerators.	6
3.1	Kinetic energy at different stages in the beam acceleration.	30
4.1	Overview of parameters the FMD sub-detectors.	40
5.1	Various parameters of the VA1 ₃ and silicon strips used in the calculation of the ENC.	54
5.2	Lines in the two bus cables. The width is in number of bits.	60
5.3	Low-voltage power supply	61
5.4	FMDD specific registers in the BC	67
5.5	Commands implemented in the FMDD BC. Note, that FMD specific commands is executed by writing a value to the register <code>fmdd_cmd</code>	67
5.6	Monitor ADC related registers.	70
5.7	Meaning of upper 20bits of <code>bd</code> during register input/output and command execution.	71
6.1	Commands and services provided by the ICL. The <code><icl_name></code> is typically something like <code>ztt_dimfed_server</code> , but can be configured.	84
D.1	Silicon sensor design parameters.	139

Part I

Introductory Material

Chapter 1

Introduction




Figure 1.1: Section through ALICE. Arrows indicate positions of the various detectors.

When the *Large Hadron Collider* (LHC) turns on early 2008, a new era of heavy-ion physics is begun. Heavy-ions, such as lead, are collided at velocities very close to the speed of light, in an effort to create a state of matter known as the *Quark Gluon Plasma* (QGP). In the QGP, the smallest known particles, *quarks*, are expected to become free and can be studied by experiments.

A *Large Ion Collider Experiment* (ALICE, see Figure 1.1) will study the results of the heavy-ion collisions produced by the LHC. The aim is to study the *strong* force — the force that binds the quarks in the nucleons — by creating a hot and dense state of quarks, in which a transition to the QGP is expected to be found.

ALICE consists of many detectors, as indicated on Figure 1.1. One of these detectors is the *Forward Multiplicity Detector* (FMD) built at the Niels Bohr Institute, and is the subject of this thesis.

This chapter gives a brief overview of the kind of beasts in the particle zoo we are dealing with in heavy ion physics, and what governs their behaviour. The idea of a phase transition from normal nuclear matter to a plasma of free quarks and gluons is presented. Presented is also an

overview of how we use relativistic heavy ion collisions to search for such a phase transition.

In Chapter 2 a few such signatures relevant to the ALICE FMD are described in more detail. From then on, the text will focus on ALICE in general (Chapter 3) and the FMD specifically. Chapter 4 gives an overview of the ALICE FMD, and Chapter 5 describes the hardware in details. Chapter 6 describes the data acquisition system and the detector control system used to operate ALICE FMD. In Chapter 7 the test beam results are presented with emphasis on the signal-to-noise ratio, the efficiency, and the general performance of the system. Chapter 8 shows the use of the FMD in the ALICE data analysis.

1.1 Quarks and Gluons — The Strong Interaction

The current framework of modern physics is called the *Standard Model*¹. In this theory, matter and interactions are described as particles. How particles interact with each other depends on the *charge* of the particles involved. A charge for each kind of interaction is assigned to each particle. For the electromagnetic interaction, it is the familiar electric charge, and for the strong interactions it is the *colour* charge.

Quantum *ChromoDynamics*, or QCD for short, is the theory that describes how *coloured* objects, such as the quarks interact. In QCD, the coupling between the coloured objects are mediated by the gauge boson called the *gluon*. Gluons themselves are coloured, which means that gluons can interact, or couple to one another. This property of QCD, makes it radically different from the rest of the theories describing interactions between particles. In particular, the interaction of gluons gives rise to what is known as *confinement*.

As quarks are moved farther and farther apart, it costs more and more energy (see Figure 1.2(a)), since the force² is non-vanishing. The textbook analogy is that of a rubber band — stretching a rubber band requires increasingly more energy. Just as a rubber band will eventually break if stretched too far, so will the interaction of two coloured objects, eventually ‘break’. What happens is, that at a certain distance enough energy will be present in the colour field to form a new quark–antiquark pair i.e., a meson.

This is what is meant by confinement. The quarks in a baryon or meson may never escape the other quarks. This has the unfortunate effect that we can not hope to study the properties of a free i.e., non-interacting, isolated quark [8]. The fact that no coloured particles have been observed in nature, supports this view. A net charged coloured object would allow us to study QCD in more details, however it is as if the rubber band breaks because the colour charge can not exist on localised objects.

This is in contrast to for example *Quantum ElectroDynamics* (QED), that concerns interaction of electric charged particles. In QED, the field boson, the photon (γ) is not itself electric charged, and γ 's therefor can not couple to one another.

The forces of QED (and more generally *Quantum ElectroWeakDynamics*, or QWD), and the much familiar force of gravity, lessens with the distance between the interacting particles. For example, it get easier and easier i.e., does not cost as much energy, to remove a rock from the gravitational field of another (larger) body e.g., the Earth.

In fact, the rubber band analogy can be extended even further. If the two ends of a rubber band are moved close to one another, the tension in the band disappears, and the two ends do not feel each other. For quarks this is also true. If two quarks are close to one another, the interaction between them is small — if moved really close, they essentially do not interact at all — in a sense

¹*Model* is an unfortunate term, in that it suggests that the theory of the Standard Model is a conceptual idea of what nature is like without any real empirical evidence. However, nothing could be further from the truth. The standard model has successfully described many phenomena observed in nature to a great precision. One example is the measurements of electrons dipole magnetic moment (the so-called $g/2$ experiments) which gives values for the electromagnetic coupling α that agrees with the standard model predictions a precision of $< 0.0000001\%$ [5, 6].

² $F = -\frac{dV}{dr} = -k \approx 1 \text{ GeV/fm}$




Figure 1.2: (a) The heavy quark–quark potential, from Lattice QCD, roughly behaving as $a/r + br$. K is the string tension, a theoretical parameter. From [7]. For comparison, (b) show an $1/r$ potential.

they become free, which is termed *asymptotic freedom* [8, 9]. This is shown in Figure 1.3. The strength of the strong coupling α_s is plotted as a function of the momentum transfer³.

1.2 The QCD Phase Transition and Chiral Symmetry

As hadronic matter is heated and compressed, the interactions between the quarks become stronger, and the typical momentum transfers become very large on very short distances. At a critical temperature T_c the quarks are expected to become deconfined, due to the asymptotic freedom of QCD. In such a state of very dense and very hot quarks, it is expected that quarks and gluons essentially do not interact strongly anymore.

Lattice QCD (LQCD) calculations show a sharp increase in the energy density ε , and a more modest growth in pressure P , motivating the idea of a phase transition from hadronic, or confined, matter to a deconfined state. See also Figure 1.4(a). Note however, that LQCD calculations have until recently only been possible at zero baryon chemical potential $\mu_b = 0$ i.e., an equal amount of baryonic and anti-baryonic matter is present. Recently Fedor and Katz [11] have been able to extend LQCD into $\mu_B \neq 0$.

In QCD, the symmetry of *chirality* is broken both explicitly and spontaneously. Chirality is the *handedness* of a particle, or the sign of the projection of the spin vector onto the momentum vector of the particle. If the projection of the spin is aligned to the momentum, the particle is said to be right-handed, while if the projection of the spin is anti-aligned, the particle is said to be left-handed.

In QCD, chiral symmetry exists as an approximate explicitly broken symmetry in the limit of massless, or noninteracting quarks (see also Section 1.5.3). The explicit breaking occurs because the quark masses are not exactly zero. However, when the quarks do interact and become confined to the nucleon or meson, the chiral symmetry is spontaneously broken.

The fact that chiral symmetry is broken means, that the π mesons have very low, but distinct masses for the π^0 and π^\pm . However, if the nuclear matter is heated above a critical temperature T_c and/or compressed above a certain density ρ_c , chiral symmetry is expected to be restored (see Figure 1.5). The effects of restoring the symmetry are that the π mesons have degenerate mass.

³Note, large momentum transfer corresponds to short range due to the Heisenberg uncertainty relations $\Delta x \Delta t > \hbar$.




Figure 1.3: The strong coupling constant α_s as a function of the momentum transfer Q [10].




Figure 1.4: (a) Lattice QCD calculations of energy density ε and pressure P as a function of temperature T , with only two flavours of light quarks (u and d quarks). The critical temperature T_c is here ≈ 150 MeV. Adapted from [12]. (b) Lattice QCD calculation of the quark condensate $\langle \bar{\psi}\psi \rangle$ for various number of quark flavours (N_f) and masses (m_q). Adapted from [13].

The measure of chirality is the expectation value of the *quark-condensate*⁴ $\langle\bar{\psi}\psi\rangle$, or the content of quarks and antiquarks in the vacuum of the nucleon. For low temperatures and low densities $\langle\bar{\psi}\psi\rangle$ is high, meaning that the nucleon vacuum is highly filled with virtual quarks that interact with the regular (or valence) quarks of the nucleon, but it is expected to drop above the critical temperature T_c [14] (see also Figure 1.4(b)). Section 1.5.3 deals with the possible chiral symmetry restoration in relativistic heavy ion collisions.




Figure 1.5: $\langle\bar{\psi}\psi\rangle$ as a function of density ρ and temperature T .

The onset of deconfinement and chiral symmetry restoration is expected to occur at roughly the same temperature [15].

If such a phase transition from chiral asymmetric and confined matter does exist, and it can be measured in experiments, it could tell us a lot about the relatively unknown QCD. Therefore hadron colliders are of great interest in these years. The idea is to squeeze the nucleons or nuclei close together to form a hot fireball where the quarks are essentially free to move in a chiral symmetric soup of quarks and gluons — a *Quark-Gluon Plasma*, or QGP for short.

1.3 Relativistic Heavy-ion Collisions

By colliding heavy ions ($A \gg 1$, where A is the atomic weight of the ion) one hopes to achieve energy densities ε , temperatures T , and matter densities ρ above the critical values for the formation of the QGP. Such experiments have been carried out at the Alternating Gradient Synchrotron (AGS) at Brookhaven National Laboratories (BNL), Super Proton Synchrotron (SPS) at European Organisation for Nuclear Research (CERN), and are currently being carried out at the Relativistic Heavy Ion Collider (RHIC) at BNL. From 2008, will the Large Hadron Collider (LHC) at CERN collide lead ions at $\sqrt{s_{NN}} = 5500$ GeV. Table 1.1 tabulates various characteristics of the four accelerators. Figure 1.6 shows the trajectories of matter in phase diagram of nuclear matter at the various accelerators .

$\sqrt{s_{NN}}$ is the amount of energy per nucleon available in the centre of mass system of the collision (see also Appendix A). In a collider such as RHIC and LHC it is just the sum of the energy of the two colliding beams, divided by the number of the particles in the beam.

⁴The vacuum expectation value of a quark-antiquark system. ψ is the quark field, and $\bar{\psi}$ is the anti-quark field.

Accelerator	Location	Length	$\sqrt{s_{NN}}$	Type
AGS	BNL	800 m	5.2 GeV	fixed-target
SPS	CERN	7 km	17 GeV	fixed-target
RHIC	BNL	4.6 km	200 GeV	collider
LHC	CERN	27 km	5.5 TeV	collider

Table 1.1: Relativistic heavy-ion accelerators.

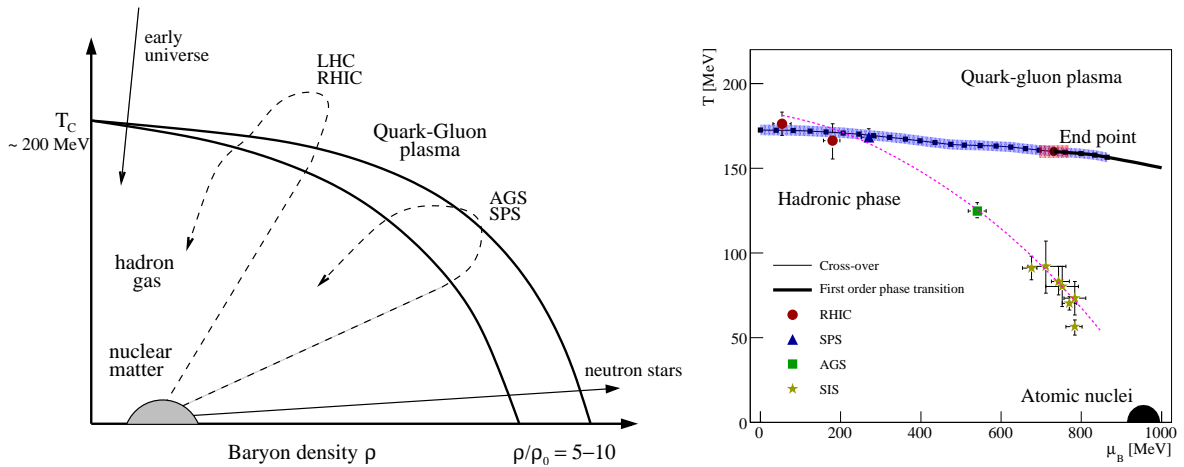


Figure 1.6: (a) The postulated phase diagram of nuclear matter. The arrows indicate how different colliders and natural phenomenon probes the QGP phase. (b) Lattice QCD calculations of the baryon chemical potential μ_B and temperature T at the phase transition. The black dot indicates the end point for the 1st-order phase transition. The dashed line represents a constant baryon density of 0.12 fm^{-3} . Adapted from [11, 16].

At RHIC there are 4 experiments dedicated to heavy ion physics: BRAHMS⁵, PHOBOS, STAR, and PHENIX. At the LHC the only dedicated heavy ion experiment is ALICE (see Chapter 3.2). The remaining 3 LHC experiments, ATLAS, CMS, and LHCb are dedicated to particle physics. However, there are emerging heavy ion subgroups within the ATLAS and CMS collaborations.

1.4 Collider Physics

In a collider experiment, two beams of charged particles are accelerated in opposite directions, and brought to collide at fixed locations. Most colliders are ring shaped with one or two beam pipes, but some colliders have other shapes⁶. Often, the beams consists of like particles, so at the collision point there is no net momentum in the laboratory frame. This means that all of the energy of the colliding beams can in principle be used for particle production, etc.. This should be compared to *fixed target* experiments, where a single beam is accelerated and shot at a stationary target. In that case, there is a net momentum in the laboratory frame, which can not be used in processes.

When the beams consists of nucleons, such as n or p one speaks of N+N collisions, and if the the beams consists of nuclei, like lead (Pb) or gold (Au) one speaks of A+A collisions. In high energy particle physics collisions of e^- and e^+ have been the system of choice for the past 10 or so years at LEP at CERN due to its relative cleanness, since the e^- and e^+ are believed to be atomic — that is, has no constituents. However, to probe new regions of QCD, particle physics will start doing $p + p$ collisions at LHC in the year 2008.

Collisions are characterised by the amount of energy available in the centre of mass frame of the collision, denoted by \sqrt{s} . For symmetric collisions in a collider $\sqrt{s} = \sqrt{s_1^2 + s_2^2}$ (the root of the sum of the square kinetic energy of the two colliding particles). In A+A one often uses centre of mass energy per nucleon $\sqrt{s_{NN}} \equiv \sqrt{s}/2N$ where N is the number of nucleons in each of the colliding nuclei.

1.4.1 Impact Parameter, Participants, and Spectators

When colliding particles of non-vanishing size, like nuclei, the two colliding particles can have varying size of overlap zones. The area of the overlap is determined by the *impact parameter* \mathbf{b} . \mathbf{b} is the vector that connects the two centres of the colliding nuclei, and its length b is measured in fm, as shown in Figure 1.7.




Figure 1.7: (a) Impact parameter in heavy ion collisions. (b) Spectators and participants.

⁵The author is a member of the BRAHMS collaboration.

⁶For example SLAC in Stanford, California, USA.

Because the radius of a ^{207}Pb nucleus is approximately 8 fm [17], one finds that $b \in [0, 16]\text{fm}$ in Pb–Pb, where $b = 0\text{ fm}$ corresponds to a head-on, or central, collision, and $b \gtrsim 10\text{ fm}$ corresponds to grazing, or peripheral, collisions.

As can be seen from Figure 1.7 the colliding nucleons may be categorised into two groups: the *participant* and the *spectator* nucleons. The participants are the nucleons that scatter, or interact in the collision, while the spectators essentially go on unaffected by the collision, and do not contribute to the particle production.

The impact parameter is not a directly measurable quantity — instead one characterises the overlap of the collisions partners by *centrality* (see Section 2.1). However, the number of participants as a function of b can be estimated using various models of the nuclear density distribution. Figure 1.8 shows the number of participants in Pb–Pb as a function of the impact parameter from the HIJING event generator [18, 19].




Figure 1.8: Number of participants (a) and produced charged particles (M_{ch}) (b) versus b estimated from HIJING [18, 19] in Pb–Pb at $\sqrt{s_{NN}} = 5500\text{ GeV}$.

1.4.2 Transparency and stopping

Often the colliding particles are accelerated to speeds that are a considerable fraction of the speed of light, and so the collisions must be treated relativistically. At LHC, for example the Pb nuclei travel at a speed of approximately 99.999993 % of the speed of light, corresponding to $\gamma = 2750$. That has the effect that distances in the direction of the velocity, as seen from the laboratory frame and the rest frame of the opposing nucleus, become shorter than in the rest frame of the Au nucleus by $\frac{1}{\gamma}$. Hence the Au nuclei, as seen from the centre of mass system, become flat like pancakes; they are Lorentz-contracted.

Due to baryon-conservation, the net number of baryons after the collision, is the same as before. However, during the collision, the colliding baryons and quarks may transfer momenta to one another, or form new particles. That means, that the phase space distribution of the baryons need not be the same after the collision as before.

This leads to two scenarios for the collision systematics. In the first scenario applicable for low energies ($\sqrt{s_{NN}} \approx 10\text{ GeV}$), the constituent particles of the colliding nuclei are completely stopped at the collision point. This is called *full stopping*, and is depicted in Figure 1.9. In this

hot and dense fireball the nuclei will essentially overlap and possibly undergo a phase transition to deconfined and chiral symmetric matter.

The degree of stopping is normally defined as the average rapidity loss of the nuclei's nucleons $\langle y - y_{beam} \rangle$, where y is the rapidity of the nucleons, and y_{beam} is the rapidity of the beam nucleons before the collision. Transparency, explained below, can then be thought of as 1 minus the degree of stopping.




Figure 1.9: Stopping in heavy ion collisions.

However, at higher energies, the Lorentz-contraction of the colliding nuclei limits the amount of nuclear matter the colliding nucleons (and their constituent quarks) will see during the collision. The time t needed in the rest frame of the nucleon to produce a particle follows from $\Delta E \Delta t > \hbar$ and gives $t \approx 1$ fm. In the laboratory frame this is Lorentz-dilated to $t' = \gamma t$. Due to the extremely large γ in relativistic collisions ($\mathcal{O}(1000)$) $t' \approx 1000$ fm; large enough that the two colliding nuclei has long since passed each other before the constituent particles can interact a second time. This is called *transparency* and is depicted in Figure 1.10.




Figure 1.10: Transparency in heavy ion collisions.

So, in effect the two nuclei become transparent to each other, and each nucleon will only interact a few times. Their constituent quarks will drag out their interaction gluons. Eventually the gluonic zone between the colliding nuclei will have such a high energy density that the colour field will break (like rubber-bands in the allegory introduced in Section 1.1) and produce $q\bar{q}$ pairs in a hot and dense zone, possibly forming deconfined and chiral symmetric matter.

The degree of stopping can be studied via the ratio of \bar{p} to p as a function of the rapidity y . In the fully transparent scenario, the ratio becomes unity over the full range up the rapidity of the beams, due to charge, baryon, and isospin conservation in the collision. In a full stopping scenario, the ratio is affected by the collision nucleons left over in the central region ($y \approx 0$).

BRAHMS at RHIC has studied this behaviour for $\sqrt{s_{NN}} = 62.4$ GeV, 130 GeV, and 200 GeV and finds a \bar{p}/p of 0.63 [20] and 0.75 [21] at $\eta \approx 0$ respectively (see also Figure 1.11), suggesting that full transparency has not been reached at RHIC energies. The AMPT model from Texas A&M [22] suggests that some rescattering takes place in the initial stages of the collision transporting beam nucleons into the central region, and thus giving a ratio \bar{p} to p less than 1.




Figure 1.11: Ratios of π^-/π^+ , K^-/K^+ , and \bar{p}/p as a function of collision energy (left panel) and as a function of rapidity at $\sqrt{s_{NN}} = 200$ GeV (right panel) [23].

BRAHMS has also studied the average rapidity loss δy as a function of the collision energy. Figure 1.12 shows the net protons ($N_p - N_{\bar{p}}$) as a function of rapidity for several collision energies and the average rapidity loss as a function of $\sqrt{s_{NN}}$. Clearly the net proton (baryon) poor region becomes wider as the collision energy increases. Unfortunately, it will not be possible for the ALICE experiment at LHC to measure the average rapidity loss, since ALICE does not have coverage and particle identification capabilities in the forward rapidity region.

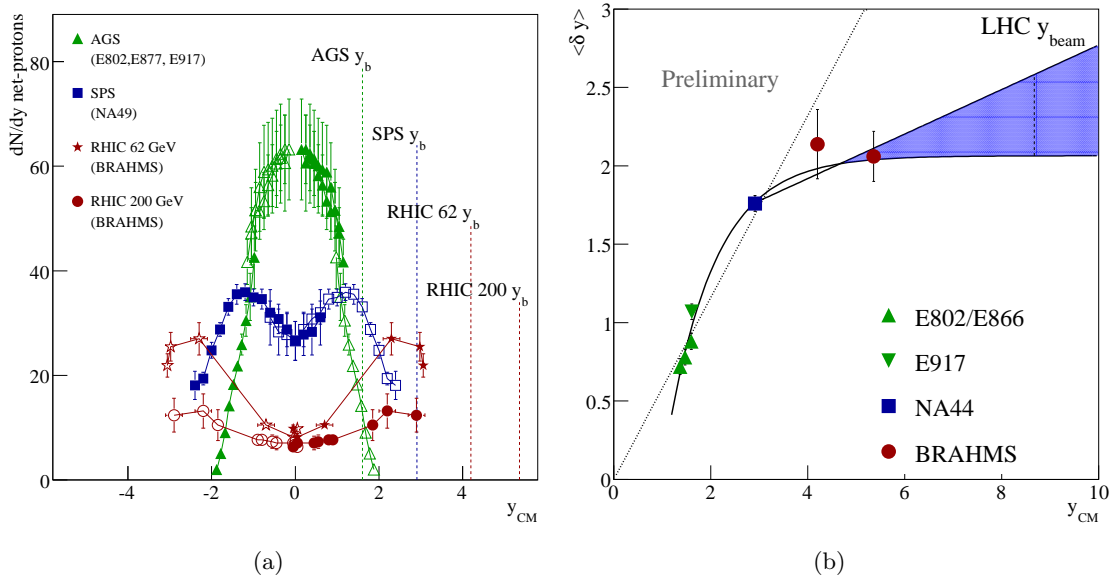


Figure 1.12: Net protons as a function of rapidity (a) and the average rapidity loss as a function of collision energy (b) (adapted from [20, 24, 25]).

1.4.3 Chemical and Thermal Freeze-out

In relativistic heavy ion collisions an abundance of particles are created in the initial stages of the interaction. Therefore there is ample opportunity for the particles to reinteract to produce new particles or transfer momenta.

When the particles no longer reinteract to produce new particles, the system is said to have reached *chemical freeze-out*. That is, the chemical potential of the various species, like protons, pions, kaons, and so on, does not change any further — except via secondary interactions in material out of the collision zone (such as the beryllium beam pipe or iron magnets), or decays of short lived particles (such as Λ 's).

When the system has reached chemical freeze-out the particles may still interact to transfer momenta. When the particles from the collision essentially no longer transfer momenta, the system is said to have reached *thermal freeze-out*. That is, the mean velocity (or temperature) of the system does not change any more.

1.5 Signatures for QGP

The detection of QGP in relativistic heavy-ion experiments is unfortunately not a straight forward matter. The number of produced charged particles in gold on gold collisions at $\sqrt{s_{NN}} = 200$ GeV is around 5000 [26], and in $\sqrt{s_{NN}} = 5500$ GeV Pb–Pb around 40000 produced charged particles is expected. Most of the particles have interacted after the initial production, thereby washing out many of the characteristics of the original production in the collision. A number of *signatures* for the detection of QGP, have been proposed.

1.5.1 Kinematics

Temperature The phase transition can be studied by the inversion of Figure 1.4(a), and the associations $T \leftrightarrow \langle p_{\perp} \rangle$, and $\varepsilon \leftrightarrow \frac{dN}{dE_{\perp}}$. The temperature is often determined as the inverse slope of the single particle spectrum [27]

$$\frac{d^3N}{dp^3} = \frac{Y}{4\pi T m^2 K_2\left(\frac{m}{T}\right)} e^{-\frac{E}{T}} \quad , \quad (1.1)$$

where Y is the total yield, m the rest mass of the particle, and p is the momentum of the particles. The idea is to study this inverse slope parameter as a function of varying collision energy. However, the extraction of T and interpretation of this parameter as the temperature of the collision, is strongly dependent on the assumption that the hot fireball of the collision is in thermal and chemical equilibrium. This assumption is far from trivial, as it is questionable whether the system has the time to become equilibrated. However, if the critical temperature is high enough (> 200 GeV), the thermalisation time maybe as short as 0.4 fm, giving ample time for reaching equilibrium [9]. The STAR experiment at RHIC has published result from the $\sqrt{s_{NN}} = 130$ GeV run [28], BRAHMS has published p_{\perp} spectra for Au–Au and d–Au at $\sqrt{s_{NN}} = 200$ GeV [24,29–31].

ALICE will be able to measure the temperature and yields of various particles near mid-rapidity with a very high precision with ITS, TPC, TRD, and TOF.

Particle Ratios The ratio of anti-particle to particle number, like $\frac{\bar{p}}{p}$, $\frac{K^+}{K^-}$ and $\frac{\pi^+}{\pi^-}$ as a function of rapidity, provides extensive information on the kinematics of the collision. Due to charge, baryon, isospin, strangeness, etc. conservation, particles are assumed to be produced in particle anti-particle pairs. Therefore, an excess of particles in these ratios signals that something other than pure kinematics, like pair production, is at play in the collision. The $\frac{\bar{p}}{p}$ near $y = 0$ provides information on how much of the colliding nuclei is left over in the various rapidity regions of the collision zone. All four RHIC experiments have published particle ratios from the $\sqrt{s_{NN}} = 130$ GeV

run [20, 32–34] and BRAHMS has published results from the $\sqrt{s_{NN}} = 200$ GeV run [23, 31], see Figure 1.11.

1.5.2 Dynamics

Strangeness Enhancement Due to the abundance of quark anti-quark pairs produced close in phase space in the heavy ion collisions, the Fermi sea for u and d quarks is expected to be quickly filled, leaving it energetically favourable to produce $s\bar{s}$ pairs. Hence the observation of strangeness enhancement, is a good probe for the formation of a hot and dense phase of nuclear matter, in particular the ratios $\frac{N_{K^-}}{N_{\pi^-}}$ and $\frac{N_{K^+}}{N_{\pi^+}}$ are of particular interest. BRAHMS has published results on these ratios for Au–Au at $\sqrt{s_{NN}} = 200$ GeV [21, 30, 31].

On the other hand the production of $c\bar{c}$ (charmonium) is expected to be strongly suppressed in the deconfined phase. This shows up in the suppression of the formation of J/Ψ particles. NA50 of CERN has published results to that effect [35], and it is the intend of PHENIX to look for this signal at $\sqrt{s_{NN}} = 200$ GeV too. ALICE has a large dedicated μ spectrometer to detect dileptons, which will measure the amount of J/Ψ suppression in $\sqrt{s_{NN}} = 5500$ GeV Pb–Pb collisions.

Jet Quenching Another hadronic signature of QGP that recently received a lot of attention is *jet quenching*, which shows up as a suppression of high p_{\perp} ⁷ particles. From particle physics it is well known that jets can be produced by gluon radiation. In the hot and dense medium formed in heavy-ion collisions, it was not expected to see jets in any significant amount. However, not only are jets observed, but an interesting phenomenon is observed in jet production. When a jet is formed, one expects, due to momentum conservation, to see a corresponding jet in the opposite direction — the so called away-side jet. However, at RHIC, the away-side jet almost disappears in central Au–Au collisions. This has been investigated in detail by STAR [36–39]. The broadening of the away-side jet is seen as an indication of the quarks and gluons interacting strongly with the hot and dense medium [40], see Figure 1.13.

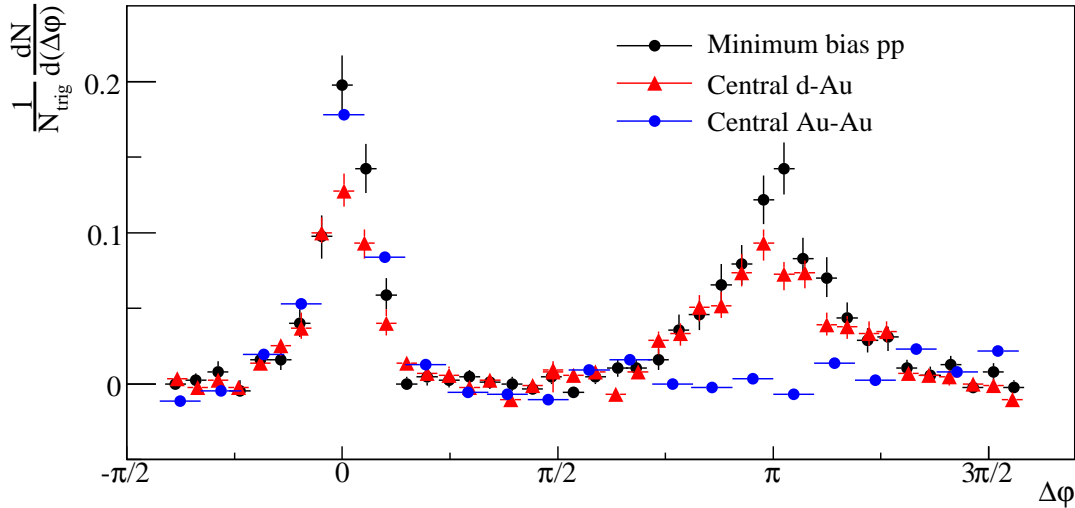


Figure 1.13: The disappearance of the away-side jet in central Au–Au at $\sqrt{s_{NN}} = 200$ GeV as published by STAR. $\Delta\phi$ is the angle between any two particles. Adapted from [37].

Jet quenching has also been studied by the other 3 experiments at RHIC, via the so-called

⁷See Appendix A.

nuclear modification factor [41]

$$R_{AA} = \left\langle \frac{\sigma_{inel}}{N_{binary}} \right\rangle \frac{d^2 N_{AA}/dydp_{\perp}}{d^2 \sigma_{pp}/dydp_{\perp}} ,$$

where σ_{inel} is the total inelastic cross section, N_{binary} is the number of binary collisions, $d^2 N_{AA}/dydp_{\perp}$ is the yield observed in A–A, and $d^2 \sigma_{pp}/dydp_{\perp}$ is the pp differential cross section. The study of this ratio using BRAHMS Au–Au and Cu–Cu data at $\sqrt{s_{NN}} = 200$ GeV and 62.4 GeV has been the subject of two fellow Ph.D. students [41, 42].

Azimuthal Anisotropy Another observable that has generated excitement is the *azimuthal anisotropy* (or v_2) of the particles emitted from the hot and dense fireball. More will be said about v_2 in Chapter 2.3, but for completeness a few words will be said here. The relatively large anisotropy observed in non-central collisions, can, in the framework of thermal models, be interpreted to mean that the quarks and gluons thermalises rather quickly. The thermalisation may happen as quickly as 0.6 fm [43], thereby preserving the initial anisotropy of the collision. This, on the other hand, indicates that the fireball can be seen as a perfect liquid of quarks and gluons moving around relatively freely. Together with the jet quenching results, this has prompted people to formulate the idea of sQGP — a *strongly* interacting perfect fluid of deconfined quarks and gluons.

A number of detectors in ALICE will be used for the measurement of the azimuthal anisotropy in various way. The ITS and TPC will measure the transverse momentum and azimuthal distribution of charged particles, and can determine the anisotropy around $y \approx 0$. The FMD can measure the azimuthal distribution of charged particles in the forward regions, and determine the anisotropy in that region. The FMD can also be used to determine the event plane independently of the ITS and TPC, and thereby reduce the uncertainty in the v_2 measurements by these detectors.

1.5.3 Chiral Symmetry Restoration — Disoriented Chiral Condensates

It is expected that extremely hot and dense nuclear matter, such as in QGP, will undergo (partial) chiral symmetry restoration. Figure 1.14 shows the potential in the linear sigma model [44] under various conditions. Initially, the quark matter is in the asymmetric state of spontaneous and explicitly broken symmetry, as seen in Figure 1.14(a). As the quarks become deconfined, the spontaneous broken symmetry is restored, while the explicit persists, giving a potential as shown in Figure 1.14(b). When the system cools down and becomes less dense, the symmetry is again spontaneously broken, and the potential is once again that of Figure 1.14(a). For reference, Figure 1.14(c) shows the fully symmetric potential.

This has some rather interesting experimental effects. As the symmetry is partially restored, the π and ρ meson masses will change. Since the ρ meson decays into dileptons, which do not interact strongly with the surrounding nuclear matter, one can hope to measure the change in the ρ mass, by measuring the invariant mass spectrum of the dileptons. PHENIX at RHIC has not observed any significant change to the ρ mass in the $e^- + e^+$ or $K^- + K^+$ decay channels [45].

Another interesting phenomenon, is *Disoriented Chiral Condensates* (or DCC) [46]. When the symmetry is again spontaneously broken, the π masses initially ‘sit’ on top of the local maximum of the potential, and suddenly acquire mass by seeking toward the global minimum. Since the π masses (one for π^{\pm} and one for π^0) can get to the minimum by a variety of paths down the slope of the potential, there is no reason to believe that π mesons of the 3 different charges will be produced in equal numbers, as is the case in normal nuclear matter. One could therefore expect to see

$$\frac{\pi^0}{\pi^+ + \pi^- + \pi^0} \neq \frac{1}{3} , \quad (1.2)$$

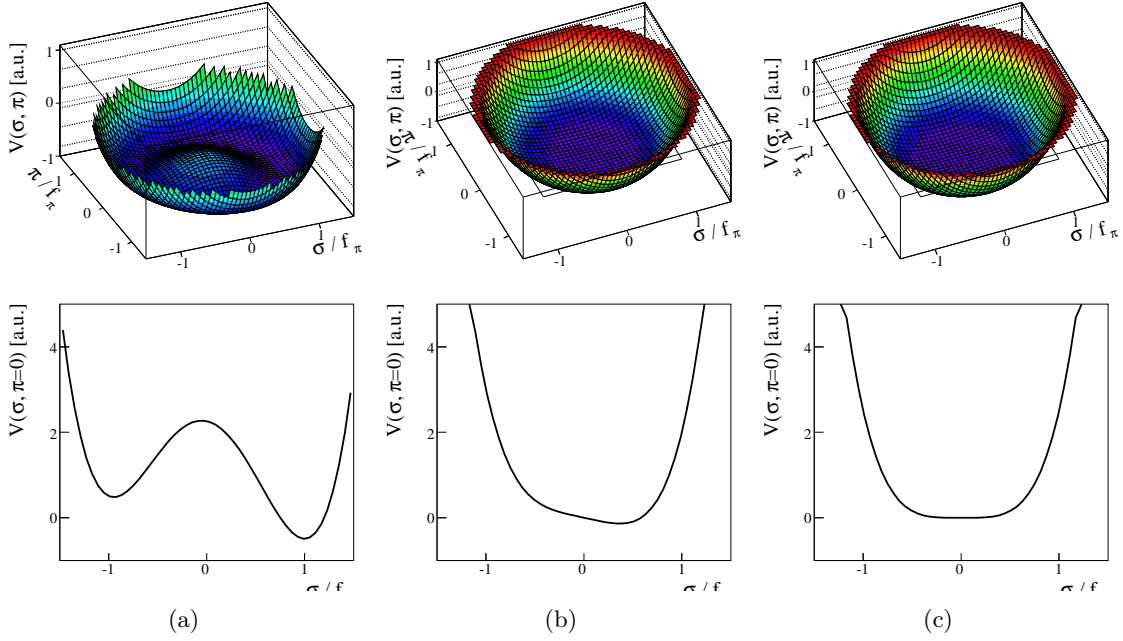


Figure 1.14: The potential V between particles in a linear σ model [44]. (a) The explicitly and spontaneously broken symmetry, (b) The explicitly broken and spontaneous symmetry (partial restoration), (c) fully symmetric. The lower row are the same as the upper row, with $\pi = 0$. Note that the temperature increases from left to right. The potential is shown with arbitrary units.

giving a clear signal of chiral symmetry restoration [14]. This can be studied by looking at fluctuations in electric charge asymmetries, such as $\frac{\pi^0}{\pi^\pm}$ ratios on an event by event basis.

Although chiral symmetry restoration should occur if quarks become deconfined, it has not been observed in the overwhelming amount of data produced at RHIC. It is hoped that clear signals will be found by ALICE at the LHC.

Chapter 2

Observables Relevant to the ALICE FMD

In this chapter, a number of observables relevant to the ALICE FMD will be given some theoretical consideration. In particular, the charged particle density and azimuthal anisotropy will be described.

2.1 Impact Parameter and Centrality

s

The impact parameter \mathbf{b} , introduced in Section 1.4.1, is an important parameter in heavy-ion collisions. For example, in the HIJING model, as shown in Figure 1.8, the number of interacting particles, and consequently the number of produced particles, depends strongly on the overlap of the colliding nuclei. Hence, the physics of the heavy-ion collisions can be rather different from a head-on collision ($b \approx 0$) to mere grazing strokes of large impact parameter.

As mentioned in Section 1.4.1, the length b of the impact parameter \mathbf{b} is not directly measurable by experiments. Instead, experimentalists use a measure known as *centrality*. The exact definition of centrality varies slightly from experiment to experiment, however; it usually involves a cut in the total observed charged particle multiplicity, M_{ch} (see Figure 2.1(a)). The cuts define centrality *classes* like 0-5 %, 5-10 %, 10-20 %, and so on.

The mapping from impact parameter b to centrality is generally not one-to-one, as shown in Figure 2.1(b). Therefore, the inverse mapping, from M_{ch} to b , although monotonic, is ambiguous, and generally difficult to do. Instead, the number of participants N_{part} and number of binary collisions N_{bin} is calculated for a given centrality class, using a Glauber model (see for example [47, 48]). The procedure is to use the Wood-Saxon parameterisation as an approximation for the nuclear density

$$\rho(r) = \frac{\rho_0}{1 + e^{\frac{r-R_A}{a}}} \quad ,$$

to define the thickness function

$$T_A(\mathbf{s}) = \int dz \rho(\mathbf{s}, z) = \int dz \frac{\rho_0}{1 + e^{\frac{\sqrt{s^2+z^2}-R_A}{a}}} \quad ,$$

where \mathbf{s} spans the transverse (x, y) plane. The number of participants and binary collisions is then given by

$$N_{\text{part}}(b) = \int d^2\mathbf{s} T_A(\mathbf{s})(1 - e^{-\sigma_{NN}T_B(\mathbf{s})}) + \int d^2\mathbf{s} T_B(\mathbf{s} - \mathbf{b})(1 - e^{-\sigma_{NN}T_A(\mathbf{s})}) \quad , \quad (2.1)$$

$$N_{\text{bin}}(b) = \int d^2\mathbf{s} \sigma_{NN} T_A(\mathbf{s}) T_A(\mathbf{s} - \mathbf{b}) \quad , \quad (2.2)$$

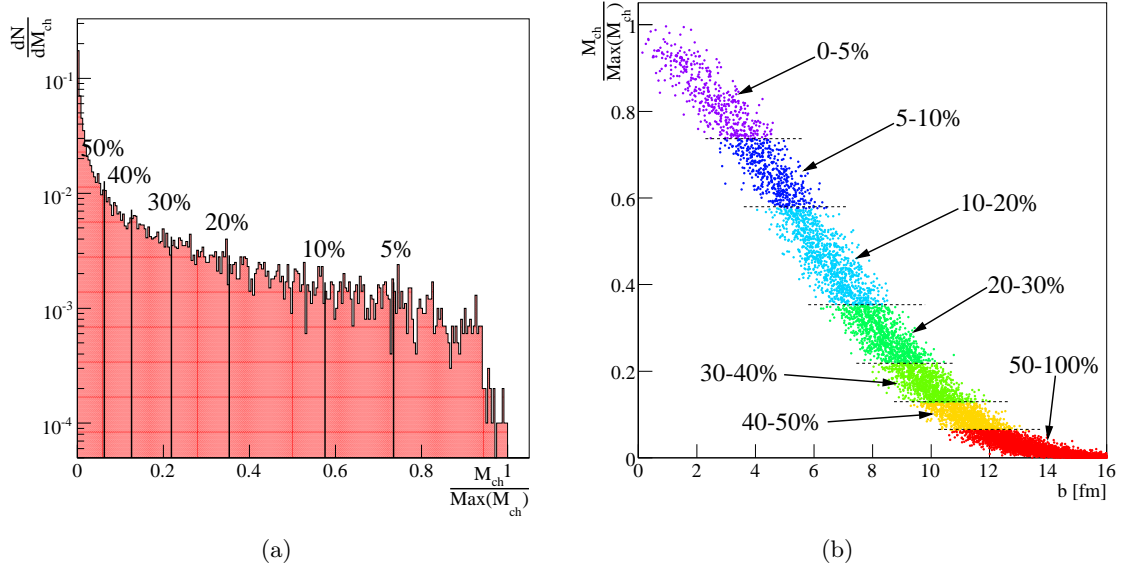


Figure 2.1: (a) Definition of centrality cuts from observed total charged particle multiplicity. The data in the figure is from a 10,000 HIJING Pb–Pb at $\sqrt{s_{NN}} = 5500$ GeV. The x axis has been scaled to the maximum observable total charged particle multiplicity. (b) Comparison of centrality classes to impact parameter. Clearly, the mapping from impact parameter b to centrality class is not one-to-one, and the mapping is therefore not invertible.

where σ_{NN} is in the total inelastic nucleon–nucleon cross-section. One can then relate the observed centrality classes to N_{part} and N_{bin} by evaluating the full cross section and calculating N_{part} and N_{bin} for each sample to get a distribution and cuts like in Figure 2.1(a).

Good determination of a given events centrality requires a relatively good measurement of M_{ch} . That is, the experiment should strive measure all charged particles produced in the collision, to remove uncertainties in the centrality determination. The ALICE FMD and ITS together will measure $\approx 83\%$ of the total number of produced charged particles, providing an excellent centrality determination.

2.2 Charged Particle Density ($dM_{ch}/d\eta$)

The charged particle density $dM_{ch}/d\eta$ is the number of charged particles (p , \bar{p} , π^\pm , K^\pm , and so on) per unit of pseudo-rapidity η . The measurement of the $dM_{ch}/d\eta$ is a baseline measurement, in that it gives provides a first-level discriminatory test of the various models of heavy-ion collisions. Figure 2.2(a) illustrates this for central Au–Au collisions at $\sqrt{s_{NN}} = 200$ GeV. The points represents various predictions for the charged particle *rapidity* density, while the vertical line and band shows the observed scaled¹ pseudo-rapidity charged particle density. Clearly, most models over estimated the charged particle production at $\sqrt{s_{NN}} = 200$ GeV, and expectations for Pb–Pb have consequently been downscaled.

If a model cannot reproduce the observed $dM_{ch}/d\eta$, it is likely that the model is at best insufficient, at worst poor, in describing the mechanism at play in heavy-ion collisions.

$dM_{ch}/d\eta$ together with the average m_\perp of the produced particles as a function of η , can be

¹ [49] scales $dM_{ch}/d\eta$ by a factor of 1.1 to get $\frac{dM_{ch}}{dy}$.

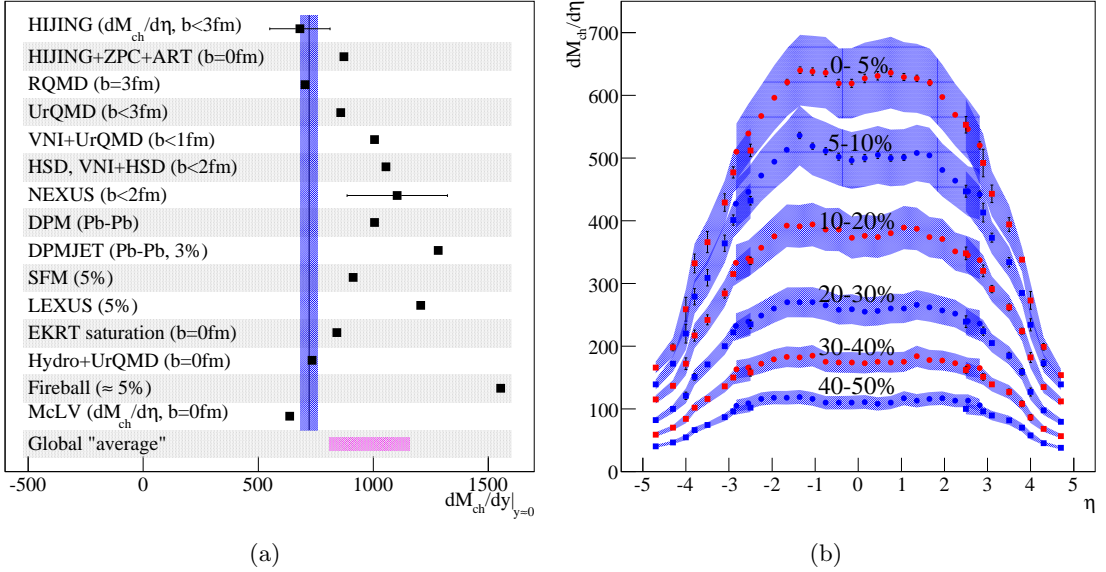


Figure 2.2: (a) Comparison of various model predictions of the charged particle density in Au–Au at $\sqrt{s_{NN}} = 200$ GeV compared to the PHOBOS observation, indicated by the vertical line. The vertical band represents the systematic errors on the PHOBOS data, while the horizontal band is the average of all the predictions. Adapted from [49]. (b) Charged particle density in Au–Au at $\sqrt{s_{NN}} = 200$ GeV at different centralities, as measured by BRAHMS. Adapted from [50].

used to estimate the energy density of the hot and dense fireball, using Bjorkens formula [51]

$$\begin{aligned} \epsilon &= \frac{1}{\pi R^2 \tau} \frac{d\langle E_{\perp} \rangle}{dy} \\ &= \frac{1}{\pi R^2 \tau} \frac{\langle m_{\perp} \rangle dN}{dy}, \end{aligned} \quad (2.3)$$

where R is the radius of the overlap zone, and τ is the formation time.

The BRAHMS experiment at RHIC has published results on the energy density based on these $dM_{ch}/d\eta$ (see Figure 2.2(b)) measurements in Au–Au at $\sqrt{s_{NN}} = 200$ GeV, and finds an energy density of ≈ 5 GeV/fm³ [50] assuming a formation time of 1 fm.

2.2.1 Overview of previous $dM_{ch}/d\eta$ measurements

It is instructive to compare the measurement of the charged particle density at mid-rapidity at the top RHIC energy to previous measurements at lower energies. Figure 2.3(a) shows $dM_{ch}/d\eta$ at mid-rapidity $|\eta| < 1$ from previous measurements in Au–Au at AGS, Pb–Pb at the SPS, and in Au–Au at RHIC at all four energies. The density has been scaled down by the average of half the number of participants $\langle N_{part}/2 \rangle$ in the collisions.

The data suggests that the bulk of the charged particle distribution scales logarithmically with the beam energy. This is further underlined by the overall shape of the charged particle density distribution as shown in Figure 2.3(b). The shape does not change significantly as the beam energy increases from the lowest RHIC energy at $\sqrt{s_{NN}} = 19.6$ GeV to the highest at $\sqrt{s_{NN}} = 200$ GeV.

The rather slow increase in the charged particle production is surprising compared to the expectations from perturbative QCD (pQCD) calculations. From the factorisation theorem of

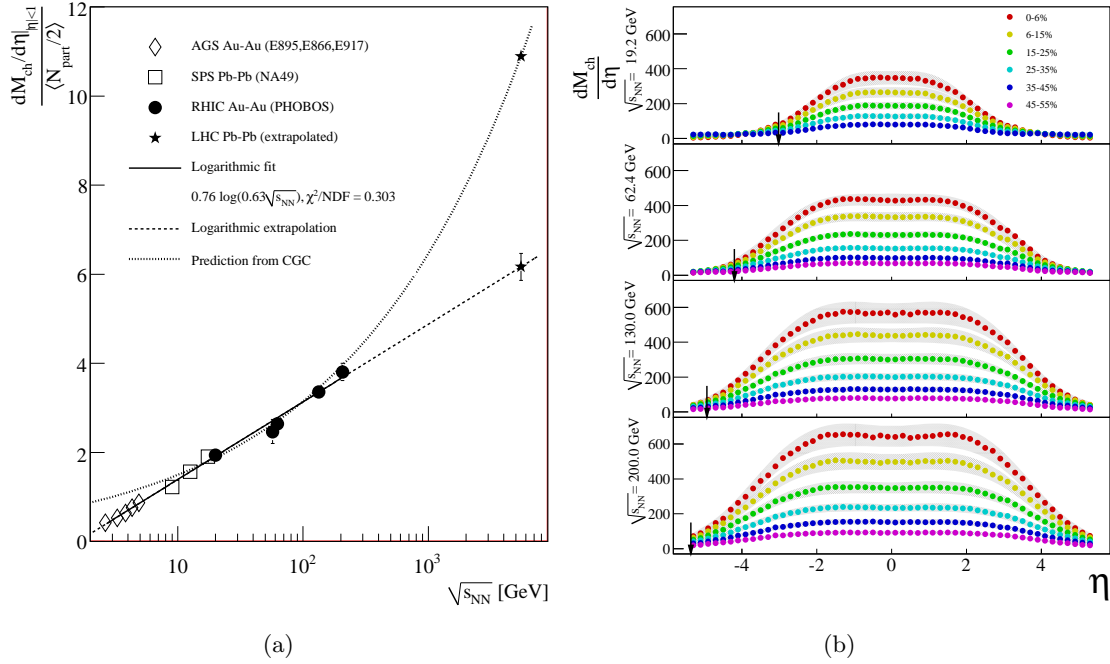


Figure 2.3: (a) $dM_{ch}/d\eta|_{|\eta|<1}$ in central events, scaled by the average N_{part} , as a function of $\sqrt{s_{NN}}$. The dashed line shows a logarithmic extrapolation of the lower energy data to LHC energies. Extrapolating to LHC energies gives a value of $\approx 6.2 \pm 0.31$. With an average number of participants of ≈ 369 at $\sqrt{s_{NN}} = 5500$ GeV, this gives an estimate of circa 1140 ± 57 produced charged particles in the mid-rapidity region of ALICE. The dotted line shows the prediction from CGC (see (2.5)). Adapted from [49, 52]. (b) $dM_{ch}/d\eta$ as a function of centrality, at the 4 RHIC energies $\sqrt{s_{NN}} = 19.6$ GeV, $\sqrt{s_{NN}} = 62.4$ GeV, $\sqrt{s_{NN}} = 130$ GeV, and $\sqrt{s_{NN}} = 200$ GeV, as measured by the PHOBOS experiment. Grey bands mark the systematic errors, while the statistical errors are smaller than the points, and the arrows indicate the beam rapidity. Adapted from [3, 4, 49].

pQCD [53], one expects that the hard processes² will lead to a charged particle production that scales with the number of binary collisions N_{bin} [54].

The N_{part} scaling can be understood in the framework of *Color Glass Condensate* (CGC), which describes the initial state³ of the colliding nuclei. Larry McLerran describes the CGC as [55]

Color The gluons that make up this matter [CGC] are colored.

Glass The gluons at small x^4 are generated from gluons at larger values of x . In the infinite momentum frame, these larger momentum gluons travel very fast and their natural time scales are Lorentz time dilated. This time dilated scale is transferred to low x degrees of freedom which therefor evolve very slowly compared to natural time scales. This is the property of a glass.

Condensate The phase space density

$$\rho = \frac{1}{\pi R^2} \frac{dM_g}{dyd^2p_\perp}$$

is generated by a trade off between a negative mass-squared term linear in the density which generates the instability, $-\rho$ and an interaction term $\alpha_S \rho^s$ which stabilizes the system at a phase space density $\rho \approx 1/\alpha_S$. Because $\alpha_S \ll 1$, this means that the quantum mechanical states of the system associated with the condensate are multiply occupied. They are coherent, and share some properties of a Bose condensate. The gluon occupation factor is very high, of order $1/\alpha_S$, but it is only slowly (logarithmically) increasing with the energy, or decreasing with the transverse momentum. This provides saturation and cures the infrared problem of the traditional BFKL approach.

The motivation of the CGC came from observations of the gluon density $xG(x, Q^2)$ in the hadron at the HERA experiments [55], shown in Figure 2.4(a). Here, x is the fraction of the total hadron momentum carried by a parton i.e., a quark, anti-quark, or gluon, and Q is the momentum transfer in the $e^- + p$ collision.




Figure 2.4: (a) The gluon density $xG(x, Q^2)$ at different momentum transfers Q . (b) Saturation of the gluon density in CGC.

The gluon density is seen to scale logarithmically with $1/x$ at higher Q , while at the same time, the cross-section only grows slowly with energy. There must therefor be some saturation scale Q_s where the density of gluons become saturated [56]. Again quoting Larry McLerran [55]:

²Large momentum transfers.

³I.e., before the collision.

⁴ x is the fraction of the total nucleon momentum carried by a parton — i.e., a quark, anti-quark, or gluon.

This led to the conjecture that the density of gluons should become limited, that is, there is gluon saturation. Actually, we argue that as one goes to higher energy, a hadron becomes a tightly packed system of gluons larger than some size scale. For smaller gluons, there are holes. As one increases the energy, one still adds in more gluons, but these gluons are small enough that they fit into the holes. Because in quantum mechanics, we interpret size as wavelength which is inversely proportional to momentum, at higher energies, the gluons are tightly packed below some momentum, the saturation momentum, which characterizes the filling. This saturation momentum increase as the energy increases, so the total number of gluons can increase without bound.

As the gluons are more tightly packed, the interaction between them becomes weaker and weaker due to the asymptotic freedom in QCD. That means that a probe with a low wavelength, or equivalently at higher momentum, sees a denser and denser medium of gluons, which partially shadow each other. The probe therefor interacts coherently with the gluons, and the interactions, up to some saturation scale, can be described in terms of classical fields. This allows one to analytically calculate the particle production as a function of rapidity below the saturation scale, while the for higher momentum pure pQCD can be used. One finds that for the gluon production the following approximations [57, 58]

$$\frac{1}{\pi R^2} \frac{dN_g}{d^2 p_\perp dy} \approx \begin{cases} \frac{1}{\alpha_S} \frac{Q_S^4}{p_\perp^4} & \text{above saturation} \\ \frac{1}{\alpha_S} Q_S^2 & \text{below saturation} \end{cases}, \quad (2.4)$$

where πR^2 is the nuclear overlap. Since the saturation scale should be proportional to $N_{part}^{1/3}$, this leads to

$$\frac{1}{N_{part}} \frac{dN_g}{dy} \approx \frac{1}{\alpha_S} \approx \log N_{part} \quad .$$

That is, the number of gluon produced scales with the number of participants — not the number of binary collisions. Kharzeev and Levin [59] have used these result to extract the $dM_{ch}/d\eta$ as predicted by the CGC initial state (see Figure 2.5) and find good agreement with data.




Figure 2.5: $dM_{ch}/d\eta$ at different beam energies and centralities. The data is the same as in Figure 2.3(b), and the lines are the predictions from CGC at the same beam energies and centralities [59].

Another indication of the validity of CGC is provided by what is known as *Limiting Fragmentation*. If one plots the charged particle density, not as a function of η but as a function of

the scaled $\eta' = |\eta| - y_{beam}$ where y_{beam} is the beam rapidity, one sees that the shape of charged particle production at a given centrality is roughly the same for all energies (see Figure 2.6 and Figure C.1).

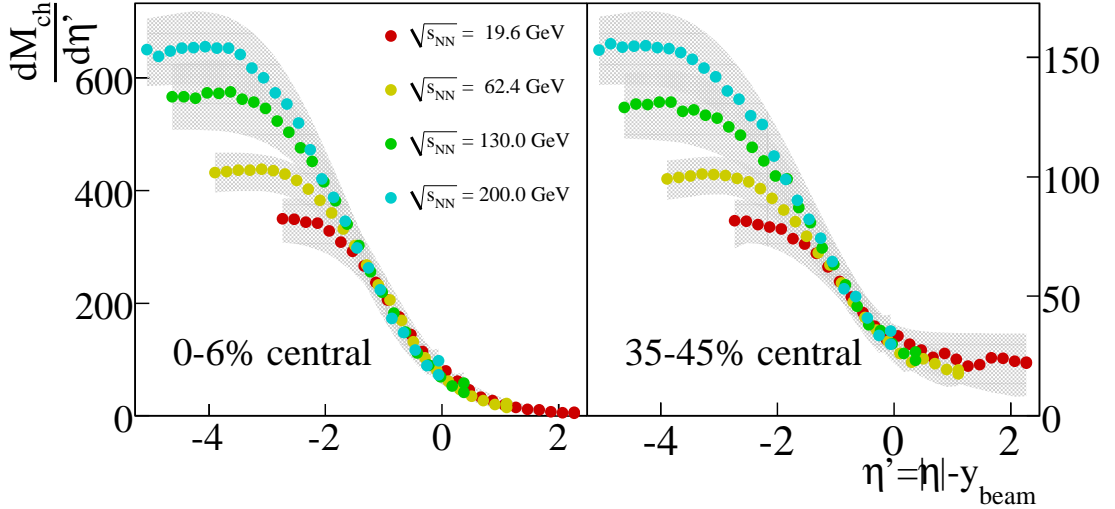


Figure 2.6: The charged particle rapidity as a function of the scaled pseudo-rapidity $\eta' = |\eta| - y_{beam}$. The statistical and systematic errors are same as in Figure 2.3(b) i.e., the statistical errors are smaller than the points, and the systematic errors are marked with a gray band. Adapted from [3, 4].

As we approach higher (pseudo-) rapidity, it corresponds to probing at lower values of x . Hence, in the forward region, there are fewer and fewer high x sources [55] feeding the particle production so that the effect is that the production becomes limited and largely independent of the collision energy in that region.

The limiting fragmentation observed at RHIC by BRAHMS [26], and by PHOBOS [3] also indicates that there is no central plateau in the charged particle production. Instead, the collision can be seen as leaving the two colliding nuclei in an excited state from which they radiate particles near, or very near, to nucleic pseudo-rapidity range.

Note, that the CGC describes the initial state of the colliding nuclei. To calculate observables, like $dM_{ch}/d\eta$, one need to feed that initial state into a model of gluon jet production [59] to calculate the properties of the produced particles, and then feed these into a fragmentation and evolution calculation to get the final state observable particles. One popular method is to use the CGC calculations as an initial state input to a hydrodynamic calculation [58] which gives very reasonable results (see for example Figure 2.2(a) — the lowest point marked “McLV” is an example of such a calculation.)

2.2.2 Lessons to be learned from the $dM_{ch}/d\eta$ measurements

Extrapolating the simple-minded logarithmic parameterisation in Figure 2.3(a) to LHC energies ($\sqrt{s_{NN}} = 5500$ GeV), we find that the charged particle multiplicity per participant pair at mid rapidity should be $\approx 6.2 \pm 0.31$. With an average number of participants of ≈ 396 in Pb-Pb at $\sqrt{s_{NN}} = 5500$ GeV, this gives an estimate of

$$\left. \frac{dM_{ch}}{d\eta} \right|_{|\eta| < 1} \approx 1140 \pm 57$$

produced charged particles in the mid-rapidity region of ALICE. This is far lower than most early expectations for the LHC. Of course, since the extrapolation is over an order of magnitude in

$\sqrt{s_{NN}}$, and that the logarithmic parameterisation is without prior assumptions, it should not be overly emphasised.

On the other hand, the results of the CGC calculations are rather robust, and indicate a simple N_{part} and $\sqrt{s_{NN}}$ scaling [52]

$$\begin{aligned} \left. \frac{dM_{ch}}{dy} \right/ \left. \frac{N_{part}}{2} \right|_{y=0} &= 0.87 \left(\frac{\sqrt{s_{NN}}}{\sqrt{s_{NN,0}}} \right)^{\tilde{\lambda}} \log \left(\frac{Q_s^2(A, \sqrt{s_{NN}}, y=0)}{\Lambda_{QCD}} \right) \\ &= 0.87 \left(\frac{\sqrt{s_{NN}}}{130} \right)^{0.252} \left(3.93 + 0.252 \log \left(\frac{\sqrt{s_{NN}}}{130} \right) \right) , \end{aligned} \quad (2.5)$$

see also Figure 2.3(a). Here, Q_S is the saturation scale, A the atomic weight of the nucleus, and the parameter $\tilde{\lambda}$ is fixed from fits to HERA data. Multiplying by $N_{part}/2 = 369/2$ gives an estimate of ≈ 2000 charged particles in the mid-rapidity region of ALICE.

CGC also predicts little change in the shape of the charged particle density as shown in Figure 2.7, which in turn means that limiting fragmentation should prevail at LHC energies.

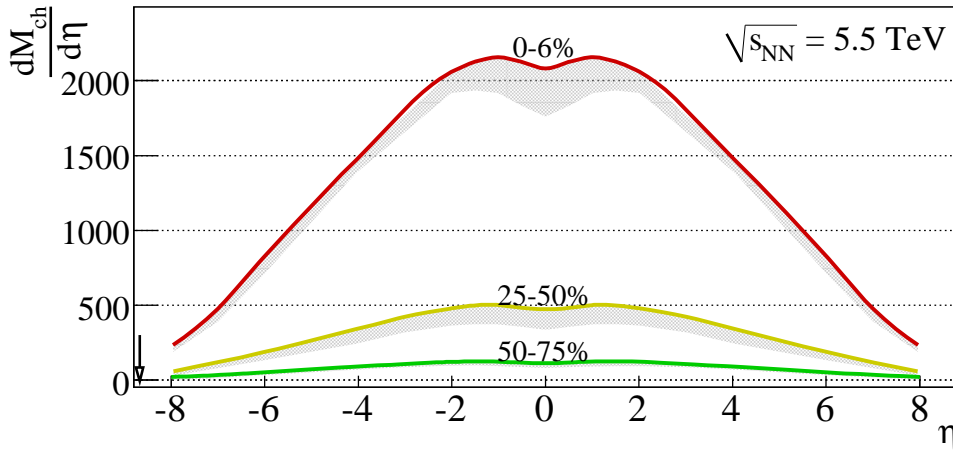


Figure 2.7: CGC prediction on $dM_{ch}/d\eta$ at different centralities in Pb-Pb at $\sqrt{s_{NN}} = 5500$ GeV. The arrow indicates the beam rapidity. Adapted from [52].

2.3 Azimuthal Anisotropy (v_2)

In non-central heavy-ion collisions, the overlap region of the colliding nuclei has a strong asymmetry, simply due to the geometry of the collision (see Figure 2.8). This spatial asymmetry can, within hydrodynamic models [43], propagate into the momentum-space of the produced particle around mid-rapidity.

2.3.1 Fourier decomposition of the azimuthal particle distribution

We define the *event plane* P_E as the plane spanned by the collision axis (normally defined as z axis) and the impact parameter vector \mathbf{b} . As mentioned earlier, \mathbf{b} is the vector that connects the centres of the colliding nuclei. The angle Ψ_R is then the plane's inclination from the horizontal (x) axis.

The azimuthal particle distribution is then expanded into a Fourier series around Ψ_R [60,61],

$$E \frac{d^3 N}{dp^3} = \frac{1}{2\pi} \frac{d^2 N}{p_{\perp} dp_{\perp} dy} \left(1 + \sum_{n=1}^{\infty} 2v_n \cos(n(\phi - \Psi_R)) \right) . \quad (2.6)$$




Figure 2.8: Event P_E plane spanned by the colliding nuclei.

The sine terms in the Fourier expansion disappear due to the reflection symmetry through the event plane. The coefficients v_n of the harmonics in the Fourier expansion are determined as

$$v_n = \langle \cos(n[\phi - \Psi_R]) \rangle \quad , \quad (2.7)$$

where the average $\langle \rangle$ runs over all particles in all events. When defined this way, the v_1 is called the direct flow, and v_2 is called the elliptical flow [60, 61], and

$$v_1 = \left\langle \frac{p_{x'}}{p_{\perp}} \right\rangle \quad (2.8)$$

$$v_2 = \left\langle \left(\frac{p_{x'}}{p_{\perp}} \right)^2 - \left(\frac{p_{y'}}{p_{\perp}} \right)^2 \right\rangle \quad (2.9)$$

where x' is in the event plane, and y' is perpendicular to the event plane. Geometrically v_1 is the offset from the centre of azimuthal distribution $r(\varphi)$, and v_2 is the ellipticity of the distribution.

The inclination Ψ_R of the event plane to the global coordinate frames x axis is not directly measurable by experiments. Instead, one uses the estimated m^{th} order event plane defined by

$$\Psi_m = \frac{1}{m} \tan^{-1} \left(\frac{\sum_i w_i \sin(m\phi_i)}{\sum_i w_i \cos(m\phi_i)} \right) \quad ,$$

where the sums run over all particles in an event, and w_i is some weight assigned to the i^{th} particle of the event. Any weight can in principle be used, subject to the experimental limitations, but [62] recommends using $w_i = p_{x',i}/p_{\perp,i}^2$. Obviously this requires that the detector can determine the momentum of the i^{th} particle, which is not the case for the FMD, and we are therefore forced to set $w_i = 1$. Note, that one should use an event plane angle Ψ_m such that $n = km$ for $k \in \mathcal{N}$ [61], and that Ψ_m has a periodicity of $2\pi/m$. The angle Ψ_1 is in general the angle at which the azimuthal distribution is offset from the centre, and Ψ_2 is the angle of the minor axis of the corresponding ellipsis⁵.

⁵The general interpretation of the Fourier coefficient v_1 as the strength of the offset, and Ψ_1 as the angle of the offset, of an azimuthal distribution, is used in other applications as well. For example, primitive heat-seeking missiles uses a similar decomposition to determine the best trajectory to the target. More fancy heat-seekers can use the higher orders v_2, \dots and Ψ_2, \dots for advanced navigation. For example, the second coefficient could be used by the seeker to determine just how abruptly the missile needs to turn to hit its target.

The pre-factor $\frac{d^2N}{p_\perp dp_\perp dy}$ in (2.6) allows us to study v_n as a function of transverse momentum, rapidity, or centrality. That is, v_n can be determined from particles in different p_\perp , y , and or centrality bins. This allows experiments to study the flow properties of the heavy-ion collisions in very different regimes, for example in the soft/hard part of the produced particles, in the near baryon-free mid-rapidity region versus the more baryon-rich forward regions, and as a function of the number of binary collisions or participants in the heavy-ion collisions.

In [60] the authors write

It is known that the reconstruction of the reaction plane event-by-event is very ambiguous due to the fluctuations caused by finite multiplicity.

Hence, in [61] the same authors suggest that the observed azimuthal anisotropy v_n^{obs} with respect to the $m = n/k$ (for some $k \in \mathcal{N}$) event plane is corrected by the factor $R_k = \langle \cos(km[\Psi_m - \Psi_R]) \rangle$, such that

$$v_n = \frac{v_n^{obs}}{R_k} = \frac{\langle \cos(n[\varphi - \Psi_m]) \rangle}{\langle \cos(km[\Psi_m - \Psi_R]) \rangle} . \quad (2.10)$$

Since Ψ_R is not directly observable, one cannot evaluate the denominator directly. Instead, considering the distribution of event plane angles $m(\Psi_m - \Psi_R)$ [61],

$$\frac{dP}{dm(\Psi_m - \Psi_R)} = \int \frac{v'_m dv'_m}{2\pi\sigma^2} e^{-\frac{v_m^2 + v'_m{}^2 - 2v_m v'_m \cos(m[\Psi_m - \Psi_R])}{2\sigma^2}} , \quad (2.11)$$

$$(2.12)$$

where $\sigma^2 = \frac{1}{2N} \frac{\langle w^2 \rangle}{\langle w \rangle^2}$ is the standard deviation in the weights from the N observed φ 's. Following [61], the resolution R_k can then be expressed as

$$\begin{aligned} R_k(\chi) &= \langle \cos(km[\Psi_m - \Psi_R]) \rangle \\ &= \frac{\pi}{\chi} e^{-\chi^2/4} \left(I_{\frac{k-1}{2}}(\chi^2/4) + I_{\frac{k+1}{2}}(\chi^2/4) \right) , \end{aligned}$$

where $\chi = \frac{v_m}{\sqrt{2\sigma}}$, and $I_n(x)$ is the modified Bessel function of the first kind⁶. There are different strategies to evaluate this function [61], but usually involves correlating the event plane angle from two or more independent sub-events. See also Section 8.10.1 for a suggestion on how to evaluate the function for FMD data.

2.3.2 Overview of previous v_2 measurements

Figure 2.9 [63] shows previous measurement of v_2 as a function of the charged particle multiplicity $dM_{ch}/d\eta$ at $\eta \approx 0$ and scaled to the initial overlap area A_\perp . In the figure, v_2 has been scaled to participant eccentricity ε to factor out effects from centrality selection and geometric effects. Also shown in Figure 2.9 are two hydrodynamic calculations based on two different equation-of-state (EoS). The full-drawn line comes from a hadron gas EoS and the dashed line from a full QGP EoS. Clearly, v_2 increases as a function of $\sqrt{s_{NN}}$ and $dM_{ch}/d\eta$, with the system size playing a secondary role.

All four RHIC experiments have reported results on the azimuthal anisotropy in Au–Au measured in various ways [40, 49, 50, 64], including as a function of p_\perp , η , centrality (or number of participants), and particle species. A few summary plots is shown in Figure 2.10 and 2.11.

The BRAHMS result for π^\pm v_2 shown in Figure 2.10(a) indicates that the anisotropy is constant over η , though the PHOBOS results for unidentified particles shown in Figure 2.10(b), shows a clear fall-off at higher pseudo-rapidities. However, [65] attributes the discrepancy to relatively

⁶The Bessel functions comes from the integration of (2.11).

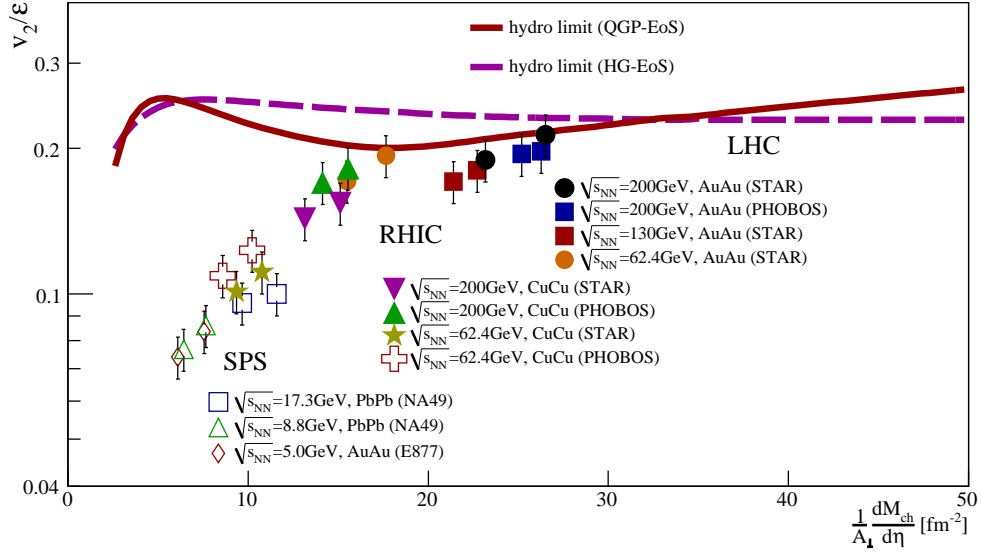


Figure 2.9: v_2/ε as a function overlap normalised $dM_{\text{ch}}/d\eta$. ε is the eccentricity and A_{\perp} is the area of the initial overlap region. The full-drawn line is a hydrodynamic calculation with a hadron-gas equation-of-state (EoS), while the dashed line is a hydrodynamic calculation with a QGP EoS. Adapted from [63].

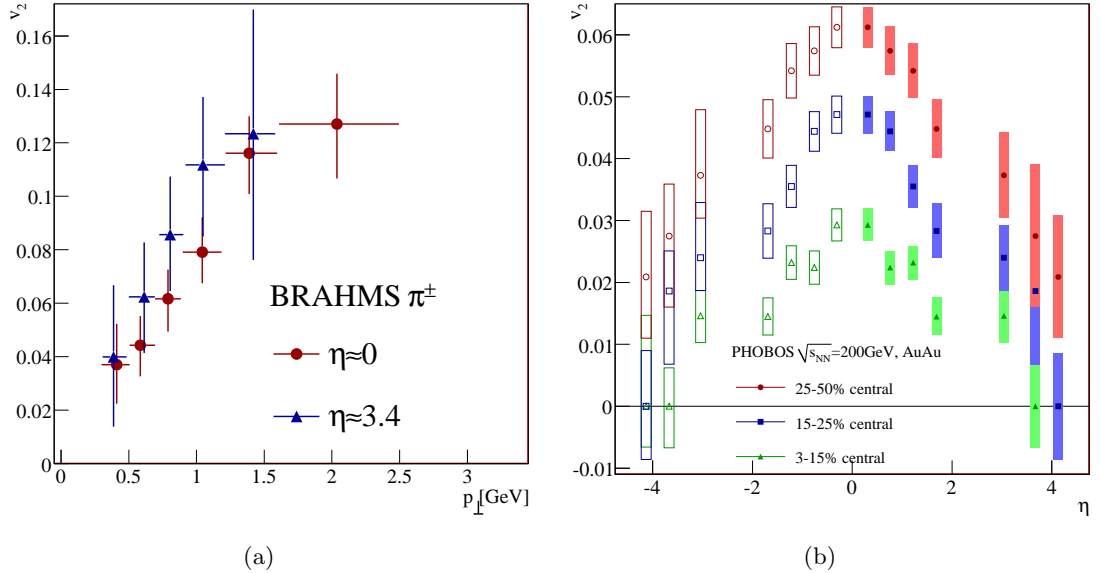


Figure 2.10: (a) BRAHMS measurement of p_{\perp} dependent v_2 of pions in various pseudo-rapidity bins. Adapted from [65]. (b) PHOBOS measurement of M_{ch} dependent v_2 in a broad range of pseudo-rapidity bins at 3 different centralities. Adapted from [66].

large systematic errors when extrapolating to low p_{\perp} . The results does, however, indicate that the azimuthal anisotropy may be very different at forward regions away from the almost baryon-free region (see Figure 1.11 and 1.12). If the almost baryon-free region does not extend up to $\eta = \pm 5$ at the LHC (although some results gives reason to believe that it may [25]), then the FMD could potentially provide significant insight into the flow profile of the collisions.

The centrality dependence seen in Figure 2.10(b) is not entirely surprising. When the two nuclei collide at very low impact parameter b , the interaction zone is almost spherically symmetric around the directions of the beams (the z axis by convention). Hence, there is very little initial spatial anisotropy turn into momentum space anisotropy. For very peripheral collisions, there may not be enough produced particles to build up enough pressure to sustain any initial anisotropy. The strongest flow signal is therefore expected to be found at mid-central to mid-peripheral collisions.

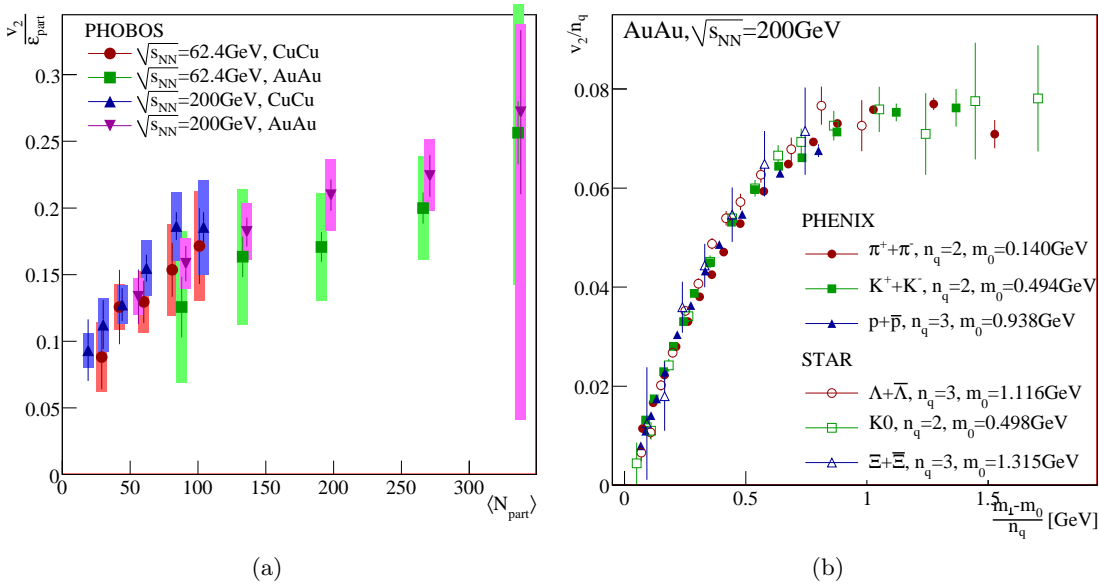


Figure 2.11: (a) PHOBOS measurement of the integrated v_2 in $|\eta| < 1$. Here, v_2 is scaled by the eccentricity of the participants ε_{part} and plotted as a function of the number of participants in an attempt to find scaling properties of v_2 . Adapted from [67]. (b) PHENIX and STAR measurement of p_{\perp} dependent v_2 for various species, plotted as a function of $KE_{\perp} = m_{\perp} - m_0$ (where m_0 is the rest mass of the particle), and scaled by the number of constituent quarks n_q . Again, some scaling seems likely, though the pions seems to break it. Adapted from [68].

Figure 2.11(a) and 2.11(b) represents attempts at finding scaling properties of v_2 . PHOBOS shows the multiplicity v_2 , scaled by the eccentricity of the participants ε_{part} versus the number of participants in the collision. Within error bars, there is certainly a clear scaling from lower energy/less number of participants to higher energies or more participants. The authors of [67] take this to indicate that the geometric properties of the collision plays a significant role in the build up of v_2 .

On the other hand, the clear scaling of v_2 with number of constituent quarks n_q as shown in Figure 2.11(b) tells us that more than geometry is at play. Although the scaling is not perfect, in particular the higher $KE_{\perp} = m_{\perp} - m_0$ pions seem to not follow the general trend, it is striking. The authors of [68] interprets this result as an indication that quark-like degrees-of-freedom are dominant in the initial stages, where the v_2 is build up, of the collision. That the scaling breaks down for pions over $KE_{\perp} \sim 1$ GeV is indicative of the hydrodynamic evolution of the

system, which is only applicable for low momentum particles. For higher momentum particles, the hydrodynamic description breaks down since these particles are not equilibrated (even locally) with the rest of the system [43].

2.3.3 Lessons to be learned from the v_2 measurements

Assuming that hydrodynamics is a valid description of the bulk of the matter created in the heavy ion collisions, one can calculate limits for how large the azimuthal anisotropy can be, as indicated in Figure 2.9. The equation-of-state used for a QGP assumes an ideal fluid — that is no viscosity and local thermal equilibration [43]. The observed azimuthal anisotropy sets a limit for how far back in time the hydrodynamic description is valid, as the hydrodynamic expansion of the fluid under the internal pressure will generally suppress the flow. Calculations find that almost all of the anisotropy observed at RHIC is formed in the first 3-4 fm after the collision, corresponding to a very early equilibration time of 0.6 fm [43].

The early equilibration time, and corresponding high pressure, together with the clear constituent quark scaling observed (Figure 2.11(b)), indicates that the matter formed at RHIC has colour degrees-of-freedom, and that particles created interact strongly with coloured matter. This is seen as supporting the idea of a new state of matter, the sQGP — a *strongly* interaction perfect fluid of deconfined quarks and gluons.

Chapter 3

The LHC and the ALICE Experiment

3.1 The Large Hadron Collider

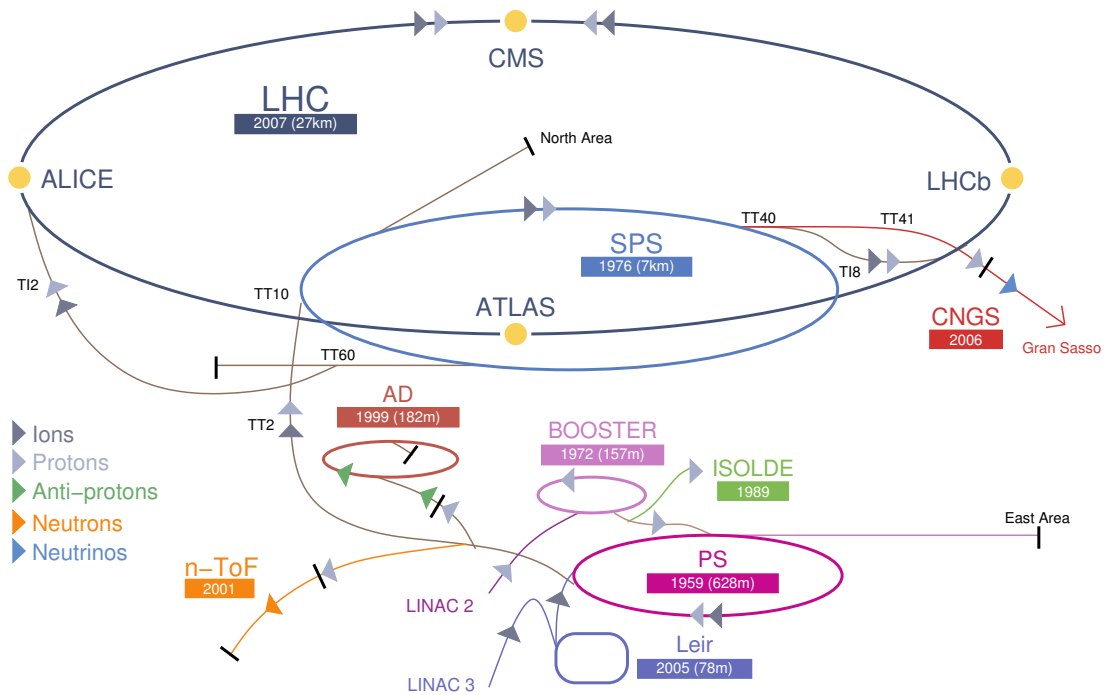


Figure 3.1: The CERN accelerator complex. Adapted from [69]. The figure is not to scale.

The Large Hadron Collider is the next CERN collider. It is 27 km in circumference (≈ 4.3 km in radius) and is equipped with some 1200 super-conducting magnets. The LHC can accelerate protons to $\sqrt{s} = 14000$ GeV and lead ions to $\sqrt{s_{NN}} = 5500$ GeV.

Protons are produced at LINAC 2 and boosted in the BOOSTER. They are then sent into the PS and SPS for further acceleration up to $\gamma = 479.6$. From the SPS the proton beam is split into 2808 bunches going in either direction along the LHC main ring [70]. The bunches of protons are further accelerated to $\gamma = 7461$ and brought to collisions at 4 points along the ring, coinciding with the location of the 4 experiments ALICE, ATLAS, CMS, and LHCb. In proton runs, there will be bunch-crossings every 25 ns (or at 40 MHz). Table 3.1 summarises the kinetic energy of the protons at different stages in the acceleration.

$^{208}\text{Pb}^{+27}$ are produced at LINAC 3 and sent into LEIR. The ions are then further accelerated in the PS and SPS, until the ion beam is split into roughly 592 bunches sent either way round the LHC ring with $\gamma = 190.5$. The ions are further accelerated in the main LHC ring to $\gamma = 2963.5$,

and brought to collisions three places along the ring, coinciding with the location of the ALICE, ATLAS, and CMS experiments. The time between bunches will be 100 ns. Table 3.1 summarises the kinetic energy and charge of ^{208}Pb ions at different stages in the acceleration.

Step		LINAC2/ LINAC3	BOOSTER/ LIER	Ps	SPS	LHC
proton	E_{kin}	50 MeV	1.4 GeV	25 GeV	450 GeV	7 TeV
^{208}Pb	E_{kin}	4.2 MeV/u	72.2 MeV/u	5.9 GeV/u	176.4 GeV/u	2.76 TeV/u
	Charge	+27	+54	+82	+82	+82

Table 3.1: Kinetic energy at different stages in the beam acceleration for proton and ^{208}Pb [71]. Where there are two rows for the accelerator step, the upper refers to the proton acceleration, and the lower to the Pb acceleration.

3.2 Overview of ALICE




Figure 3.2: Section through ALICE. Arrows indicate positions of the various detectors.

ALICE is the only dedicated heavy-ion experiment at the LHC, although ATLAS and CMS will have smaller heavy-ion programs too. Although ALICE is dedicated to heavy-ion physics it will take data during the pp runs, as these measurements will provide a baseline for many interesting heavy ion physics measurements.

Dominating the ALICE cavern is the huge L3 magnet — the worlds largest conventional magnet. It is inherited from the old LEP experiment L3. It can provide a solenoidal¹ field of up to 1/2 T.

ALICE consists of three sections: The barrel detectors, including the ITS, TPC, TRD, TOF, HMPID, EMCAL, and PHOS; the forward detectors, including FMD, T0, V0, and PMD; and

¹I.e., parallel to the centre line.

last but not least, the MUON spectrometer. Below, each sector and detector will be described briefly.

3.2.1 Barrel Detectors




Figure 3.3: Section through the magnet and barrel detectors.

The barrel detectors foremost role, is to measure the momentum and identity of particles produced in the region $|\eta| < 0.9$. The 7 detectors in the barrel sector contributes to these measurements by use of diverse technologies. Figure 3.3 shows a section through the L3 magnet and the barrel detectors.

ITS The *Inner Tracking System* is a high-resolution silicon detector. With its inner radius of 4 cm, it is the detector system closest to the interaction point. It consists of 3 sub-detectors, starting from the centre and going outwards: The silicon pixel detector (SPD), the silicon drift detector (SDD), and the silicon strip detector (SSD). Each of these 3 sub-detectors have 2 layers.

The SPD is a pixel detector — that is, the active elements are small pixels on the face of a silicon sensor. It has a resolution in the $r\varphi$ -plane of $12\ \mu\text{m}$ and $70\ \mu\text{m}$ in the z direction. With its expected occupancy of 0.4% to 1.5% it is a formidable charged particle multiplicity detector in the region $|\eta| < 2.1$ (see also Figure 4.2). Furthermore, by combining all possible hits in the SPD one can get a rough estimate of the position of the primary interaction. The FMD relies on precise knowledge of the z coordinate of the interaction (v_z) to determining the η of the measured multiplicities.

The other two layers of the ITS, the SDD and SSD, have slightly less granularity than the SPD. They provide further tracking points and charged particle multiplicity measurements.

The ITS as a whole provides tracking of charged particles near the interaction point. Due to its fine granularity, it can resolve decays of short-lived particles (such as Λ_s , and Ξ_s) and determine

the point of decay². By forming tracklets in all 6 layers of the three sub-detectors one can further improve the determination of the position of the primary interaction. This information is used to restrict the global tracking of particles in the barrel detectors: Tracks that do not seem to originate relatively close to the interaction point can be discarded as background tracks from cosmic rays, scatterings in material, or other such sources.

TPC The *Time Projection Chamber* is by far the largest single detector in ALICE. It is a 80 m³ gaseous detector providing tracking and particle identification in the region $|\eta| < 0.9$. When charged particles traverse the gas volume, they ionise the gas and electrons drift toward the read-out planes on either end-cap.

The end-caps are instrumented with several thousand front-end cards with the custom built ALTR0 Analog-to-Digital-Converter (ADC) chip. This highly advanced chip has been re-used by several sub-detectors in ALICE, including the FMD, PHOS, and EMCAL. The front-end cards are grouped together in partitions and each partition is controlled by a Read-out Controller Unit (RCU) mother card. Data is transferred from the RCU directly to the data acquisition system via optical fibres — there are 280 of these in total.

The full drift-time of the TPC is roughly 88 μ s which sets the trigger scale of ALICE — that is, within a window of 88 μ s starting from a collision, there should be no other collisions. If there were collisions within the drift time, it would pollute the event being read-out. Unlike ATLAS and CMS where each read-out can be tagged with a time-stamp, the ALICE TPC cannot resolve particles from multiple interactions. The maximum trigger rate of ALICE will therefore be around 10 kHz.

Particle identification in the TPC is done by using the energy loss of particles in the gas. The information on the energy loss is picked up during the tracking and can be used to adjust the assumptions of the Kalman filter used for tracking (see below).

TRD The *Transition Radiation Detector* is located just outside the main TPC barrel (see Figure 3.3). The TRD works from the principle of transition radiation i.e., that a charged particle that crosses over the boundary between two materials with different dielectric constants will emit radiation. The amount of radiation emitted depends on the particle's γ factor in the rest frame of the materials, and the TRD can therefore be used for particle identification.

Practically, the TRD consists of layers of plastic, inter-spaced with wire chamber layers. The wire chamber layers detect the radiation emitted in the transition with a resolution of 400 μ m in the $r\varphi$ plane, and 2.3 cm in the z direction [72]. The holes in the φ coverage are there to provide open windows with little material for the PHOS and HMPID.

TOF The *Time Of Flight* detector is placed outside of the TRD (see Figure 3.3). It provides a measurement of the time it takes a particle to travel from the interaction point, through the magnetic field, to the outer rim of the barrel. Since the particle momentum p and trajectory length l is known from tracking in the ITS, TPC, and TRD, one can deduce the particle mass from the time of flight measurement t , by

$$m = p \sqrt{\frac{t^2}{l^2} - 1} \quad .$$

The TOF covers basically the same η region as the TRD. Again, there are holes in the φ coverage to minimise the amount of material in front of the PHOS and HMPID.

The *Time Of Flight* is built of Multigap Resistive Plate Chambers. In such a detector, resistive plates are put perpendicular to the particle trajectories, and spaced with gas volumes. A high voltage is applied over the full chamber so when a charged particle traverse the gas volume,

²So-called V0s.

it loses energy through ionisation and creates an avalanche. The signal from the avalanche is then detected at the anode of the detector [73–75]. This design gives a timing resolution of some 120 ps [76].

HMPID The *High Momentum Particle Identification Detector*, has, as the name indicates, as its primary role, to identify the particle type of very high momentum particles. The π/K separation goes up to 3 GeV, K/p separation up to 5 GeV.

The HMPID exploits the phenomena that charged particles emit Cherenkov radiation when the velocity of the particle is larger than the speed of light in the medium traversed $v > c/n$ (n is the index of refraction of the medium). The Cherenkov light is emitted as a shock wave at an angle to the particle track that depends on the speed of the particle. Thus, if one measures opening angle (or equivalent, the radius) of the shock wave, one can determine the velocity of the particle, relative to the velocity of the light in the medium. $\cos \vartheta_C = 1/\beta n$ [77].

Practically, the HMPID consists of a layer of radiator material (C_6F_{14}) and a Multi-Wire Proportionality Chamber (MWPC) behind to detect the radiated Cherenkov light. The light emitted will hit the MWPC and form a ring on the read-out plane. This ring is then identified and associated to a known track (from, say the TPC) via complicated algorithms, and thus the particle type can be determined.

EMCAL The *Electro-Magnetic CALorimeter* covers a large φ segment, and roughly the same η range as the TPC. Its job is to measure the transverse energy of (E_\perp), and hence p_\perp , of particles that hit the detector. The EMCAL provides p_\perp measurements in the region from 100 MeV to 100 GeV, making it an excellent detector to study jets with. In fact, the read-out of the EMCAL is fast enough to participate in the LVL1 trigger decision, and can therefore provide ALICE with a jet-trigger [78].

The calorimeter is made of lead-glass rods placed so that they point toward the nominal interaction point. Light created by traversing charged particles is collected in fibres and sent to a photo-chip for collection.

The EMCAL awaits final approval from the United States of Americas Department of Energy, and will therefore not be installed for the first few running periods of the LHC.

PHOS The *PHO*ton *S*pectrometer is a specialised detector in ALICE. It is an electro-magnetic calorimeter of lead-tungsten crystals, and will measure photons, π^0 , and η mesons up to a p_\perp of 10 GeV. These measurements can be used to do jet physics, direct γ measurements of initial temperature, and to look for signatures of chiral symmetry restoration [79].

Tracking in the Barrel Tracking in the barrel is done by seeding a Kalman filter³ with tracklets in the outer periphery of the TPC cylinder. The filter then tries to predict the next point in the volume from the prior knowledge (including magnetic field and spatial distortions) and checks if there is any hit consistent with a list of assumptions. The filter progresses inward in the detector, constantly adjusting the assumptions, until it hits the inner cylinder. From there, the ITS takes over and another Kalman filter tracks through that detector. Finally, the Kalman filter is applied in reverse, tracking from the inner most of the ITS to the outer cylinder of the TPC. These tracks are then used to seed another Kalman filter that tracks through the TRD.

3.2.2 Forward Detectors

Figure 3.4 shows the 3 forward detectors: T0, V0, and the FMD, along with the ITS. The forward detectors serve to provide over-all event characteristics, like triggering, primary vertex,

³A Kalman filter is an algorithm first used for landing aeroplanes. As stated earlier, Fourier decompositions are used in missile trackers, so in a sense the ALICE detector can be said to do everything in its power to keep aeroplanes on the ground.

and multiplicity.




Figure 3.4: The forward detectors and the ITS.

T0 The T0 detector is a high-resolution timing detector. It consists of Cherenkov radiators glued unto photo-multiplier tubes. The time resolution of the T0 is of the order of 10 ps. A coincidence between the two sides T0-A and T0-C will serve as a LVL0 trigger and early wake-up signal to other detectors.

V0 In pp where the density of charged particles is much lower than in A-A, the T0 does not have a large enough acceptance to provide a reliable LVL0 trigger. The V0 was therefore designed to have a larger acceptance to provide the first trigger in pp. The V0 can also measure the charged particle density in the rapidity $-3.6 \leq \eta \leq 1.6$ and $2.8 \leq \eta \leq 5.1$, though at a lower precision than offered by the FMD and ITS.

Furthermore, the V0-C array can be used as validation for the MUON spectrometer's separate trigger chambers.

FMD The primary role of the *F*orward *M*ultiplicity *D*etector is to measure the number of charged particles at forward (small) angles relative to the beam line in fine η and φ bins. Chapter 4 will explain the FMD in detail.

3.2.3 Muon Spectrometer

Beside the barrel detectors for tracking and particle identification, and the forward detectors for global event characterisation, there is another special purpose group of the detectors in ALICE — the MUON spectrometer. The purpose of the spectrometer is to measure dileptons, and extract all the possible physics from those measurements, including J/Ψ suppression, ρ mass broadening, and so on.

Figure 3.5 shows the MUON spectrometer. Closest to the interaction point is the front absorber. This cone-like structure is a compound of concrete, lead, and other metals, to stop and absorb all particles except for μ^\pm . The front absorber works as a filter, so that the most likely particles to observe in the rest of the spectrometer are μ^\pm . Next, just behind the absorber nose, are tracking station 1 and 2. Tracking station 1 is within the L3 volume, while station 2 is flush with the edge of the L3 magnetic field. This allows the MUON spectrometer to precisely determine where the particles left the solenoidal field. Dominating the spectrometer is the large dipole magnet, which bends the trajectory of charged particles in the yz plane. The third tracking

station is in the middle of the dipole to allow precise measurements of the angle of deflection. Tracking station 4 and 5 sit further back, on either side of another μ filter. All tracking stations in the MUON spectrometer are cathode plane detectors. Finally, behind tracking station 5 is the trigger chambers for measuring the time-of-flight of the μ particles to allow identification. The trigger chambers are resistive plate chambers, as in the TOF.




Figure 3.5: The muon spectrometer.

Part II

The ALICE Forward Multiplicity Detector

Chapter 4

Overview and Design Considerations

4.1 Overview

The ALICE FMD is a relatively finely segmented silicon strip detector. It consists of 3 sub-detectors: FMD1, FMD2, and FMD3, and each of these has 1 or 2 rings of sensors arranged around the beam-line, as shown on Figure 4.1. FMD1 consists of one ring, FMD1i, while FMD2 and 3 consist of two rings each, FMD2i, FMD2o, FMD3i, and FMD3o.

The FMD was designed at the Niels Bohr Institute at the University of Copenhagen, and culminated in the release of the *Forward Detectors: FMD, T0, and V0* technical design report [2] in the Autumn of 2004.




Figure 4.1: The forward detectors and the ITS.

4.2 Design Considerations

The primary motivation for building the FMD is to increase the charged particle coverage of ALICE into the forward regions as shown in Figure 4.2. The reasons for this, as outlined in Section 2.2, is to get a more complete view of the particle production and investigate the onset of the fragmentation region.

Another motivation is to provide independent measurements of the angle of the impact parameter \mathbf{b} to the global zx plane, and independent measurements of the azimuthal asymmetry v_2 away from the mid-rapidity region, as detailed in Section 2.3.

These goals for the FMD lead to the following requirements of the FMD design:

1. Broad coverage in η .

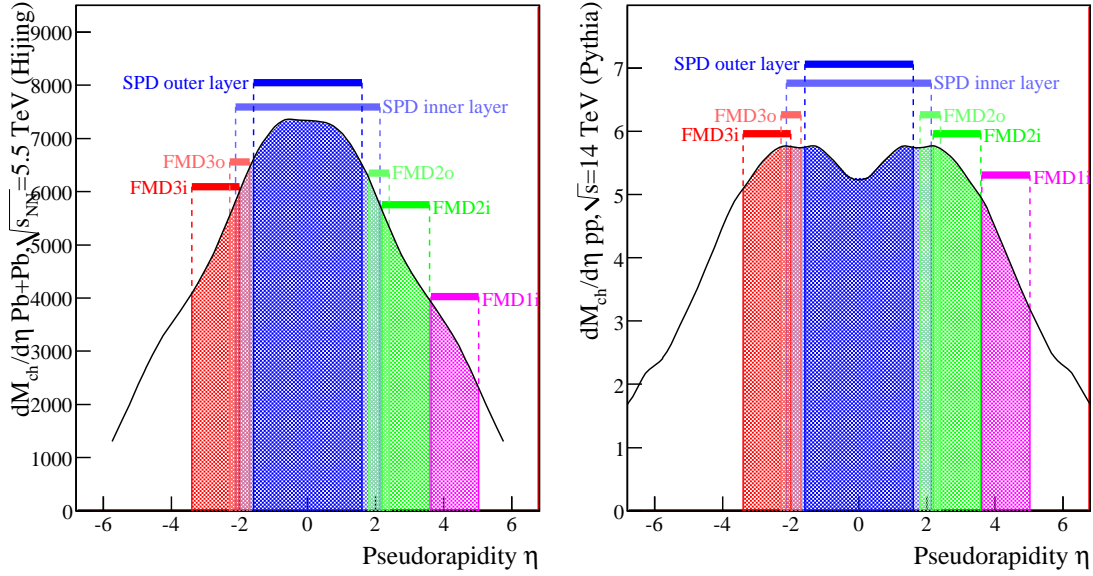


Figure 4.2: Total charged particle multiplicity coverage of ALICE, broken up by detector system. The data is from a HIJING simulation of central Pb–Pb at $\sqrt{s_{NN}} = 5500$ GeV and PYTHIA simulation of pp at $\sqrt{s} = 14000$ GeV.

2. High resolution in the measurement of charged particle multiplicity, implying a low average number of particles will hit each detection element.
3. Fine η segmentation.
4. Full φ coverage with a modest segmentation.
5. High radiation tolerance.

A tracking detector at high $|\eta|$ with a good momentum resolution would require a very long base–line¹, since particles emitted in these regions have very high momentum, and are thus deflected very little in the magnetic field. The FMD was therefore not designed with tracking capabilities in mind.

These considerations lead to a design based on silicon semi–conductor sensors with a fine radial segmentation, organised in rings around the beam–line.

The overall placement and parameters of the 3 sub–detectors are listed in Table 4.1.

Sub–detector	Ring	# φ segments	# r segments	z [cm]	r_{in} [cm]	r_{out} [cm]
FMD1	i	20	512	320.0	4.2	17.2
FMD2	i	20	512	83.4	4.2	17.2
FMD3	o	40	256	75.2	15.4	28.4
	i	20	512	-62.8	4.2	17.2

Table 4.1: Overview of parameters the FMD sub–detectors.

¹Distance to the interaction point in z .

Chapter 5

Sensors and Electronics

This chapter will describe the physical parts of the FMD, starting from the sensors that register the charged particles produced in the collision, and following the stream of the data through the *Front-End Electronics (FEE)*, all the way up the data acquisition system. The electronic system involved is shown schematically in Figure 5.1¹

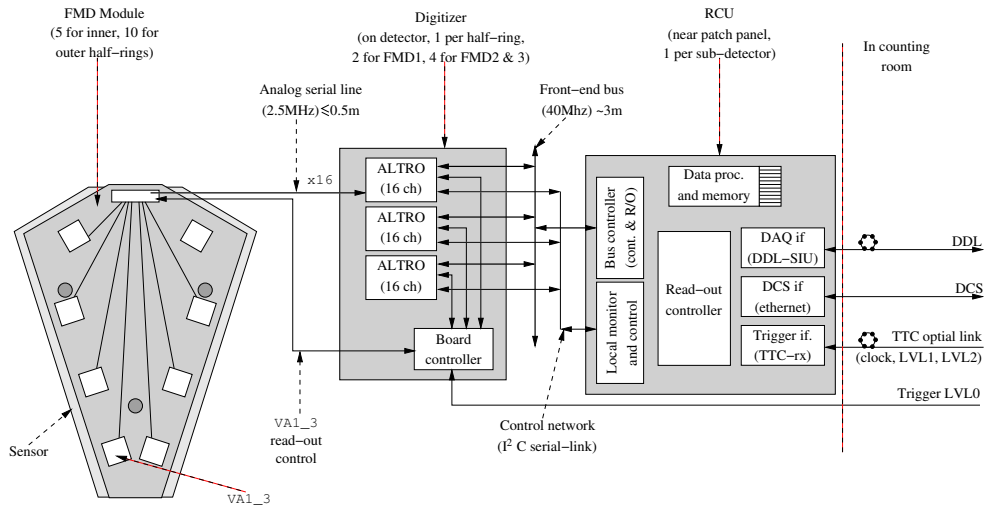


Figure 5.1: The ALICE FMD electronics chain. From the left, the sensor-hybrid module, the digitiser card, and the read-out controller unit.

Section 5.1 will describe the silicon sensors, and Section 5.2 outlines the basic functionality of the hybrid card, with emphasis on the VA1₃ pre-amplifier. In Section 5.3 FMD digitiser (FMDD) card is described, emphasising the board controller and ALTRO analogue-to-digital converter chip. Finally, in Section 5.4 the read-out controller unit (RCU) will be briefly explained, along with its two daughter cards, the DCS card and the SIU.

5.1 The FMD Sensors

The basic building blocks of the FMD are two kinds of silicon sensor, the *inner* and *outer* types, as shown in Figure 5.2. Both types are cut from 300 μm thin wafers with a diameter of 6”.

The sensor are arranged in two rings, as shown in Figure 5.3. Every second sensor is staggered 8 mm in z to allow room for the bonding pads, and at the same time have full 2π acceptance in azimuth. At the time of the design of the FMD 6” silicon wafers was the biggest possible to manufacture at the high quality needed for a detector. That means that some corners of the

¹In Figure F.1 is a photo of a sensor, and Figure F.2 shows a photo of the full front-end electronics chain.

silicon sensors had to be sacrificed to fit the sensor into the wafer. This in turn means an $< 2\pi$ azimuth acceptance for the most peripheral strips, as can be seen on Figure 5.3.




Figure 5.2: Inner (left) and outer (right) sensor types cut from 6" silicon wafers. Each sensor is divided azimuthally into 2 sectors, and radially divided into 512 strips for the inner type, and 256 strips for the outer types, as suggested on the figure. The insensitive area (dark gray) along the edges are used for bonding pads, and guard rings.




Figure 5.3: Arrangement of inner (left) and outer (right) sensor into rings. There are 10 inner sensors (20 sectors) to a ring, or 20 outer sensors (40 sectors) to a ring.

5.1.1 Semi-conductor detectors

Silicon is a *semi-conductor* material. This means that the band-gap between the *valence* and *conduction* band is very small, and it therefore takes little energy to excite an electron from the valence band into the conduction band, leaving a *hole* in the valence band. This is in contrast to insulators and metals, where in the former it takes comparably much more energy to excite an

electron into the conduction band, and in the latter where it takes no energy since electrons can readily dissociate from the parent nucleus and roam the crystal freely (see Figure 5.4).




Figure 5.4: Valence and Conduction band gap in insulators, semi-conductors, and metals. Adapted from [77].

In a silicon crystal the electron bonds are covalent as shown in Figure 5.5(a). Because of the small band-gap, even small thermal excitations are enough to bring electrons from the valence band into the conduction band as shown in Figure 5.5(b). To increase the number of charge carriers (electrons *and* holes) one can create a so-called *doped* crystal. This is done by replacing a small number of the silicon atoms with some pentavalent² atoms. The extra electron (see Figure 5.5(c)) will introduce an allowed energy state in the band-gap around 0.05 eV below the conduction band [77]. The extra electron can easily be excited into this state, and thereby increase the conductivity of the crystal.

Doped crystals Such a crystal, where pentavalent atoms (called *donors*) have replaced some of the silicon atoms, is called a doped *n-type* crystal. One can of course also replace silicon atoms with trivalent³ atoms, forming a *p-type* crystal. *n-type* crystals have extra holes and an allowed energy state in the band-gap near the valence band into which electrons can easily be excited, leaving an excess of holes in the crystal again, increasing the conductivity but with the holes being the majority carries.

Generally, the concentration of donors is very small ($\propto 10^{-9}$). Increasing the concentration of donors dramatically in an n-type crystal forms a so-called *heavily* doped crystal, and is normally denoted n+. Similarly a heavily doped p type crystal is labelled p+. Heavily doped crystals are very conductive, and is therefore often used for electrical contacts to read-out electronics.

Silicon crystals as detectors The small band-gap in semi-conductors makes them excellent detectors for ionising radiation. When a charged particle traverses the crystal, it will kick electrons from the valence band into the conduction band, leaving a trail of electron-hole pairs behind it. If an external electric field is applied over the crystal, then the excited electrons will drift towards the anode, while the holes left behind will drift towards the cathode, thereby setting up a current over the crystal.

However, more is needed to make an effective sensor. Since thermal excitations can excite electrons into the conduction band, a pure n or p type crystal would be very noisy. For detectors, one exploits a semiconductor junction to set a very sensitive low noise sensors.

Semiconductor junction and depletion depth The idea is to make a junction between an n-type and a p-type crystal⁴. Since there is a higher concentration of electrons in the n-type, and a higher concentration of holes in the p-type (see Figure 5.5), there will be a diffusion of electrons from the n side to the p side, and likewise a diffusion of holes from the p side to the n side of the junction. Since both sides of the junction are initially neutral, this diffusion will charge the n side positively and the p side negatively, thus creating an electric field over the junction. This

²5 electrons in the valence band.

³3 valence electrons

⁴Making the junction is an art form in it self, and will not be discussed here.




Figure 5.5: Bindings in the silicon crystal. (a) Covalent bonds in a Si crystal at low temperatures. (b) At non-zero temperatures, electrons are excited by thermal energy and pushed up into the conduction band, leaving a hole in the covalent bonds. (c) Adding a pentavalent impurity to the crystal frees an electron from the covalent bond, introducing allowed states in the gap between the valence and conduction band.

contact potential will at some point bring the diffusion to a stop, since the electrons and holes have to perform work to overcome the electric field, and the remaining charge carriers are left immobile in what is known as the *depletion* region around the boundary of the junction. Any electron-hole pairs created in the depletion region will immediately be swept out — the electron to the p side and the hole to the n side.

The depth of the depletion zone is for n-type silicon is given by [77]

$$d \approx 0.53 \sqrt{\rho_n V_0} \mu\text{m} \quad (5.1)$$

where ρ_n is the crystal resistivity, and V_0 is the contact potential. Taking $V_0 = 1$ V, and $\rho_n \approx 5000 \Omega\text{cm}$ (see Table D.1) gives a depletion depth of $\approx 35 \mu\text{m}$ — about 2500 times smaller than a radiation length in silicon. This small depth makes it less likely that an impinging radiation will create any substantial number of electron-hole pairs. Furthermore, since capacitance of the sensor, is given by $C = \epsilon \frac{A}{d}$ where ϵ is the dielectric constant of silicon, A the crystal area, and d the depth of the depletion region, and that the noise of the sensor is proportional to the capacitance of the crystal, it is desirable to increase the depletion depth.

The depletion depth can be extended up to some $100 \mu\text{m}$ in depth by applying a *reverse bias* voltage. That is, a negative voltage applied to the p side of the junction. This will drag the excess holes on the p side and excess electrons in n side further away from the junction, thus increasing the depletion depth. The depletion depth is still given by (5.1) except that the reverse bias voltage V_b should be substituted for V_0 . The maximum voltage, and hence maximum depletion depth, is determined by the resistivity of the crystal — at some point the voltage will be high enough to overcome the resistance, and the junction breaks down. With a bias voltage of 70 V, we get a depletion depth of roughly $300 \mu\text{m}$, and a corresponding drop of almost an order of magnitude in the capacitance.

The capacitance of the sensor falls off until full depletion depth is reached. When this happens depends on the purity of the crystal, the geometry of the sensor, and how doped the n and p side is.

Leakage current A silicon crystal with a semiconductor junction is ideally not conducting. However, due to various reasons, a small current may nevertheless flow across the sensor when a reverse bias voltage is applied. This is called the *leakage current*, and contributes to the noise of the detector. There are various sources of the leakage current [77].

- Holes in the n side attracted to the p side, or electrons on the p side attracted across the junction to the n side. This generally is a very small contribution on the order of a few nA/cm^2 .
- Thermally created charge carriers in the n and p region that leak into the depletion zone. These are the result of impurities in the silicon that may introduce so-called trapping and recombination centres. This contribution is generally of a few $\mu\text{A}/\text{cm}^2$.
- The largest contribution comes from surface effects, and depends on geometry, contaminants, atmosphere, to name but a few. These effects distort the electric field, creates current conduits, and so on. The leakage current from these is very complex and can not be calculated.

Sensor geometry Figure 5.6 shows a schematic cross section through a silicon sensor. The sensor is n-type silicon with p+ type implants to form the semiconductor junction. On the side facing the interaction point, there is a thin layer of aluminium, which provides the high-voltage connection plane. A bias ring running all along the edge of the detector, coupled to each strip via a $20\text{ M}\Omega$ bias resistor provides the other side of the electric field.




Figure 5.6: Cross-section through a silicon sensor. Note that the left and side is a section along the direction of the strips (a.c.-pad).

Running along the entire periphery of the sensor is a floating guard ring to ensure homogeneous electric field across the bulk volume, so that all electrons are collected on the strips (see for example [80]). Direct-current pads (d.c. Figure 5.6), in direct contact with p+ implant that make up the strips, provides a testing point for the silicon. The analogue current pads (a.c.-pads on Figure 5.6) are linked via an ohmic coupling through the top level SiO_2 layer. The signal collected by the sensor is picked up through micro-bonds to the a.c.-pads.

Test results from Hamamatsu The manufacturer of the silicon sensors, Hamamatsu, performed a number of tests on the sensors before they were shipped to Copenhagen. Figure 5.7 shows the capacitance (left axis) and leakage current (right axis) as a function of bias voltage for arbitrary inner and outer sensors (the full set of curves can be found in Figure D.1 and Figure D.2 in the appendix).

The figure shows that the full depletion is reached around 70 V for inner detectors, and 120 V for outer detectors, as indicated by the vertical dotted lines. Since the leakage current continues




Figure 5.7: Capacitance (left axis) and leakage current (right axis) as a function of bias voltage for arbitrary inner and outer sensors.

to rise with increasing bias voltage, the optimal value of the bias voltage is just after full depletion is reached.

Figure 5.8(a) shows the measured depth of the depletion zone. For both kinds of detectors, the depletion zone is $325 \pm 3 \mu\text{m}$ deep. Figure 5.8(b) shows the operating stability of the sensors. The leakage current is measured over 170 hours, and shows very little variation.

The test from Hamamatsu shows that the raw material of the FMD — the silicon sensors are of very high quality. Independent tests of the sensors have shown that there are $< 1\%$ faulty strips in the sensors.

5.1.2 Energy loss in Silicon

The primary mechanism through which particles are detected in silicon sensors is through the production of electron–hole pairs as outlined in Section 5.1.1. When a particle traverses the sensor, electron–hole pairs are created when the particle lose energy to excitation of the crystal.

Generally, the average energy loss per length (or stopping power) is described by the Bethe-Bloch equation

$$-\frac{dE}{dx} = 4\pi N_A r_e^2 m_e c^2 z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\log \left(\frac{2m_e c^2 \beta^2 T_{max}}{I^2} \right) - \beta^2 - \left\{ \frac{\delta(\beta\gamma)}{2} + \frac{C}{Z} \right\} \right] , \quad (5.2)$$

where

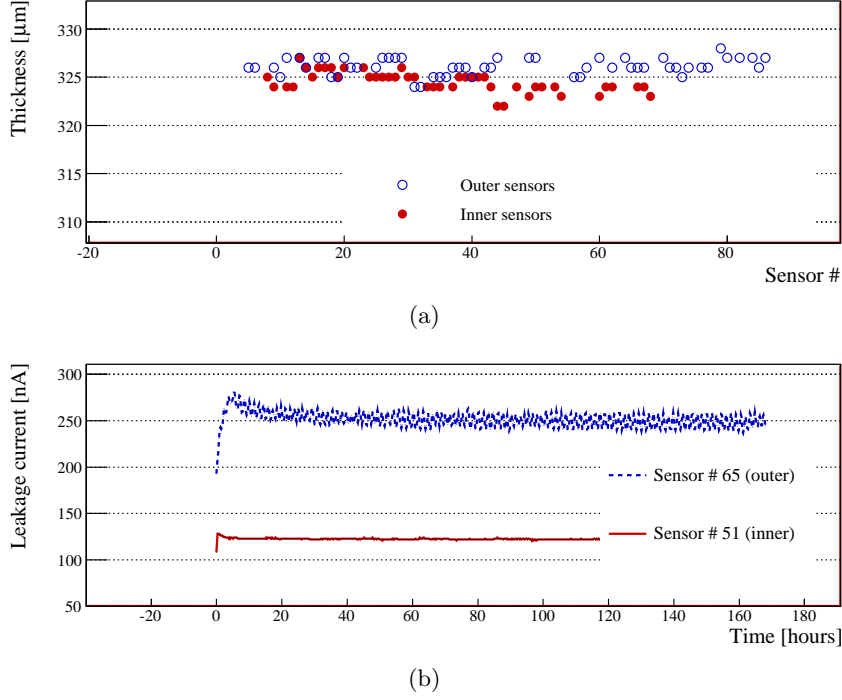


Figure 5.8: (a) Thickness of the depletion zone of all sensors. (b) Stability of the leakage current at optimal operating bias voltage for two arbitrary sensors.

N_A	Avogadro's number	$6.0221415(10) \times 10^{23} \text{ mol}^{-1}$
r_e	Classic electron radius	$2.817940325(28) \text{ fm}$
m_e	Electron mass	$0.510998918(44) \text{ MeV}$
c	Speed of light	
z	Charge of incident particle	e
Z	Atomic number of material	Si: 14
A	Atomic mass of material	Si: 28.0855
β	$= p/m$ of incident particle	
γ	$= 1/\sqrt{1 - \beta^2}$ of incident particle	
T_{max}	Maximum energy transfer in a single collision	MeV
I	Mean excitation energy	eV
$\delta(\beta\gamma)$	Density effect correction to ionisation energy loss	
C	Shell correction	

T_{max} is the maximum energy transferred in a single interaction between the incident particle and the electrons in the material. It is given by

$$T_{max} = \frac{2m_e c^2 \beta^2 \gamma^2}{1 + 2\gamma m_e/m + (m_2/m)^2} \xrightarrow{2\gamma m_e/m \ll 1} 2m_e c^2 \beta^2 \gamma^2$$

where m is the mass of the incident particle. The mean excitation energy I is determined from experiments, and vary greatly with Z [81], but the overall behaviour can be parameterised as

$$\frac{I}{Z} = \begin{cases} 12 + \frac{7}{Z} & Z < 13 \\ 9.76 + 58.8Z^{-1.19} & Z \geq 13 \end{cases},$$

in units of eV

The equation (5.2) (except the last term in the square brackets) was derived by Bethe and Bloch from a quantum mechanical description of charged particles scattering softly with the

electrons in the matter. The term $\frac{\delta(\beta\gamma)}{2}$ stems from density effects. As the momentum of the incident particle becomes larger, then the long range interactions increase as $\log \beta^2\gamma^2$. However, at very incident particle momentum, the material becomes polarised and the increase is suppressed to order $\log \beta\gamma$. The shell correction term $\frac{C}{Z}$ arises when the velocity of the incident particle is comparable to the orbital velocity of the electrons in the absorber. In this regime, the assumption done by Bethe and Bloch that the electrons are at rest relative to the incident particle is no longer valid, and an additional correction is needed. The shell correction is however only valid for very small energies.

Figure 5.9 shows the stopping power of μ in copper as a function of momentum and $\beta\gamma = p/m = \beta/\sqrt{1-\beta^2}$. The Bethe–Bloch equation is in general only a function of $\beta\gamma$ of the incident particle, except for incident electrons and light materials, such as He. The Bethe–Bloch description breaks down for $\beta\gamma < 0.1$, and other models are normally used in that regime.




Figure 5.9: Bethe–Bloch equation evaluated for μ on copper

At very high energies, the incident particle may lose energy through radiative effects, such as bremsstrahlung and Cherenkov radiation. In that region, the Bethe–Bloch description breaks down, and other models must be used. However, the radiative energy loss is only valid for thicker absorbers⁵, since thin absorbers do not provide enough material to induce significant radiative effects.

At higher energies, the incident electron may transfer enough energy to the material electrons that these electrons may create secondary interactions or, for thin absorbers, escape the absorber completely. These electrons are called δ -electrons and lead to a lower average energy loss at higher momentum, as shown on Figure 5.9 by the dash-dotted line.

Minimum Ionising Particles The Bethe–Bloch equation has a minimum at $\beta\gamma_{MIP} \approx 3.5$, as can be seen in Figure 5.9. A particle incident with a $\beta\gamma$ of this value is said to be a *minimum ionising* particle (or MIP)⁶

Energy Loss Straggling Since the process of energy loss is a statistical process, we cannot expect that a particle with given $\beta\gamma$ will always lose the amount of energy prescribed by the Bethe–Bloch equation. What we can expect to see is a *distribution* around the a mean as given by the Bethe–Bloch equation. This is known as energy loss *straggling*.

⁵Thicker means of thickness comparable to the mean interaction length of particles in the material.

⁶Sometimes, the term MIP is also used for particles with $\beta\gamma \geq \beta\gamma_{MIP}$. For some materials, like Silicon, these particles are almost minimum ionising and the terminology appropriate.

For thick absorbers, where the incident particle suffers many collisions, the central limit theorem finds application, and the distribution will be Gaussian distributed [77].

$$\frac{dN}{d\Delta} \propto \exp -\frac{(\Delta - \bar{\Delta})^2}{2\sigma^2} \quad \text{with} \quad \sigma^2 = \frac{1 - \frac{1}{2}\beta^2}{1 - \beta^2} 4\pi N_A r_e^2 (m_e c^2)^2 \rho \frac{Z}{A} x \quad .$$

Here, Δ is the energy loss, $\bar{\Delta}$ its mean, x is the thickness of the absorber and ρ its density, N_A, r_e, m_e, Z , and A as defined at (5.2).

In thin absorbers, the incident particle does not normally suffer enough collisions for the central limit theorem to hold. The calculation of the energy loss distribution is complicated by the relatively large probability for single large energy transfers, up to a limit of T_{max} , which tends to skew the distribution and give rise to a very long high energy loss tail, as illustrated in Figure 5.10(a). Landau, Symon, and Vavilov have done various calculations of the energy loss distribution in thin absorbers, see Figure 5.10(b). Their applicability are different, and are characterised by the quantity

$$\kappa = \frac{\bar{\Delta}}{T_{max}} \quad \text{where} \quad \bar{\Delta} \approx \xi = 4\pi N_A r_e^2 m_e c^2 \frac{Z}{A} \frac{x}{\beta^2} \quad ,$$

Here, ξ is an approximation to the mean energy loss as given by the Bethe–Bloch equation.

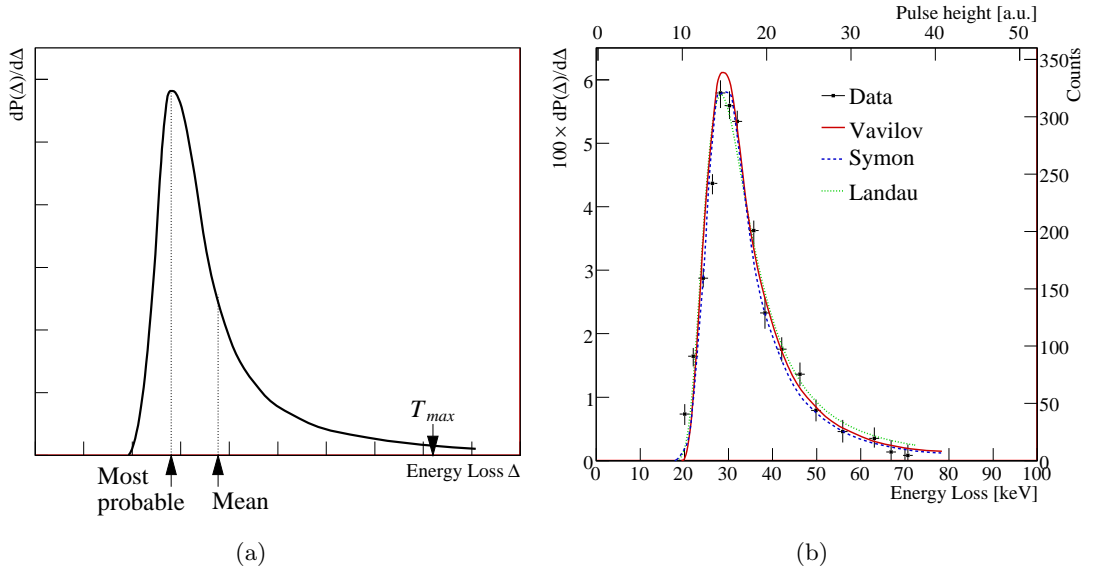


Figure 5.10: (a) Typical energy loss distribution in thin absorbers. Notice, that the most probably value, and the mean is not the same, and also note the long tail at high energy loss. Adapted from [77]. (b) Calculations of the energy loss distribution by Landau, Syman, and Vavilov compared to data. Adapted from [77].

When calculating the energy loss distribution, Landau assumed that

1. the T_{max} can be infinite, meaning $\kappa \rightarrow 0$,
2. the energy transfer in individual collisions is large enough that the excited target electrons can be treated as free, and
3. the incident particles do not lose any significant kinematic energy.

With these assumption, Landau derived the formula

$$\frac{dN}{d\Delta} = \frac{1}{\xi} \frac{1}{\pi} \int_0^{\infty} du e^{-u \log u - u\lambda} \sin \pi u \quad (5.3)$$

where

$$\lambda = \frac{1}{\xi} \left[\Delta - \xi \left(\log \xi - \log \varepsilon + 1 - C \right) \right]$$

$$\log \varepsilon = \log \left[\frac{(1 - \beta^2)I}{2mc^2\beta^2} \right] + \beta^2$$

Here, $C = 0.577$ is Eulers number, and other quantities as before. ε is essentially the minimum energy transfer as per assumption number 2 above [77]. The shape of (5.3) is shown in Figure 5.10(a). The long high energy loss tail forces the average away from the peak of the distribution. It is therefore more appropriate to parameterise the distribution in terms of the *most probable value* (or MPV), given by

$$\Delta_p = \xi \left[\log \left(\frac{2mc^2\beta^2\gamma^2}{I} \right) + \log \frac{\xi}{I} + j - \beta^2 - \delta(\beta\gamma) \right] \xrightarrow{\beta\gamma \geq 100} \xi \left[\log \left(\frac{2mc^2\xi}{(\hbar\omega_o)^2} \right) + j \right] \quad (5.4)$$

where $\hbar\omega_p = 28.816 \text{ eV} \sqrt{\rho \frac{Z}{A}}$ is the plasma energy, and $j = 0.2$. From the high energy limit of (5.4), one can see that the most probable energy loss flattens out over the MIP energy. So in thin absorbers we can expect a near constant energy loss if the incident particle has a $\beta\gamma > \beta\gamma_{MIP}$, motivating footnote⁶.

Vavilov extended Landau's calculations by essentially doing away with assumption 1 above. That means that Vavilov treated the maximum energy loss correctly. Although Vavilov's treatment of energy loss allows for better calculability and increased precision, the extra complexity makes it much more cumbersome to deal with, and the more simple Landau description is used.

Summing up, the energy loss in Silicon is Landau distributed around some most probable value that depends little on $\beta\gamma$ of the particle, when $\beta\gamma > \beta\gamma_{MIP}$, as shown in Figure 5.11. Also shown in the figure for reference are various other calculations of the mean energy loss in Silicon.

5.2 The Hybrid cards and VA1 pre-amplifier

Since the current signal from a silicon sensor is very small — on the order of a few Coulomb, it needs to be amplified immediately it is available on the bonding pads. For this purpose, a custom *hybrid* card was designed for the FMD. A schematic drawing of an inner-type hybrid is shown in Figure 5.12. The VA1₃ chips on the hybrid cards are *shaping* and *pre-amplifier* circuits to integrate and amplify the signal from the silicon sensor.

The hybrid card is glued to the back of the silicon sensors via 3 ceramic spacers, using non-conducting glue. The spacer is glued to the hybrid with epoxy glue, while the other side of the spacer is glued with thermal glue to the sensor. This allows for easy replacement of a sensor should it break. Each of the silicon strips is then bonded up to the pitch adaptors on the hybrid cards. The pitch adaptors are ceramic plates with 128 implanted gold lines, and serve to adapt from the relatively wide pitch on the silicon (250 μm on the inner, and 500 μm on the outer silicon) to the relatively small 50 μm pitch on the VA1₃ chip. Each of the 128 lines on the pitch adapter are then bonded to the 128 inputs on the VA1₃ chips, as shown on the right in Figure 5.12.

The shape of the hybrid cards is designed to match the shape of the silicon sensors, though about 5mm narrower on either side to leave room for bonding the silicon strips to the pitch adapters. Thus there are two types of hybrids: one that matches the inner type sensors and one that matches the outer type sensors. The *Printed Circuit Board* (PCB) of the hybrids are made




Figure 5.11: Mean and most probable energy loss in $300\ \mu\text{m}$ Si. GFIMATE is the routine used in the GEANT simulation program and RPP refers to the *Review of Particle Physics* [81].




Figure 5.12: Schematic drawing of an inner-type hybrid. The right-hand side shows a schematic section through the stack at A-A. Note, that the detail is not to scale.

slightly longer — about a millimetre or two — to provide protection for the fragile silicon during mounting and transportation.

Since there are 512 strips on either side of the inner sensors, there are 4 VA1₃ pre-amplifier chips on either side of the inner hybrids, giving a total of 8 VA1₃'s on an inner hybrid. Because of the hexagonal geometry of the silicon, it is not possible to place the top-most pitch-adapter near the edge of the hybrid PCB. Instead, the pitch adapter is placed further back on the board, and the signal lines run over the PCB (see also Figure 5.12). This should in principle not make any difference to the performance of these strips, but in practice it turns out that these lines may indeed pick up more noise than those for which there is no such lines on the PCB (see Section 7.2.3). There are only 256 strips on either side of an outer sensor, so only 2 VA1₃ pre-amplifiers are needed, giving a total of 4 VA1₃'s on an outer card.

Near the middle of the hybrid cards is a connector for the bias voltage to the silicon, and a large capacitor to decouple the silicon from the input voltage. The bias voltage ground is linked to this capacitance too, to have a shielded common ground for the VA1₃ and silicon.

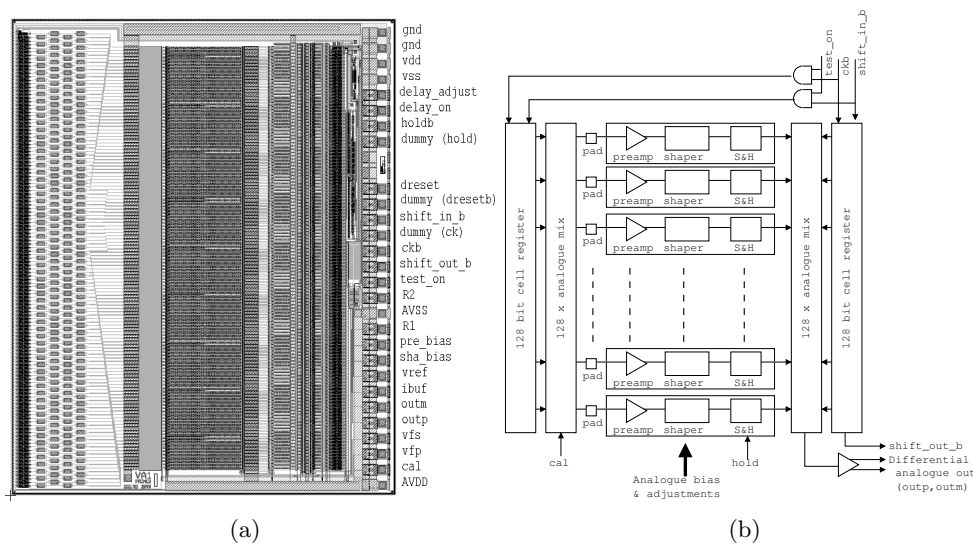


Figure 5.13: Architecture of the VA1₃ chip.

5.2.1 The VA1₃ chip

The VA1₃ chip is part of the time-tested and widely used VIKING [82] family of silicon pre-amplifier and shaper circuits. A picture of the chip is shown in Figure 5.13(a). It is, as the hybrid PCB⁷, designed and manufactured by IDEAS in Norway. Its main features are low noise, high gain, constant shaping time, and radiation tolerance.

Functional description Functionally, the VA1₃ chip is composed of a pre-amplifier, followed by a shaping circuit, and ending in a sample-hold circuit for each input, as shown in Figure 5.13.

Pre-amplifier: Amplifies the relatively weak current signal from the silicon. The level of amplification (of order $10 \mu\text{A}/\text{fC}$) can be set by the inputs *vfp* and *prebias*. On the FMD hybrid cards, the *prebias* is fixed via resistors on the hybrid card, while the *vfp* can be controlled externally.

Shaper: Integrates the current signal from the pre-amplifier to attain a constant shaping time. The shaper is a CR-RC circuit, which means it differentiates the input, and then re-

⁷Printed Circuit Board

integrates with the net effect that high-frequency noise is filtered and detector performance improved. The shaping circuit is controlled by the input voltages `vfs` and `sha.bias`. These voltages essentially determine the time constants of the shaping circuit.

Sample–hold: The sample–hold circuit holds the charge output by the shaper until it is told to release it by the external signal `hold`. That is, as long as `hold` is asserted, then the sample–hold circuit will accumulate the charge. Then `hold` is released, the circuit will bleed–off the charge to the subsequent network.

Noise The noise of the VA1₃ chip depends largely on the capacitance of the bulk silicon crystal C_b , the internal feed–back capacitor C_f , the leakage current I_l from the silicon, and the capacitance of the shaping filter (see for example [77, 82, 83]). IDEAS has measured the noise (in terms of electrons) as a function of the simulated full strip capacitance C_s , strip leakage current I_s , and shaping time t_p (essentially the capacitance of the shaper). The results are shown in Figure 5.14.




Figure 5.14: Noise (in electrons) of the VA1₃ chip. The left panel shows the noise as a function of the peaking time of the chip, the middle as a function of detector capacitance, and to the right as a function of leakage current. For comparison, one should note that a minimum ionising particle in 300 μm silicon will produce a 22400 electron signal. Adapted from [84].

The noise in units of Equivalent Noise Charge (ENC) from various the pre–amplifier (n_{pa}), leakage current (n_l), bias resistor (n_R), and bulk resistance (n_b) is given by [82, 83]

$$\begin{aligned}
 n_{pa} &= f(C_s) & C_s &= C_b A_s / A + C_l l_s \\
 n_l &= \frac{e}{q} \sqrt{\frac{q I_s t_p}{4}} & I_s &= I_l A_s / A \\
 n_R &= \frac{e}{q} \sqrt{\frac{t_p k T}{2 R_p}} & R_p &= \frac{R_f R_b}{R_f + R_b} \quad , & (5.5)
 \end{aligned}$$

where

$f(C_s)$	$258 + 6C_s$	First order polynomial fit to data in middle panel of Figure 5.14
C_i	1.2 pF/cm	Inter strip capacitance per unit length
C_b		Total capacitance of the sensor (see Figure 5.7)
A_s	$w_s \times l_s$	Surface area of a strip of width w_s and length l_s
A		Surface area of sensor
C_s		Combined strip and inter-strip capacitance
e	2.718281...	Base of the natural logarithm
q	$1.60217653(14) \times 10^{-19}$ C	The charge of an electron
I_l		Total leakage current of sensor (see Figure 5.7)
I_s		The leakage current per strip
t_p	1.35 μ s	Peaking time of the shaper circuit
k	$1.3806505(24) \times 10^{-23}$ J/K	Boltzmanns constant
T	~ 300 K	Temperature
R_f	∞	Pre-amplifier feed-back resistance. This is controlled by the voltage input <code>vfp</code> . When <code>vfp</code> is set correctly, then R_f is in principle infinite, and R_p is simply R_b .
R_b	20 M Ω	Bias resistance
R_p	R_s	Parallel combination of strip resistance R_s and pre-amplifier feed-back resistance R_f

The total noise is then given by adding the contributions in quadrature

$$n = \sqrt{n_{pa}^2 + n_l^2 + n_R^2} \quad . \quad (5.6)$$

From (5.5) and that $R_f \gg R_b$ such that $R_p \equiv R_b$, it is clear that n_R — the contribution from the parallel resistors — does not depend on any sensor or strip specific parameters, and evaluates to 200.60. (5.5) and (5.5) are evaluated in Table 5.1 for the shortest and longest strips in both inner and outer sensors.

Sensor parameters					
		Inner		Outer	
C_b	pF	~ 2400		~ 2100	
I_l	nA	~ 5300		~ 4300	
A	mm ²	8292		8317	
w_s	mm	0.25		0.50	
Strip parameters					
		Short	Long	Short	Long
l_s	mm	14.2	50.7	12.1	21.2
Calculated strip quantities					
I_s	nA	0.08	0.29	0.61	1.07
C_s	pF	2.60	9.29	5.96	1.04
Noise contributions					
n_{pa}		273.8	314.0	294.0	320.8
n_l		35.6	67.2	97.5	129.1
n		341.3	378.6	369.0	399.8
Signal-to-noise		65:1	59:1	61:1	56:1

Table 5.1: Various parameters of the VA1₃ and silicon strips used in the calculation of the ENC.

A MIP produces ~ 22400 electrons in silicon of 300 μ m thickness. This is used in Table 5.1 to give an estimate of the absolute best signal-to-noise ratio possible with the FMD sensors and

the VA1₃ pre-amplifier. The best possible signal-to-noise ratio is in the range from 59:1 to 65:1 for the inner sensors, and 56:1 to 61:1 for the outer sensors.

Readout and calibration pulser Common to all inputs is an analogue shift-register. The shift-register controls a multiplexer, or many-to-one, by acting as a cursor that selects the current input. Asserting the signal `shift_in.b` together with `shift_clk` resets the cursor to point to the top of the shift register, while each subsequent `shift_clk` moves the cursor one down the shift register, selecting the next channel for input. Thus, the 128 channels of the VA1₃ can be read out into a single output line (`outp-outm`) via the sequence shown in Figure 5.15.

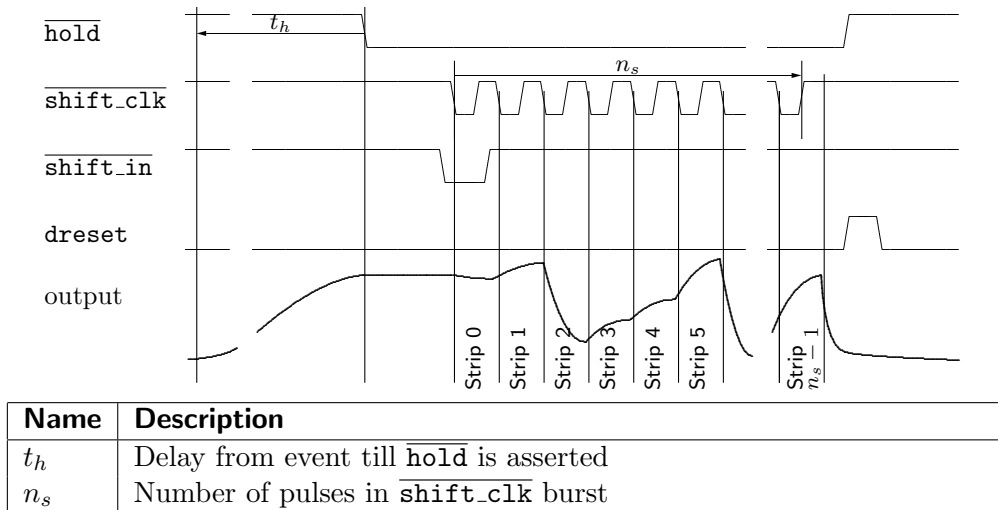


Figure 5.15: Read-out of all 128 channel in a VA1₃ chip. `shift_in` together with `shift_clk` resets the shift register to point at the first channel, while each subsequent falling edge of `shift_clk` makes the shift register point at the next channel. At the end of the read-out, a single pulse on `dreset` can be sent to clear the digital parts of the chip, though it is generally not needed.

Alternatively, the cursor can be preset to point at the first channel before the event, and then after the full read-out, be set back to the beginning. This sequence, as shown in Figure 5.16 improves the quality of the read-out of the first channel, and makes the signals less susceptible to disturbances from the logically circuits.

The output of the multiplexer is passed through differential splitter to the output pads `outp` and `outm`. The signal is therefore differential, meaning that `outm` is the negative reflection of `outp`, and the signal value is given $s = \text{outp} - \text{outm}$. The signal is transmitted differentially to reduce pick-up of external sources of noise. That is, since the two lines are drawn close to each other, outside sources of noise will most likely affect both lines equally, meaning that their individual values may change but the difference will not. The differential lines are then drawn closely across the hybrid to the connector at the top of the PCB.

Parallel to the output shift register and output multiplexer is an input shift register and input multiplexer. This input is enabled by the external signal `test_on` (see Figure 5.13(b)). The input multiplexer routes a current signal generated on the dedicated line `cal`. Internally to the VA1₃ is a capacitor connected to this input, so that a voltage step on `cal` will induce a proportional current on the input multiplexer. By enabling `test_on` and providing a varying voltage step on `cal`, one can determine the internal amplification of the VA1₃. This is used by the FMD to periodically do gain calibrations of each strip. Note, since the input shift register works in complete parallel to the output shift register, it is only possible to induce a charge on a single VA1₃ channel at a time. Therefore, to do a gain calibration of say the 3rd channel of the VA1₃,




Figure 5.16: Read-out of all 128 channel in a VA1₃ chip. The channel is primed before the event to have the cursor at the first channel. Subsequent falling edges of `shift_clk` make the shift register point at the next channel. At the end of the read-out, a single pulse on `dreset` can be sent to clear the digital parts of the chip, and `shift_in` together with `shift_clk` resets the shift register to point at the first channel.

the sequence of signals shown in Figure 5.17 must be followed.




Figure 5.17: Setting up for gain calibration of the 3rd channel of a VA1₃

In-Laboratory Tests of Sensor-Hybrid Modules All sensor-hybrid compounds (called modules), were extensively tested in the laboratory before being released for testing with rest of the system. In particular, the modules were tested for broken or shorted bonds, dead channels and so on. Figure 5.18 shows an example of the test results for a single module. The results were catalogued and used later for validation and selection of sensors for installation.

5.3 The Digitiser Cards

Once the charge for the silicon has been amplified and shaped by the VA1₃'s on the hybrid card as outlined in the previous section, the charge needs to be turned into digital numbers that can be stored in computer storage and processed by computer software. This is the primary function of the FMD digitiser (FMDD) card. Each of the 10 half-rings has an associated FMDD, placed on the opposite to the hybrids on the honey-comb support plate. Figure 5.19 shows the two types of FMDD cards — inner and outer — and the principal components on the cards.

The FMDD is responsible for

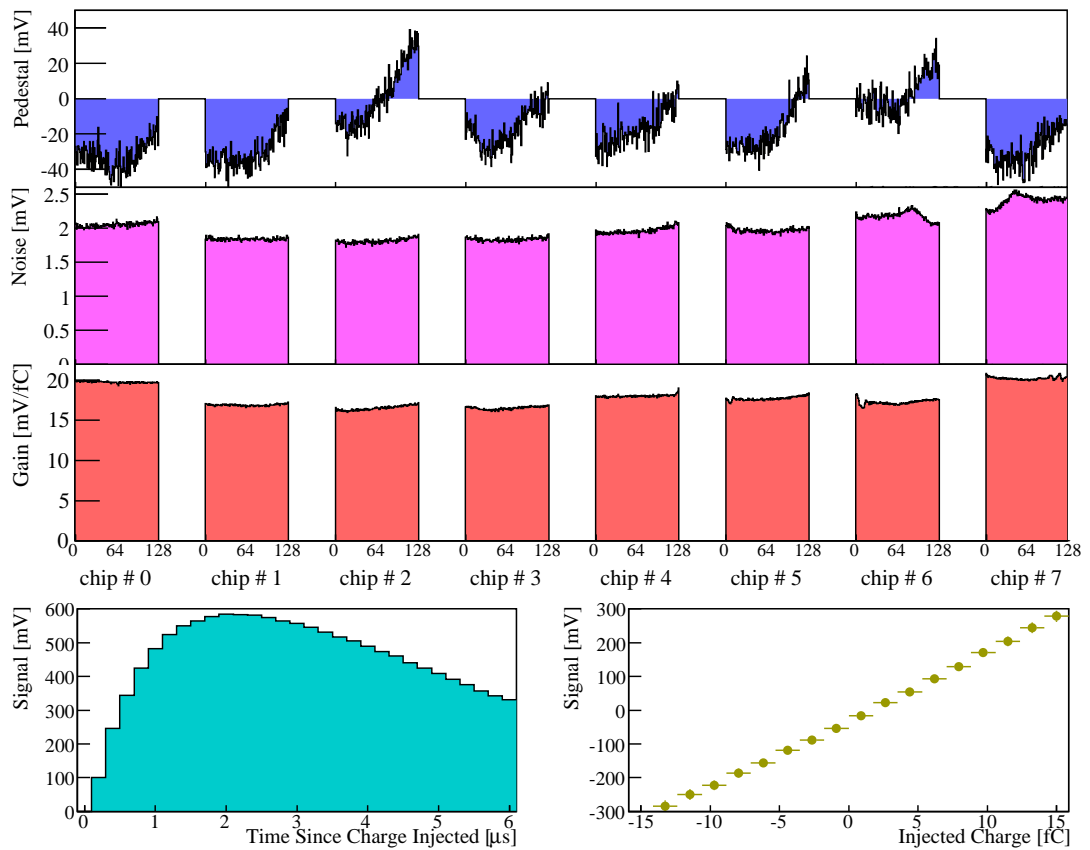


Figure 5.18: Testing of a single Sensor-Hybrid compound. From the top is shown the pedestal (base-line signal), noise, and gain of a single module. The two bottom plots show the signal as a function of the hold time (probing for the optimal shaping time) and the signal as a function of injected charged for a single strip.

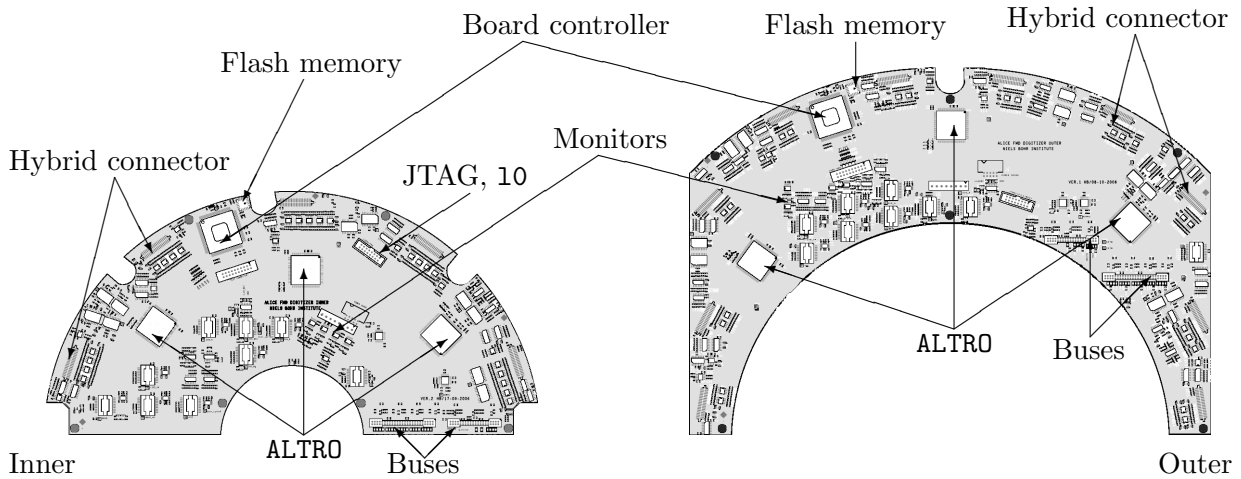


Figure 5.19: The two types of FMD Digitiser (FMDD) cards. There are 3 ALTROs and 1 BC on either card, as indicated on the figure. On the inner type (left) there are 5 connections to the hybrid cards, while on the outer type (right), there are 10 connectors. The bus cables are attached in two slots on either card, as indicated on the figure. Not indicated in the figure, are the low-voltage, JTAG, and LVL0 trigger connections — one on either card. The shape of the outer card is dictated by the maximum size of card that the component mounting company could handle. The boards are shown to scale.

- Converting the analogue input from the VA1₃'s to digital signals. This is done by the ALTRO chip [85], inherited from the ALICE TPC project. The choice of this chip for the FMD analogue-to-digital converter (ADC) dictates many other choices in the front-end electronics chain [86], since the ALTRO has specific requirements of the surrounding environment.
- On request, transfer the processed data from the ALTROs to the read-out-controller unit.
- Trigger handling for the VA1₃ pre-amplifier and ALTRO chips. The Board Controller (BC) processes the 3 levels of triggers and times the read-out of the VA1₃'s and conversion by ALTRO accordingly.
- The BC also controls the read-out sequence of the VA1₃, as outlined in the previous section.
- Control of the gain pulser calibrations. On request, the BC can set up conditions for iterating through the 128 channels of the VA1₃, and for each VA1₃ channel, increment the gain pulser to record data for gain calibrations.
- Control voltage settings for the VA1₃'s. The BC sets the VA1₃ input voltage v_{fp} which determines the feed-back resistance of the pre-amplifier. The VA1₃ input voltage v_{fs} and current sha_bias , controlling the shaping of the VA1₃, is also controlled from the BC.
- There are 4 ADCs on the FMDD for monitoring the voltages, currents, and temperature on the card. The BC periodically read these values to assert the operating conditions of the board.

The design of FMDD card is based on the TPC FEC card, but modified at the Niels Bohr Institute to accommodate the features needed by the FMD — for example the *Digital-to-Analogue Converters* (DAC) to set VA1₃ biases and make calibration pulses.

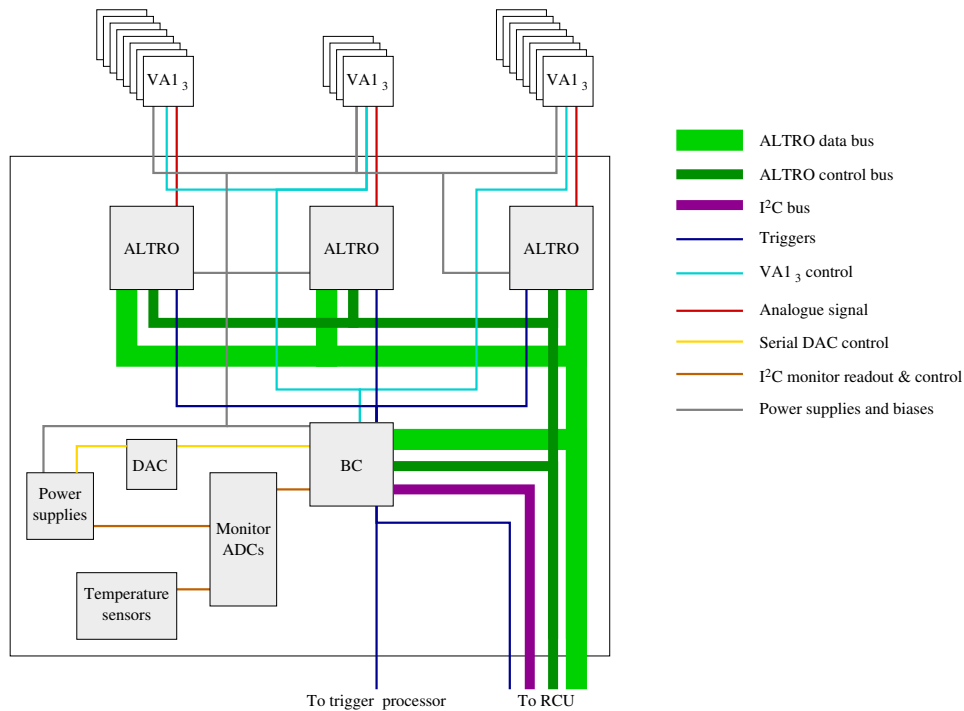


Figure 5.20: Schematic diagram of FMDD BC cards. Note, that both inner and outer cards have, in total, 20 VA1₃'s connected.

Hybrid connections Short 40 pin flat-cables connect each hybrid card to a connector on the FMDD. 24 of these lines serve to bring low-voltage, biases, and control signals from the board controller on FMDD to the VA1₃'s on the hybrid card. The remaining 16 lines bring the differential signal from the VA1₃'s to the ALTRO analogue-to-digital converter chip on the FMDD. The signal from the VA1₃'s is a current signal, so to convert it to a voltage signal, as accepted by the ALTRO, a series of receivers is placed adjacent to the hybrid connectors. These are small operational amplifiers that converts the signal and shifts the level to the input level of the ALTROs differential analogue inputs.

Since there are 8 VA1₃'s on the inner hybrid cards, and 4 on the outer hybrid cards, the number of data line pairs on each connector varies. So on the outer FMDD four data line pairs are drawn from each connector to 4 inputs on the ALTRO, while on the inner cards 8 data line pairs are drawn to 8 inputs on the ALTRO. Since each ALTRO has 16 inputs, there are two hybrids connected to an ALTRO on the inner type FMDD, while there are 4 hybrids connected to an out ALTRO on the outer type FMDD. Note however, that the middle ALTRO on both types of cards only has half the number of inputs connected to a hybrid (1 on the inner, 2 on the outer), so that only 8 input channels are used on this ALTRO.

Buses The 3 ALTROs and the BC of each card all sit on a common bus. The bus protocol is the custom ALTRO-bus protocol. It consist of 40 bits of data and 5 control bits. This bus is controlled by the RCU, except during transfer of the processed data, when the ALTRO takes control. Several FMDD cards can be attached to the bus, while there can only be one RCU. All control of the FMDD components go via this bus, though some control of the BC can be done via an I²C bus. The I²C bus is used to monitor the status of the FMDD card, leaving the ALTRO bus free to do data transfer.

The bus lines all run in 2 twisted-pair flat cables connected to the card as shown in Figure 5.19. The bus lines are listed in Table 5.2. Since the bus cables are rather long (~ 3 m) all lines are transferred in one line of a twisted pair, while its partner is held at 0 V relative to ground. In this

way, the twisted pair is relatively close to being a transmission line — i.e., with little reflection and no inter-pair couplings. This was found to be necessary to obtain proper signal quality on the FMDD.

ALTR0 bus			
<i>Name</i>	<i>Width</i>	<i>Direction</i>	<i>Description</i>
bd	40	RCU↔FMDD	Data and address
write	1	RCU→FMDD	Write enable
cstb	1	RCU→FMDD	Control strobe
ackn	1	RCU←FMDD	Acknowledge
error	1	RCU←FMDD	Error
trsf	1	RCU←ALTR0	Data transfer enable
dstb	1	RCU←ALTR0	Data strobe
l1	1	RCU→FMDD	Trigger LVL1
l2	1	RCU→FMDD	Trigger LVL2
rstb	1	RCU→FMDD	Reset
clk	1	RCU→FMDD	General clock
I²C bus			
<i>Name</i>	<i>Width</i>	<i>Direction</i>	<i>Description</i>
scclk	1	RCU→BC	I ² C clock
scdin	1	RCU→BC	Data
scdout	1	RCU←BC	Data
Other			
<i>Name</i>	<i>Width</i>	<i>Direction</i>	<i>Description</i>
cardsw	16	RCU→FMDD	Card switch, one per possible FMDD
interrupt	1	RCU←BC	Interrupt set in case BC sees a monitored value is out-of-range

Table 5.2: Lines in the two bus cables. The width is in number of bits.

The signals on the bus is transmitted according the GTL standard, which means that there are a number of GTL drivers on the FMDD. These sit next to the bus connectors. To avoid reflection on the bus lines, the last card on the bus cables must terminate the bus by providing a termination voltage through fixed resistors. These are also placed on the card next to the FMDD. These termination resistors are mounted in such a way, that they can easily be removed or added if a card is placed in some other location on the bus. This allows us to have only two types of cards for the whole FMD.

Power connectors and regulators Placed almost in the middle of both cards is the low-voltage power supply with 4 lines, as detailed in Table 5.3. Note, that all grounds of the 4 power lines are tied together at the connector on the FMDD. The power supplies in the other end should therefore be *floating*, meaning that they do not refer to some common ground. If not, there will be ground loops which can pick up radio noise, and affect the operation of the front-end electronics. All power lines are distributed to the components on the card via one or more power regulators to ensure stability and decoupling. The output of the power regulators, with the exception of one, is controlled by the **cardsw** line on the bus cables. That is, when **cardsw** is low, then there is no power on the components and the card is off. When **cardsw** is high, then the power regulators supply all components on the card. The one power regulator that is not controlled by the **cardsw** is the power regulator that supplies the bus drivers with power — with out this power, the card would never feel a change on the **cardsw** line. The RCU, which controls the **cardsw** line, can therefore selectively turn on and off FMDD cards.

The lines V_a^- and V_a^+ are used to supply the analogue part of the card and VA1₃'s. They are

therefor separated from the rest of the power on the card. On the card, there are separate power regulators for each use of these power lines. Due to the relatively large power consumption on the V_a^- line, and consequently large heat dispersion from the power regulators on this line, there are water cooling pipes attached to the power regulators that are fed by V_a^- . Typical temperatures in the laboratory without cooling of the surroundings are 45 °C, and the board should operate up to 55 °C.

The rest of the lines — V_d^+ and V_d^2 — only supply digital components on the chip, including the ALTROs and the BC.

Name	Voltage ⁸ [V]	Current [A]	Description
V_a^-	-3.0	2.5	Analogue negative power, used for VA1 ₃ and receivers
V_a^+	+2.5	0.7	Analogue positive power, used for VA1 ₃ and receivers
V_d^1	+3.5	1.1	Digital power, used by ALTRO and BC
V_d^2	+4.3	0.1	Digital power, used by ALTRO and BC

Table 5.3: Low-voltage power supply

JTAG and flash memory The BC is implemented in a *Field-Programmable Gate Array* [87] (FPGA), which is a programmable chip for doing logic and simple operations. The particular chip used for the FMDD BC can be re-programmed, either from an on-board chip or remotely via a JTAG connection. The FMDD are fitted with a JTAG connection and a flash-based memory chip, since the FMDD cards are buried deep inside ALICE with no easy access. The JTAG connection can be used to re-program the BC FPGA directly, or to change the configuration stored in the flash memory. On the next power-up, the flash memory will automatically re-program the BC FPGA with its current configuration. This makes it easy for the FMD team to easily upgrade the behaviour of the BC in case of problems or added functionality is needed.

Trigger input/output Since the level 0 trigger (LVL0) is not available from the RCU, as the rest of the triggers LVL1 and LVL2 are, the FMDD is fitted with a separate LVL0 cable input. This signal is needed by the BC to assert $\overline{\text{hold}}$ at a reasonable time after the collision has happened ($\sim \mu\text{s}1.5$). The out-going signal *busy* is needed, as the BC needs to alert the Central Trigger Processor (CTP) of ALICE that the FMD is about to handle a trigger and should not get new triggers until the FMD is ready for it.

5.3.1 The ALTRO Chip

The ALTRO (*ALice Tpc Readout*) chip [85, 88] is the analogue-to-digital converter (ADC) developed for the ALICE TPC. Beside the TPC, the following detectors in ALICE also make use of the ALTRO: FMD, EMCAL, and PHOS. It is a 16-channel 10bit parallel ADC, meaning that 16 analogue inputs can be sampled at a time, with a precision of $1/2^{10} = 1/1024$. The main features of the ALTRO chips is

- Parallel processing of 16 input channels. Each of the 16 inputs have a separate chain for processing the input, including advanced signal filtering. The digitised data is also stored in separate event buffers for each channel. This allows for a fast read-out of even large detectors such as the ALICE TPC.
- High radiation tolerance. The chip has undergone extensive irradiation tests by the CERN ALICE TPC group, and found to be very radiation tolerant.

The ALTRO also has a digital part that handles common logic, such as register input/output, trigger handling, and the interface to the ALTRO bus. Figure 5.21 shows the blocks of the ALTRO schematically.

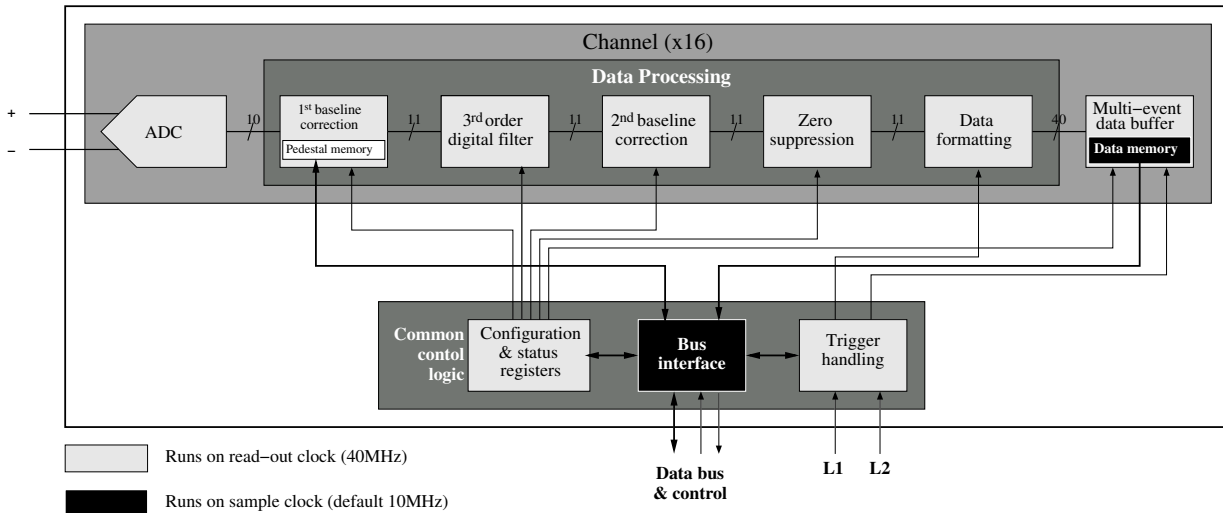


Figure 5.21: Schematic architecture of the ALTRO chip. The processing chain for all 16 channels is the same — First, the analogue signal is converted to 10 bits. Then a series of digital processing units are invoked in turn, if so enabled in the configuration registers. Finally, the processed data is formatted into 40bit words and stored in the multi-event buffer. All this is done at the speed of the sample clock (default is 10 MHz). On request, the data is strobed out over the bus. The bus interface and the data memory is working at the speed of the read-out bus (40 MHz).

There are two clock inputs on the ALTRO, the *read-out* clock (or `clk`) and the *sample* clock (or `sample_clk`). The `clk` is global to the whole front-end electronics chain, and is distributed to the FMDD from the RCU. The `sclk` is generated on the FMDD by the BC, allowing us to control the sample speed of the ALTRO channels relative the the read-out speed of VA1₃'s.

One VA1₃'s output is connected to one channel of the ALTRO, and since the VA1₃'s multiplex 128 strips into one output, each ALTRO channel will see data from 128 individual strips, one after another. However, since the ALTRO may sample its input at a higher rate than the signals is clocked out from the VA1₃'s, it is possible to get more than one sample per strip. This is shown in Figure 5.22, and is known as the *oversampling rate* ν_{over} . The oversampling rate is fixed by the ratio between the `sample_clk` and the `shift_clk` clock. Both are generated by the BC and can be configured, including the relative phase of the two. It is important that the relative phases is adjusted such that the last sample (down-edge of `sample_clk`) of a given strip falls near the end of the pulse signal, to ensure that the signal has reached its full value. This has to be investigated using an oscilloscope, but need only be done once for any given value of ν_{over} .

Due to all kinds of dormant voltages and currents in the system and in-avoidable mismatches in input and output, the “zero” signal will not convert to a digital value of exactly 0. Rather there will be some base level of the converted number, called the *pedestal*. In fact, the pedestal has a useful purpose, in that it assures that the “zero” signal is in the range of the ADC, and one can therefore calibrate the signal to this base level.

Signal processing

As is clear from Figure 5.21, the ALTRO is a rather advanced chip, and can do a lot of processing on the collected data directly. However, not all of the shown steps are relevant to the FMD and will not be discussed here. In particular, only the 1st baseline correction and the zero-suppression are relevant.

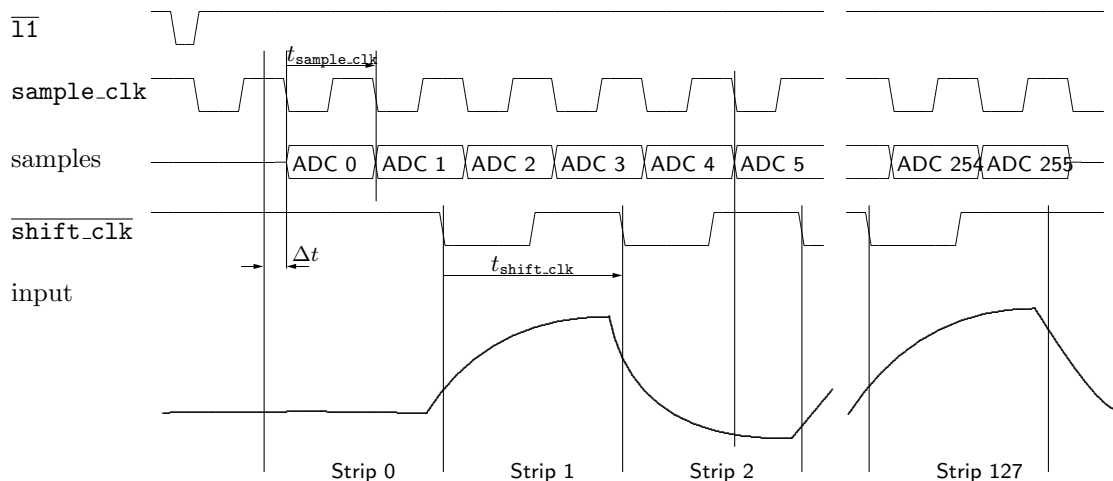


Figure 5.22: One ALTR0 channel sampling data from 128 strips. The figure shows the case of an oversampling rate of 2. Δt , the relative phase between `sample_clk` and `shift_clk` is tuned so that the last sample of each strip falls at the full value of the strip signal.

1st baseline correction or the on-line pedestal subtraction module. After the data has been digitised, it is possible to subtract from the digital value a fixed number — the so called *fixed pedestal* value (or *fpd*). The value to subtract is stored in a register on the ALTR0, and can be accessed via the ALTR0 bus.

However, there may be sample number dependent fluctuations in this base level. In the FMD this is evident as each VA1₃ channel will end up in different samples, and there is no reason to expect the pedestal to be the same for all 128 channels of a VA1₃. For the TPC, each sample corresponds to a different depth in the gas volume, and it reasonable to expect variations in the pedestal over the 2.5 m drift distance.

To compensate for this, each ALTR0 channel has a *pedestal memory*⁹ (*pmem*) into which one can upload sample dependant pedestal values to be subtracted in the 1st baseline correction module. In fact, this memory can be used in various modes — for example as a look-up table, meaning that the current sample value is looked up, and the corresponding table value is returned as the sample value. Such modes mostly find their application in testing the electronics. The mode is set in 4 lower bits of the configuration register `dpcfg`.

Zero suppression Since the event size in number of samples can be come rather big, especially for the TPC, there is a zero suppression module at the end of the processing chain. If enabled in bits 12 to 19 in the configuration register `dpcfg`, the module will subtract a fixed value, set in the configuration register `zsth`, from the sample values, and only samples that are larger than zero are retained. It is possible to include some number of samples before and after the sample values exceed the limit, in order to get more information about the signals.

Zero suppression is only mildly relevant to the FMD, since the detector will most often have a signal in a large fraction of the channels, thus nullifying the usefulness of suppression the rarer zeros. However, zero suppression is an option if the data size is later found to be too big. All software is designed with zero suppression in mind, so turning on zero suppression on the front-end should not cause major problems.

⁹Sometimes also referred to as the *pattern memory*.

Data formatting

Once the data have been processed by the various steps in the processing chain, the 10bit ADC counts are formatted into 40bit data words. In case of no zero suppression, the 10bit ADC counts are simply concatenated into one large block, and the time of the last 10bit ADC count, relative to the arrival of the $\overline{11}$ trigger in units of `sample_clk`, and the total length of the block is appended as separate 10bit words. The block is then padded with `0x2AA` (1010101010 in binary) up to the 40bit word size. Finally, a *trailer* specifying the origin of the data and the total data size is appended, inter-spaced with the fill patterns 1010, as shown in Figure 5.23(a).

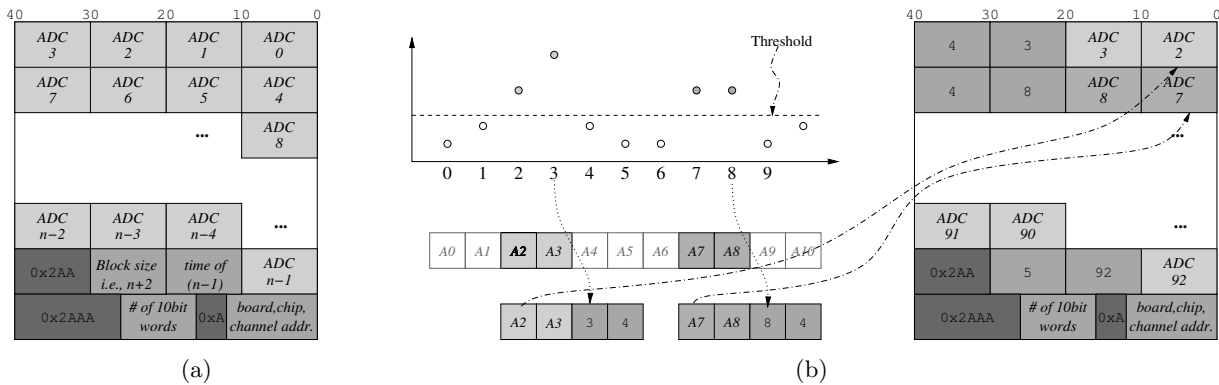


Figure 5.23: ALTRO channel data format. The trailer consist of the channel address, given as 12bits of (*board, chip, channel*). (a) In case of no zero suppression. All n ADC counts are put into one block (here starting at $ADC\ 0$ and ending at $ADC\ n-1$) followed by the time of the last 10bit word, and the total length of the block ($n + 2$). (b) In case of zero suppression, only the values over threshold are stored. If there appears holes due to ignored counts (such as counts A_3 , A_4 , and A_5), then a block trailer is appended, consisting of the time of the last ADC count, and the total size of the block. The blocks are then concatenated, and filler words are appended to pad to the 40bit word size.

In case of zero suppression, only the values over threshold are stored. If there appears holes due to ignored counts (such as counts A_3 , A_4 , and A_5 in Figure 5.23(b)), then a *block trailer* is appended, consisting of the time of the last ADC count, and the total size of the block. The blocks are then concatenated, and filler words are appended to pad to the 40bit word size. The same channel trailer is appended as in case of no zero-suppression.

Read-out of the ALTRO data

After successfully handling an event, the RCU will ask all available ALTRO channels in turn for their data. The RCU sends the command `chrdo` to the individual channels (see also Section 5.3.3) and waits for the acknowledgement. The ALTRO will then assert the bus line `trsf` (transfer), and start strobing out the 40bit words stored in the data memory, using the `dstb` signal on the bus. The sequence is shown in Figure 5.24.

Since the data is strobed out at the speed of `clk` (40 MHz) with no handshakes, it is of the utmost importance that quality of the strobe `dstb` is high. The strobe `dstb` is implemented as a simple AND between the transfer enable `trsf` and `clk`, as shown in Figure 5.25. It is therefore important that the quality of `clk`, the master clock, is very high at arrival at the FMDD. For this reason, the 3m long bus cables from the RCU to the FMDD are twisted pair cables, where one partner is held at ground and the other transmits the signal. This is to ensure that the signal line is close to being a proper transmission line. Experiments with regular flat-cables, shielded flat

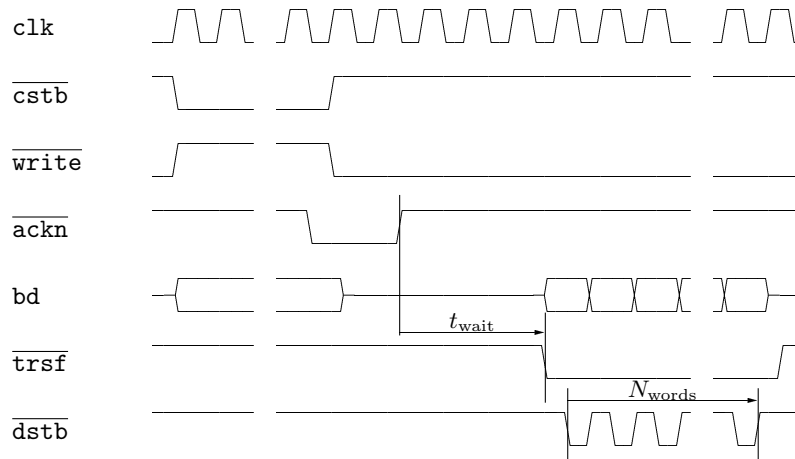


Figure 5.24: Timing of ALTR0 channel readout.

cables, and other types of cables, have shown, that there is potentially quite a lot of cross-talk over the cable. Using twisted pair cables, the cross-talk is considerably reduced since the signal lines are shielded by its ground partner from other signals.




Figure 5.25: Generation of $\overline{\text{dstb}}$. The external sample clock signal `clk` is feed to a 2-way logical AND gate, and the data transfer enable `trsf` is feed to the other input of the AND gate. The result is when `trsf` is kept low, there are appears a clock of the same frequency on the output `dstb`. Note however, that the generated clock is not as well defined as the input clock `clk` due to the AND gates internal workings.

5.3.2 The Board Controller

The FMDD board controller is implemented as firmware in an Altera ACEX1K Field-Programmable Gate Array. The firmware is based on the TPC Verilog code base, but modified for the FMD and completely re-written in VHDL code. A schematic block diagram of the FMDD BC is shown in Figure 5.26. The BC is responsible for

- Handle triggers. On receiving LVL0, LVL1, and LVL2 triggers, the BC implements the behaviour needed to get the current signals from the VA1₃'s to the ALTR0 and digitised there.
- Control the VA1₃ circuits by setting configuration voltages and bias voltage.
- Control the pulser calibration runs.
- Supply the ALTR0s with a properly tuned `sample_clk`.
- Monitor the relevant temperatures, voltages, and currents on the FMDD.

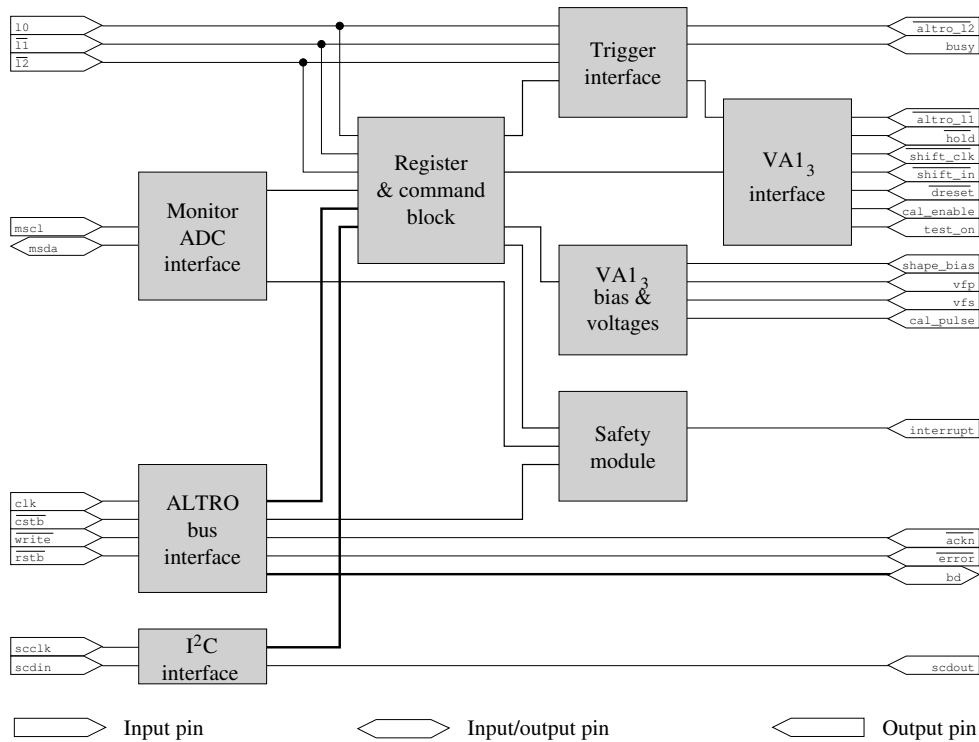


Figure 5.26: Schematic overview of the FMDD BC architecture

Register Block All aspects of the FMDD BC firmware is configurable via registers and commands. Many of the registers are inherited from the TPC BC implementation, but the FMDD BC does add some specific registers needed for the operation of the FMD detectors. These are listed in Table 5.4. All registers are accessible either via the ALTRO bus or via the I²C bus.

In addition, there is a number of commands defined for the BC. Again, some are inherited from the TPC BC while others are new to the the FMDD BC. The commands relevant to the FMD are listed in Table 5.5.

VA₁₃ interface As mentioned in Section 5.2.1, the VA₁₃ chip needs a sequence of signals to read-out the data. In particular, the $\overline{\text{hold}}$ signal needs to be timed to fit with the peaking time of the shaper of the VA₁₃, and a proper sequence of $\overline{\text{shift_clk}}$ and $\overline{\text{shift_in}}$ needs to be sent to shift the output from one input channel to the next. The VA₁₃ interface of the BC takes care of this. It gets a start signal from the trigger handling component which initiates the internal *Finite State Machine* (FSM). The state machine implements the read-out sequence shown in Figure 5.16 (except for the $\overline{\text{hold}}$ which is handled by the trigger handler) and at the same time synchronises the $\overline{\text{l1}}$ signal send to the ALTRO chips. The synchronisation is done to ensure that the ALTRO digitisation is started properly relative to when the data is expected to arrive at the input of the of the ALTRO channels.

The first and last strip to read out, and the speed and phase of the $\overline{\text{shift_clk}}$ are also adjustable via register settings (see also Table 5.4). If the first and last strip to read out are equal, then no $\overline{\text{shift_clk}}$ are sent, which means that the VA₁₃ cursor never moves and one and the same strip is read out.


Trigger handling The general trigger sequence of ALICE is shown in Figure 5.27. When a l0 trigger is received at around 1.2 μs after the two beams have collided, the trigger handler raises the external busy signal, and the FSM of the modules goes into a state where it waits for $\overline{\text{l1}}$ to validate the event. In this state, the trigger handler waits for a configurable amount time before

Address	Name	Description
Relevant BC registers inherited from the TPC BC		
0x0B	l1cnt	$\overline{11}$ trigger counter
0x0C	l2cnt	$\overline{12}$ trigger counter
0x11	csr0	Monitor enable and status bit mask
0x12	csr1	Error and status
0x12	csr2	Power enables
0x13	csr3	Monitor ADC end and clock ratio warning
FMD specific BC registers		
0x20	fmdd_status	FMDD BC status (read-only)
0x21	l0cnt	l0 trigger counter (read-only)
0x22	hold_wait	Number of clock cycles to wait from reception l0 to setting $\overline{\text{hold}}$
0x23	l1_timeout	Number of clock cycles to wait for $\overline{11}$ after receiving l0
0x24	l2_timeout	Number of clock cycles to wait for $\overline{12}$ after receiving l0
0x25	shift_div	Phase and period of the <code>shift_clk</code> in units of clock cycles
0x26	strips	First and last strip to read out
0x27	cal_level	DAC setting for calibration pulse size for stand-alone tests, and step size for automated gain calibrations
0x28	shape_bias	DAC setting for shaping bias
0x29	vfs	DAQ setting for shaping feed-back resistor voltage
0x2A	vfs	DAQ setting for pre-amplifier feed-back resistor voltage
0x2B	sample_div	Phase and period of the <code>sample_clk</code> in units of clock cycles
0x2C	fmdd_cmd	Command register
0x2D	cal_iter	Number of triggers per strip per calibration pulse size used in automated gain scans

Table 5.4: FMDD specific registers in the BC

Address	Name	Description
Commands inherited from TPC BC		
0x16	cntlat	Latch (store) trigger counters in registers
0x17	cntclr	Clear trigger counters
0x18	csr1clr	Clear error and status register
0x19	alrst	Reset all ALTR0s on FMDD
0x1A	bcrst	Reset BC
0x1B	stcnv	Single shot monitor ADC run
Commands specific to the FMD BC		
0x1	change_dac	Change DAC outputs
0x2	self_trig	Self trigger
0x4	l0_trig	Make l0 trigger
0x8	soft_reset	Reset state machines
0x10	cal_on	Enable calibration pulser
0x20	cal_off	Disable calibration pulser
0x40	test_on	Enable VA1 ₃ test mode
0x80	test_off	Disable VA1 ₃ test mode
0x100	start_cal	Start pulser calibration mode

Table 5.5: Commands implemented in the FMDD BC. Note, that FMD specific commands is executed by writing a value to the register `fmdd_cmd`.



Name	Description	Default
t_b	Time from bunch-crossing to 10	1.2 μ s
$t_{\overline{11}}$	Time from bunch-crossing to $\overline{11}$	6.5 μ s
$t_{\overline{12}}$	Time from bunch-crossing to $\overline{12}$	88 μ s

Figure 5.27: Timing of triggers

asserting $\overline{\text{hold}}$ as described in Section 5.2.1.

Upon receiving a $\overline{11}$ trigger, the trigger handler asks the VA1₃ interface to read-out the VA1₃ channels (`seq_start`), and start the ALTRO digitisation. It then waits for the VA1₃ interface to finish its job (`seq_busy`), and returns to the idle state, ready to accept more triggers (see also Figure 5.28).

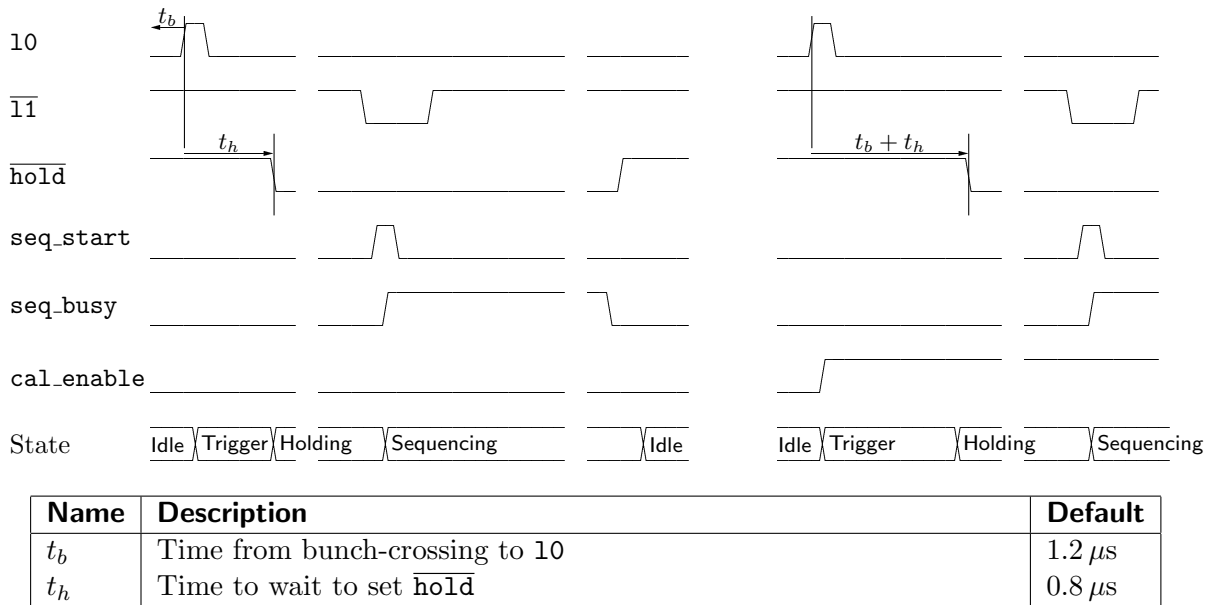


Figure 5.28: State changes in the trigger handling module. The left side illustrates normal trigger conditions, where the trigger is received 1.2 μ s after the collision. The right hand side illustrates the case of pulser calibrations, where an additional delay of t_b is introduced before asserting $\overline{\text{hold}}$.

If a $\overline{11}$ trigger is not received within a configurable time-out after reception of 10, the module returns to the idle state, and the read-out is aborted. If a $\overline{12}$ received before the VA1₃ interface module has finished, the event is aborted, the FSMs reset, and a flag set in the `fmdd_status` register.

The delay from receiving the 10 trigger to setting $\overline{\text{hold}}$, the $\overline{11}$ and $\overline{12}$ time-outs are all configurable via registers (see Table 5.4).

When in pulser calibration mode, then upon receiving 10, a pulse is made on the output `cal_enable`. This signal enables a voltage step on the `cal` inputs to the VA1₃ chips, thereby inducing a charge on the pre-amplifier input. In this mode, the trigger module waits an additional

$t_b \approx 1.2 \mu\text{s}$ before issuing hold, to ensure that the VA1₃ shaper has reached its peak value (see also right-hand side of Figure 5.28).

The trigger manager follows the following algorithm when in calibration mode

```

for  $i$  in range strip_first to strip_last do
  Set VA13 cursor to point at VA13 channel  $i$ 
  Set current pulse size cal_pulse = 0
  while cal_pulse < 256 do
    Set iteration counter  $j$  = 0
    while  $j$  < cal_iter do
      Wait for trigger
      Fire pulser
      Increment iteration counter  $j$  by 1
    done
    Step current pulser size with cal_step
  done
done
Leave calibration mode

```

In this way, the BC firmware will for each input channel on the VA1₃'s record `cal_iter` events with known inputs `cal_pulse` = 0, `cal_step`, `2cal_step`, ... In total that makes

$$N_{cal} = (\text{strip_last} - \text{strip_first} + 1) \times \text{cal_iter} \times 256 / \text{cal_step} \quad (5.7)$$

events needed. For the default settings, where all 128 strips are activated, `cal_step` = 32, and `cal_iter` = 100 events/pulser/strip, it adds up to 102 400 events needed to do a full scan. If the parameter settings are known, then data recorded can later be analysed to extract a gain for each VA1₃ input channel. The mean ADC counts recorded for each calibration pulse size is plotted against pulser size, and a straight line is fitted to the data points.

The parameters `strip_first` and `strip_last` are the upper and lower 8bits of the register `strips`, and `cal_step` is the upper 8bits of the register `cal_pulse`. `cal_iter` is a separate 16bit register (see also Table 5.4). The BC will automatically leave the calibration mode once the scan has finished, and return to normal operations. The only other way to leave calibration mode is to reset the entire FMDD card. This means, to complete a scan the FMD must be guaranteed a certain minimum number of events, as given by (5.7).

Monitor ADC interface and Safety Module This component reads the 4 monitor ADCs on the card. Each ADC has 5 10bit channels, of which one is an internal temperature. The remaining four channels are set up to measure various currents, voltages, and temperatures on the FMDD. The BC will read all 20 monitor ADC channels via an internal I²C bus, in a round-robin fashion¹⁰, and store the result in registers (see Table 5.6). The time to read-out all 20 ADC channels is about a few fractions of a second. All the monitor registers are available over the I²C bus, meaning that the state of the FMDD cards can be monitored even when data is sent of the ALTR0 bus.

Associated with each monitor register, is another threshold register. The threshold register defines an upper (or lower, depending what is monitored) limit for the operation of the FMDD. If a monitored value falls outside the range defined by the threshold register, then the BC firmware will assert the line `interrupt` on the ALTR0, alerting the RCU that something is amiss. The BC will continue operation as well as possible, and it is up to the RCU (or latter stages in the control environment) to take appropriate action.

¹⁰Read through all 20 in turn, then return to the first and read all 20 again.

Value			Threshold			Description
Address	Name	Value	Address	Name	Value	
0x06	t1	-	0x01	t1.th	40	ADC Temperature
0x07	flash_i	-	0x02	flash_i.th	-	Flash current
0x08	al_dig_i	-	0x03	al_dig_i.th	-	ALTRO digital current
0x09	al_ana_i	-	0x04	al_ana_i.th	-	ALTRO analogue current
0x0A	va_rec_ip	-	0x05	va_rec_ip.th	-	VA1 ₃ receiver current
0x32	t2	-	0x2D	t2.th	40	ADC Temperature
0x33	va_sup_ip	-	0x2E	va_sup_ip.th	-	VA1 ₃ supply current
0x34	va_rec_im	-	0x2F	va_rec_im.th	-	VA1 ₃ receiver current
0x35	va_sup_im	-	0x30	va_sup_im.th	-	VA1 ₃ supply current
0x36	flash_u	3.3	0x31	flash_u.th	3.2	Flash voltage
0x3C	t3	-	0x37	t3.th	40	ADC Temperature
0x3D	t1sens	-	0x38	t1sens.th	40	Temperature sensor
0x3E	t2sens	-	0x39	t2sens.th	40	Temperature sensor
0x3F	al_dig_u	2.5	0x3A	al_dig_u.th	2.4	ALTRO digital voltage
0x40	al_ana_u	2.5	0x3B	al_ana_u.th	2.4	ALTRO analogue voltage
0x46	t4	-	0x41	t4.th	40	ADC Temperature
0x47	va_rec_up	2.5	0x42	va_rec_up.th	2.4	VA1 ₃ receiver voltage
0x48	va_sup_up	1.5	0x43	va_sup_up.th	1.4	VA1 ₃ supply voltage
0x49	va_sup_um	-2.0	0x44	va_sup_um.th	-1.8	VA1 ₃ supply voltage
0x4A	va_rec_um	-2.0	0x45	va_rec_um.th	-1.8	VA1 ₃ receiver voltage

Table 5.6: Monitor ADC related registers.

Implementation, Performance, and Testing As mentioned above, the FMDD BC firmware¹¹ is based in the TPC BC Verilog code base. However, the FMDD BC was completely re-written in VHDL. This re-write was done to implement *test-benches* of all components in the firmware. A test-benches is a list of well-known specified input to an module, and checks on the expected behaviour of the module, implemented in VHDL too. In this way, one can simulate the behaviour of a single sub-module under various conditions, including error conditions and unexpected input.

Proper behaviour of all the FMDD BC firmware components have been verified using this technique, starting with the smallest elements and working up the hierarchy. A special module was implemented to allow simulated control via the standard FMD test tool `FmdGui`. A limited test-bench of the full FMDD card was also implemented, including the 3 ALTROs, the 4 monitor ADCS, and the digital behaviour of the attached VA1₃'s.

5.3.3 ALTRO bus protocol

All components on the FMDD are accessible via the custom ALTRO bus. It consists of the 40bit data lines `bd` and the control signals `write`, `cstb`, `ackn`, `dstb`, and `trsf`. The bus is used for both register input/output as well as data transfer from the ALTROs to the RCU. The bus protocol is based on an asynchronous hand shake of the control strobe `cstb` and acknowledge `ackn`.

When doing register input/output or command execution, then the upper 20bits of the `bd` are reserved for addressing and instructions. The 20 upper bits are interpreted as shown in Table 5.7.

Parity Bit for checking the transfer. The bus has *odd* parity, meaning that the bitwise XOR of all 20bits must evaluate to 1.

Broadcast For register input (to the FMDD components), if this bit is enabled, all components should be considered addressed. In this case, the rest of the address is ignored.

¹¹The full source of the firmware is available from http://fmd.nbi.dk/fmd/fee/fmdd_bc.

Target/Bits	39	38	37	36	32	31	29	28	25	24	20
ALTR0	Parity	Broadcast	0	Board	Chip	Channel	Instruction				
BC	Parity	Broadcast	1	Board	Instruction						

Table 5.7: Meaning of upper 20bits of *bd* during register input/output and command execution.

Bit 37 Component select. If this bit is ‘1’, then only BCs should consider the bus. If it is ‘0’, only the ALTR0s should consider themselves addressed, though the BCs may “listen-in”.

Board Board to address. Each FMDD has a hard-coded address (0 for the inner boards, and 1 for the outer boards) that must be unique on the bus. Even though the field is 5 bits wide, only 4 of them are considered by the FMDD components. Bit 36 is used only by the RCU when selecting branch (see Section 5.4).

Chip (ALTR0 only) chip on the board to address. Since there are 3 ALTR0s on the FMDD, valid addresses are 0, 1, and 2. The address of each ALTR0 is hard-wired on the board.

Channel (ALTR0 only) Some registers in the ALTR0 chips are channel specific. If the instruction is for a channel specific register or command, this field will tell which channel to address. If the instruction is not for a channel specific register or command, then this field is ignored.

Instruction The register or command to address. For the ALTR0s this is 5bits wide, while it can be up to 12bit wide for the BC, however only 7bits are used at this time.

Register input (to the FMDD components) is initiated by the RCU by asserting $\overline{\text{cstb}}$ and $\overline{\text{write}}$, setting the component and register address on the upper 20bits of *bd*, and the value on the lower 20bits of *bd*. When a component recognises an instruction it asserts $\overline{\text{ackn}}$, and the RCU removes the assertion of $\overline{\text{cstb}}$ and $\overline{\text{write}}$, and all 40bits on *bd* (see also Figure 5.29). If the RCU does not receive an $\overline{\text{ackn}}$ within a fixed number of *clk* cycles after $\overline{\text{cstb}}$ is asserted, the write times out and fails. It is possible to write in *broadcast* — that is, to all possible components on the bus. In that case, bit 38 of *bd* is asserted, and no acknowledge is asserted nor expected.

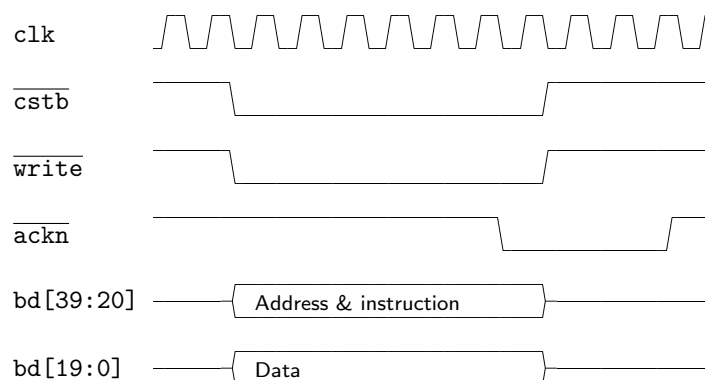


Figure 5.29: Write instruction sequence on the ALTR0 bus. Command execution is the same, except that the lower 20 bits of *bd* is left unspecified.

Commands are done in the same way as register writes, except that the lower 20bits of *bd* is not set. Commands can also be broadcasted.

Register output (from the FMDD components) is initiated by the RCU by asserting $\overline{\text{cstb}}$ and *not* asserting $\overline{\text{write}}$. The address and instruction code is set in the upper 20bits of bd , and the lower 20bits are left alone. The addressed component sets the return data on the lower 20bits and asserts $\overline{\text{ackn}}$. If no component asserts $\overline{\text{ackn}}$ within a fixed number of clk cycles after $\overline{\text{cstb}}$ is asserted, the read times out and fails. Registers *cannot* be read in broadcast.

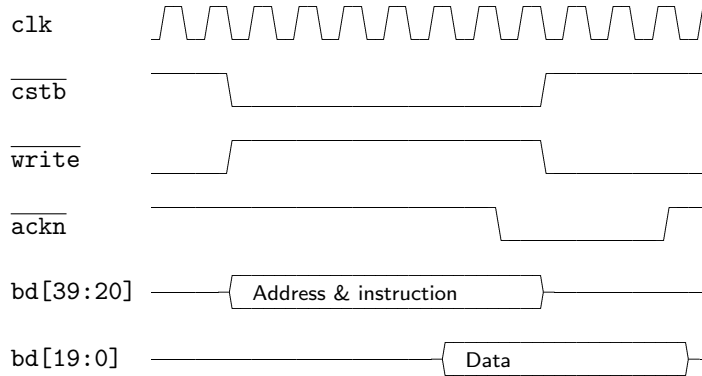


Figure 5.30: Read instruction sequence on the ALTR0 bus

The line $\overline{\text{error}}$ will be asserted, for both register input/output and command execution, if an illegal instruction is seen by any of the FMDD components.

5.4 The Read-out Controller Unit

The choice of the ALTR0 for the FMD ADC chip dictates a number of choices for the FMD front-end electronics. One, is that the implementation of the FMDD is based on the TPC front-end card, as described in the previous section. Another is the use of the TPC *Read-out Controller Unit* (RCU) [89]. The FMD uses three standard TPC RCU's with no modifications, one for each sub-detector system (FMD1, FMD2, and FMD3). The RCU card are located around 3 m from its corresponding detector, and connected to the FMDDs of each sub-detector using twisted-pair bus cables as outlined in Section 5.3. The RCU card consist of a mother board and two daughter cards, the *Source Interface Unit* (SIU) and *Detector Control System Card* (DCSC). The mother board has a relatively large FPGA mounted to implement the functionality of the card. A schematic representation of the architecture of the RCU FPGA (simply called the RCU in the following) is shown in Figure 5.31.

The RCU is responsible for

- Handle triggers $\overline{\text{I1}}$ and $\overline{\text{I2}}$ send over the optical *Trigger and Timing Control* (TTC) interface from the local trigger unit.
- Manage read-out of all ALTR0 channels.
- Push read-out data over the optical *Detector Data Link* (DDL) to the data acquisition system.
- Provide control of the FMDD cards attached to the RCU.
- Monitor the state of the attached FMDD cards.

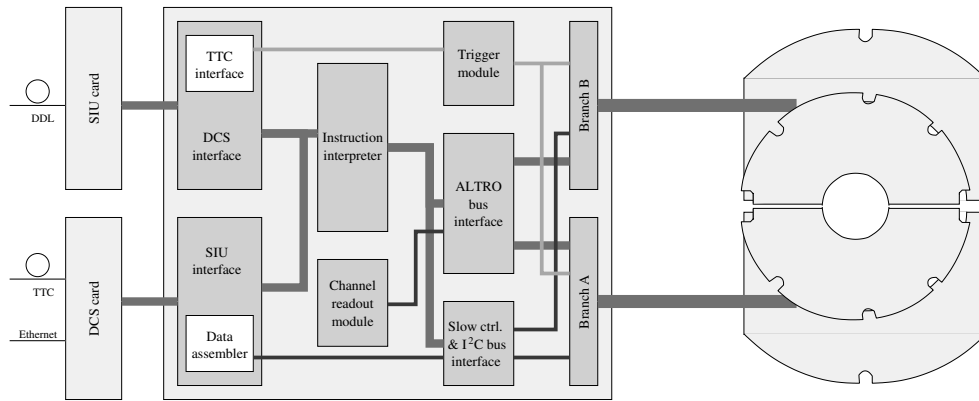


Figure 5.31: Schematic architecture of the RCU firmware with peripheral devices indicated.

Branches As shown in Figure 5.31, the RCU card has two branches, branch A and branch B. Each branch corresponds to a physical bus cable connector on the back of the RCU card, and each branch can handle up to 16 FMDD cards. It is important to note, that the branches are independent. That is, the branches can be operated by the firmware independently of each other. The triggers are the only exception — they are always send in parallel on both branches.

FMDD cards attached to branch A have a base address of 0, while cards attached to branch B have a base address of 16 (or $0x10$). It is however, only the RCU firmware that uses this base address, and the FMDDs will only ever see the lower 4bits of the address (corresponding to address in the range from $0x0$ to $0xF$). That in turn means, that cards can be switched from one branch to another with no change needed.

Individual lines on the bus go from the RCU to the FMDD cards *card switch* (`cardsw`), to enable the RCU to turn cards on and off individually. Which cards to turn on and off are set in the 32bit bit-field register `af1`¹²

The trigger handling is managed by a separate module in the RCU firmware. Aspects, such as trigger types and similar can be configured via registers in the firmware.

Data transfer from ALTR0s to the RCU When a $\overline{12}$ arrives at the RCU, validating a previous $\overline{11}$, the RCU will at the first possible time, ask each configured ALTR0 channel to hand over its data (see also Section 5.3.1). Which channels that should be read-out by the RCU is configured via the RCU `ac1`¹³ memory. The data is then stored in one of the 4 or 8 data buffers available on the RCU, and when all data is available, a *Common Data Header* (CDH) is prepended to the data block, and the data is pushed to the data acquisition system via the SIU and optical DDL fibre.

In a future version of the RCU firmware, the RCU will generate a special data-block at the start of a run — the so-called *Start-of-Data* (SOD) event. This event will contain parameters of the front-end. Exactly which parameters will be configurable via an RCU memory block. The data can for example contain the `sample_clk` and `shift_clk` frequency so that code that analyses the data can extract the oversampling rate directly from the data. Exactly how this will work remains to be specified and implemented by the CERN TPC electronics group.

Note, that transfer of the ALTR0 data to the RCU event buffer does not necessarily happen immediately after reception of $\overline{12}$. If a new $\overline{11}$ is received shortly after, the transfer is postponed.

Trigger handling and Busy box The trigger handler of the RCU firmware handles trigger data arriving from the TTC interface on the daughter DCSC, as well as take care of internally generated

¹² Active Front-end card List — also known as `actfec`.

¹³ Active Channel List

triggers — via commands or via the data acquisition SIU interface. Future trigger handling modules will be more configurable with regards to trigger masks and what triggers are passed on to the FMDD.

There is no protection against overflow of the event buffers — that is, if all 4 or 8 event buffers are filled no flag is raised to the out-side world, and further $\overline{12}$ triggers will result in a crash of the RCU firmware. To remedy this situation, a separate module in the data acquisition system was designed, the so-called *Busy-box*. The busy box keeps track of how many triggers were sent to the RCU and how many events the RCU has sent back. If the busy box sees that all 4 or 8 buffers are full in the RCU, it raises a busy flag to the central trigger to let it know that this RCU cannot handle any more events at this time. As soon as 1 or more buffers becomes free on the RCU, by the RCU pushing an event the data acquisition system, the busy box removes the busy flag, and the central trigger can send triggers to the RCU again.

Slow Control, Monitor, and I²C Interface When the interrupt mode (RCU register `intmod`) is turned, then this module will constantly monitor the BC status register `csr1` on enabled cards. The module reads this register (and in the future, possibly other registers) via the I²C bus, leaving the ALTR0 bus free to do data transfer from the ALTR0s to the RCU.

The BCs can, as explained earlier raise the `interrupt` line if a voltage falls under threshold, a current goes over threshold, or the temperature on the card is too high. If any card raises the `interrupt` line on the bus, then the monitor module will turn off that card, to ensure that the card is not damaged. A copy of the `csr1` register of the turned off card, along with the address of the card is put in the `status` memory of the RCU. The *Detector Control System* (DCS) can monitor this memory to provide status information to the human operator, or other monitoring systems. For example, DCS should, when a card is turned off due to an interrupt, turn off the bias voltage to the associated silicon sensors.

Instruction Interpreter Central to the RCU firmware is the instruction interpreter. It consists of an instruction memory `imem`, and a processing unit (PU). Code can be written to the `imem`, and executed by the processing unit via the command `exec`. The PU defines instructions to access all registers and commands of the RCU, and via the ALTR0 bus, all register and commands of the FMDD BCs, and ALTR0s. Binary code can be uploaded to the RCU `imem` via either the DCSC or SIU interface. This makes the RCU a very flexible entity, in that one can program the chip for specific behaviour.

The code uploaded to the `imem` is binary code — not easily readable and must adhere to the PU instruction set. To ease writing the `imem` code, a separate programming language and compiler was developed¹⁴. This programming language represents the code in a semi-human readable form. An example of the human-readable code and the compiled code is shown in Figure 5.32.

The compiler was used by the FMD group to verify the behaviour of the RCU and to let the group familiarise themselves with the inner workings of the RCU. In the future, the compiler may be part of specialised programs for configuring the RCUs, BCs, and ALTR0s used by the FMD.

5.4.1 DCS Daughter Card

As mentioned earlier, communication from the outside world to the RCU, and by extension the FMDD components, can go either via the SIU card of the DCSC daughter card (DCSC). The SIU interface is however used primarily for data transfer and therefore not used for slow-control of the RCU and attached FMDDs — the DCSC is used for that. The DCSC was developed at University of Heidelberg.

¹⁴The compiler and specification is available from <http://fmd.nbi.dk/fmd/fee/altrocc>.

<pre> 1 GROUP <i>din_zsp</i> : ALTRO write 1st path broadcast 1st baseline 0x7 zero suppression 0 0 0 2 ALTRO write zero suppression broadcast 0 1 <i>trigger_setup</i> : 3 ALTRO write trigger config broadcast 0 0x100 4 END </pre>	<pre> 1 0x0064000b 0x00780007 2 0x00640008 0x00700001 3 0x0064000a 0x00700100 4 0x00390000 </pre>
---	---

Figure 5.32: Example of a program for the RCU instruction interpreter. On the left is the semi-human-readable code, and on the right, the output of the compiler that can be interpreted by the RCU PU.

The main features of the DCSC card is an FPGA with an embedded processor. The processor runs a small version of the GNU/Linux operating system, and the FPGA implements peripheral devices that the operating system can communicate with. Exactly which peripheral devices are implemented depends on the use of the DCSC, but for the FMD (and other detectors that use the RCU and ALTRO chips) these include

- Ethernet connection.
- TTC interface.
- Direct write access to RCU FPGAs configuration memory.
- RCU bus interface.

The Ethernet connection on the DCSC card allows the GNU/Linux box to communicate to the outside world using standard Internet protocols, such as Secure Shell (SSH), Hyper Text Transfer Protocol (HTTP), and so on. In particular, the CERN developed protocol Distributed Information Management (DIM) [90–92] is used extensively by a server (the *FeeServer* — see later) running inside the GNU/Linux box.

The TTC interface provides an input for triggers and timing information from the general ALICE trigger and timing system via an optical fibre. On this interface, the 40 MHz clock, synchronised over all of the LHC, arrives and is distributed to the RCU and onward. Trigger signals arrive here too, and are likewise distributed to the RCU which then may or may not, depending on the RCUs configuration, distribute these further to the FMDDs.

The direct write access to the RCU FPGAs configuration memory allows the GNU/Linux box to re-program the RCU at any given time. In fact, when the GNU/Linux box boots up at power-up, one of its first chores is to program the RCU FPGA.

Via a kernel driver and an interface library, processes running in the GNU/Linux box can communicate with the RCU interface, and through that with the RCU and all associated devices like the FMDD BC and ALTROs. The interface is based on *memory mapping* i.e., the memory of the RCU is mapped into the kernel of the GNU/Linux operating system running in the DCSC FPGA.

Review

This concludes the discussion on the silicon sensors and the front-end electronics of the FMD. We have seen how the sensor detects particles, and the signal generated by this, which is then amplified and shaped by the VA1₃'s on the hybrid cards. The FMDD BC then synchronises the operation of transferring the signal to the ALTR0 analogue-to-digital converters which digitised and buffers the data. Once the event has been validated by the trigger system, the data is retrieved from the ALTR0s by the RCU, stored in buffers with proper headers, and eventually pushed out to the data acquisition system. All behaviour of the front-end is configurable via registers in the various components — the RCU, BCs, and ALTR0. Control and configuration of these components goes via a special server running in the embedded GNU/Linux box running on the RCU DCSC daughter card. The next chapter discusses the data acquisition system and the detector control system — the final link to the human operator of the experiment.

Chapter 6

Data Acquisition and Control

Data acquisition is the process of storing the digital data from the detectors, and to make sure that the proper booking is done so that later stages can reconstruct the events and analyse the data. The detector control system takes care of control the behaviour of the various detector components in the ALICE experiment.

6.1 Data Acquisition

Once the data has left the front-end electronics it is up to the data acquisition system (DAQ) to process and store the data. DATE, the ALICE DAQ, developed at CERN, builds on DIM [90, 92] for distribution of tasks, and *State Management Interface* (SMI) [93] for handling the various finite state machines of the DAQ. DATE is a three tier system, as shown in Figure 6.1.

Close to the front-end electronics of the various detectors is the *Source Interface Unit* (SIU), each linked to a *Destination Interface Unit* (DIU) via the optical DDL (*Detector Data Link*) fibre. The DIU sit on a RORC standard PCI card in a commodity computer. These computers are called the LDCs (*Local Data Concentrator*) and run a standard GNU/Linux operating system. The data of a detector or a partition of the a detector, arrives at the LDC and is collected there into a sub-event. The LDC is the first tier of the DAQ.

All LDCs are connected up to a farm of GDCs (*Global Data Concentrators*), which constitute the second tier. When possible, the LDCs ship their data to the GDC farm, where the sub-events of each LDC is collected into global events. The global events are then send to permanent data storage (the PDC) at CERN — the third and final tier.

Optionally, the LDCs can be run in stand-alone mode. In this mode, they do not hand off data to the GDCs and do not interfere with the global run. This mode can be used by detector experts to debug possible problem during normal data taking without impact on the rest of the detectors.

For special kinds of runs, the data for each LDC can be stored locally on the LDC and specialised programs can be executed on these runs. These special programs are called DAs (*Detector Algorithms*). DAs can produce data which a special program, called the *shuttle* can pick up and publish on the Internet for latter use.

For the FMD three such DAs have been designed — one for analysing pedestal data, one for analysing pulser data, and one for extracting data from the *Start of Data* (SOD) event. Both the pedestal and pulser DA will extract parameters important to the analysis — like the oversampling rate, number of events per strip per pulse size, etc. — directly from the SOD event, and is in this way completely independent of the external data.

Pedestal data Ever so often, about once or twice a day while running, special runs where the LHC will not provide collisions will be recorded. Since there are no collisions, the data read-out is what is known as *pedestal data* — that is, the detector response as it in the absence of

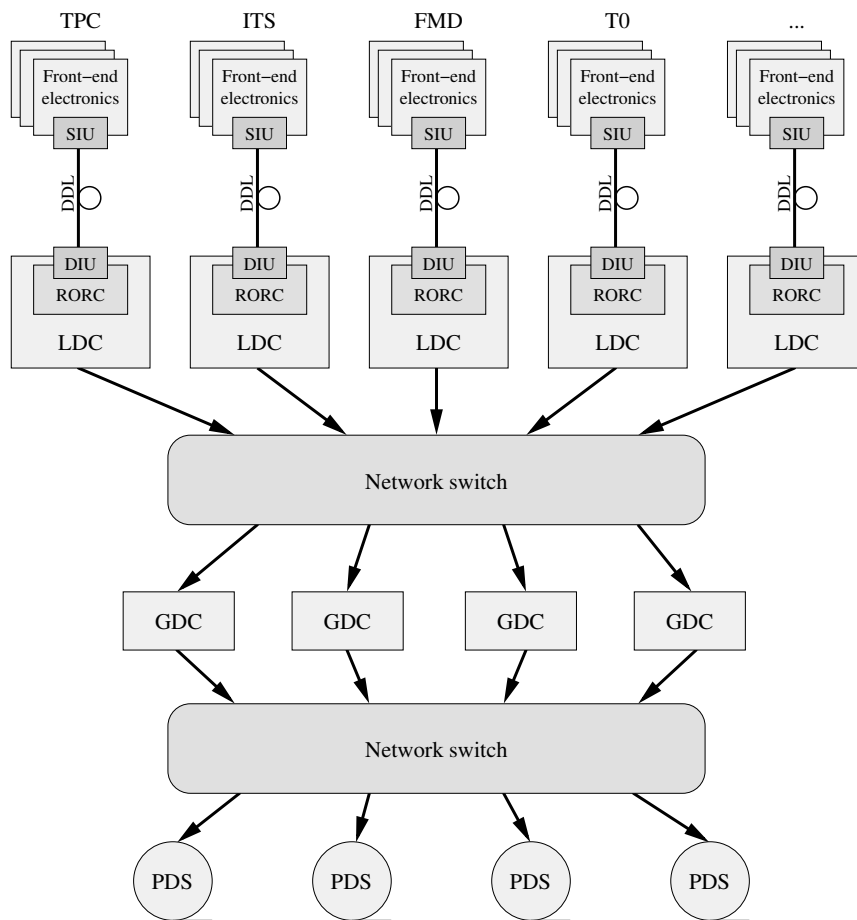


Figure 6.1: The three tier DAQ architecture.

signals. The FMD DA for these kind of runs extracts the base-line ADC values by calculating for each strip, the mean of the ADC values recorded during a pedestal run. The variance of the ADC distribution is likewise calculated and stored. The result of the analysis is copied to an intermediate storage where it will latter be read by the off-line shuttle and published on the Internet.

Pulser data As explained in Section 5.3.2, the front-end electronics of the FMDD can be put into a mode where a pulser is used to scan the gain of each VA1₃ input channel. After such data has been recorded, the FMD DA for *pulser runs* will be run. This will, for each pulse height find the mean and variance of the data, and then fit a straight line to the mean versus the pulse size. The slope of this fitted line is the gain of the input channel, and each is written to intermediate storage where they are picked up and published by the off-line shuttle.

Front-end run parameters This detector algorithm will extract important front-end electronics parameters from the RCU generated SOD event. To reconstruct the data, the off-line code needs to know the oversampling, which is simply the ratio between the `sample_clk` and the `shift_clk`. In case the ALTR0s do zero-suppression on the data, the parameters of the zero-suppression filter are also needed. This DA will therefore extract these parameters from the SOD event, put them on intermediate storage, for the shuttle to pick up and post on the Internet.

6.1.1 Event Format

ALICE events are hierarchically structured with detector data (Payload) at the bottom and a super-event at the top, as shown in Figure 6.2.

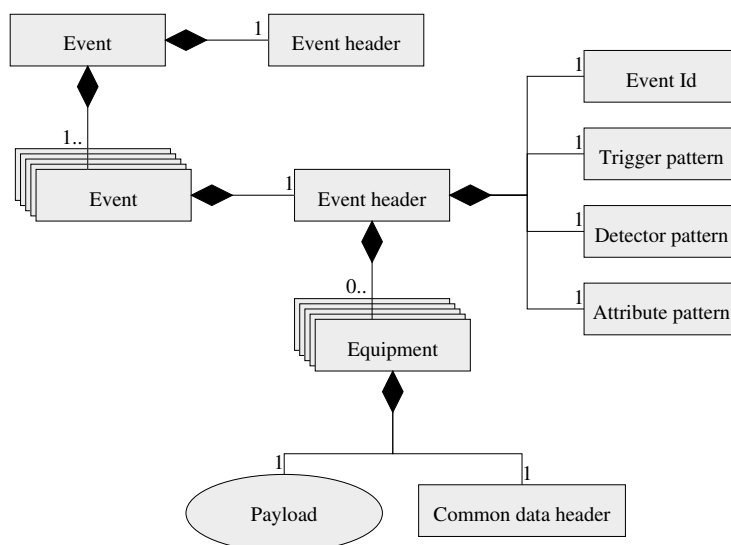


Figure 6.2: The global event structure.

The super-event contains a header that identifies the event, participating detectors, triggers, and other attributes defined by the DAQ system. The super-event contains sub-events, as produced by the LDC, again with their own header information. The sub-event headers each have a list of *equipments*.

An *equipment* is a single entry into the DAQ system, like for example a DDL line¹. Each equipment entry contain a common data header which contains timing and identification data for use by the GDCs to assembly the super-events, and to identify the data in later reconstruction

¹Other routes are VME and logical devices like the trigger and DAQ systems

passes. The equipment finally contains the actual detector data as payload. How the payload is, as the name implies, irrelevant to the DAQ system and depends entirely on the detector. For the FMD the payload is the RCU concatenated ALTR0 channel data from the ALTR0s attached to the RCU in question. Since each RCU has one DDL and there is a total of 3 RCUs (one for each sub-detector), the FMD data will be in three equipment structures.

Given the equipment identifier, the ALTR0 channel address, the sample number, as well as the oversampling rate, the corresponding physical strip in the FMD system can be identified and mapped into a physical location in terms of (φ, η) or (x, y, z) if so desired.

6.2 Detector Control System

The *Detector Control System*² (DCS) takes care of various detector tasks that are not directly related to data taking. It is the system through which the human operator can configure and control the front-end electronics, voltages, cooling, and so on, of the detectors themselves. It is also the system that provides monitoring of the state of the detectors, including temperatures, voltages, currents, pressure, and errors in the front-end of the detectors.

The ALICE DCS is based on the commercial SCADA³ system PVSSII together with a CERN developed framework for defining the detector control. DIM and SMI are also used extensively throughout the system for communication and finite state machine definitions and management.

The ALICE DCS is hierarchically structured. At the bottom there is the interface to the various pieces of hardware in the system, like front-end electronics, power supplies, valves, pressure sensors, and so on. At intermediate levels one finds components that represent a full sub-system of ALICE, be it a detector like the FMD or the cooling plant that supplies all of ALICE with cooling water. At the very top there is a general supervisory layer which the human operator interacts with. Of course human interaction at lower levels of the hierarchy is possible, but is restricted to experts of the particular system only.

Parallel to this logical and hardware hierarchy is a *Finite State Machine* (FSM) hierarchy. That means, at the very bottom, there are FSMs that take care of hardware specific chores. For example, a high voltage system may require a careful and slow increase in the (negative) voltage, or a sub-system may critically depend on the temperature be within a certain range, and so on. In any case, the FSM is expected to properly change state when changes in the hardware mandates it, or possibly at the higher levels request.

At the intermediate level, there are FSMs for full subsystems. These have a general structure that allows the top-level to ask subsystems to do certain well defined tasks, like configure your hardware, get ready for collisions, abort all operations, and so on.

6.2.1 Architecture

The FMD DCS consist of 4 hardware components:

- Low-voltage power supplies for the front-end electronics.
- Power supplies for the silicon reverse bias voltage.
- Control and configuration of the front-end electronics.
- Cooling system for the front-end electronics.

This is illustrated in Figure 6.3.

²Also known as *slow control*.

³Supervisory Control And Data Acquisition. Note, that in this context, data acquisition does not refer to collection of the primary physics data, but rather the monitored data like temperature, voltages, currents, pressure, and so on.




Figure 6.3: The FMD DCS architecture.

Also shown in Figure 6.3 are the 3 separate computers in the system. `alifmdwn001` and `alifmdwn002` are the worker-nodes, labelled for the obvious reason that they carry out the actual work, meaning running services, clients, and the control finite state machines. `alifmdon001` is the operator node. The operator node is where the human operator logs in to do FMD specific chores. It is merely a client of the worker nodes and does not execute other programs than a graphical user interface (GUI).

The FMD DCS does not directly control the cooling plant. Indeed, only the status of the cooling plant is monitored by the FMD DCS system, and no more will be said about this component.

6.2.2 Low and bias voltage system

The FMD uses CAEN power supplies for the low and bias voltages. The CAEN low voltage system was chosen because of its high radiation tolerance. The low voltage modules sit in the experimental hall and must therefore sustain more radiation than usual equipment. The choice of CAEN bias voltage supply was dictated mostly by uniformity and coherence of the full power supply system.

Both the low and bias voltage supplies are controlled and monitored through an OPC⁴ server running in the CAEN control create. The PVSS project running on the first worker node `alifmdwn001` is a client of this server and communicates over a standard computer network.

It is possible in both kinds of power supplies to set thresholds, trip (failure) conditions, and so on, directly in the hardware. The OPC server will then inform the client (PVSS) if such a condition should arise. The FSM of the PVSS will then respond by changing state accordingly.

⁴Originally OLE for Process Control, where OLE is *Object-Linking and Embedding* as developed by Microsoft — however, officially OPC is no longer an acronym.

6.2.3 Front–end Electronics Chain

The architecture of the front–end electronics detector control system is based on the TPC system, but heavily modified to suit the needs of the FMD. Starting near the detector is the FeeServer running in the embedded GNU/Linux box on the *Detector Control System Card* (DCSC). It is connected to the *Inter Com Layer* (ICL) running on `alifmdwn002`. All FeeServers connect to one and the same ICL, which then propagates monitoring information up to PVSS and forwards the control protocol to the FeeServers. This is illustrated in Figure 6.4. Communication between all components of the system is via normal Ethernet, using the DIM protocol.

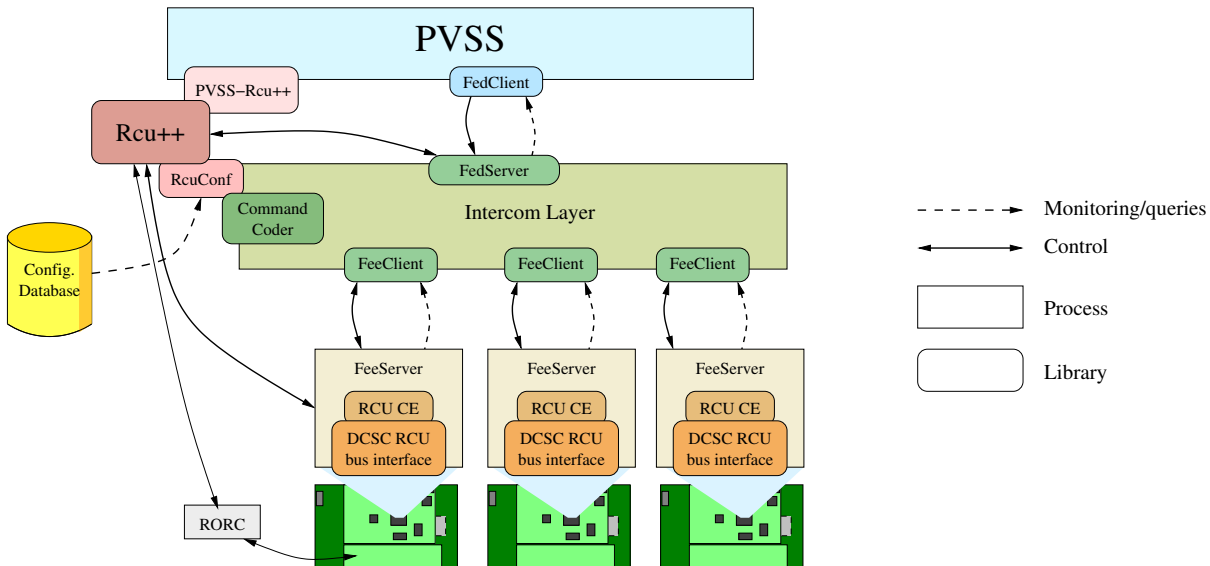


Figure 6.4: The FMD DCS software architecture.

GNU/Linux box on the DCSC The small processing unit in the FPGA on the DCSC is not an omnipotent, general–purpose kind of CPU as we know them from our workstations and laptops. It is more akin to the processing units that we find in most mobile phones these days. However, despite its short–comings as desktop machine, the GNU/Linux box does take care of crucial services in the detector control system of the front–end electronics.

At start up, or on request, the GNU/Linux box will load the RCU firmware into the RCU FPGA. Without this firmware, the chip would be nothing but inactive logical gates, and no triggers would arrive at the front–end, nor any data in the DAQ. The firmware is read from a disk mounted over the network from `alifmdwn002`, meaning that upgrades to the firmware are easily managed for all of the FMDD.

Another crucial service that the GNU/Linux box runs, is the FeeServer. This service is started at boot–time, and as for the RCU firmware, the executable image is read over the network to ease upgrades to this vital service.

The FeeServer This server, running in the embedded GNU/Linux box on the DCSC consists of four components: The core FeeServer, an interface library to the RCU bus kernel driver, RCU Control Engine (RcuCE), and a factory of front–end card representations.

The core FeeServer itself is independent of the hardware and detector it is meant to control. It provides the outside communication protocol over DIM, implemented as three channels: A command channel for sending commands to the various components of the FeeServer, an acknowledge channel to return the results of commands, reads, and writes, and finally a message channel for logging. Furthermore, the core provides mechanism for publishing other DIM services, as well as an abstract interface to a user defined control engine.

To guarantee regular updates of services and low-latency response to commands, the core FeeServer is implemented in a number of separate suspendable threads. One thread handles all commands sent to the server components, while another thread implements a watchdog on the message channel to keep the volume of repeated messages down. A third thread is meant for the user to execute queries to the hardware, and finally a fourth thread updates service values as seen by the outside world. The idea is, that even when the server is executing complex commands, or doing slow reads from the hardware, it should still be possible for the clients of the server to get regular updates, or to channel operations in case of dead-locks, and so on.

The original FeeServer, implemented in C, core was developed and designed at Zentrum für Technologietransfer und Telekommunikation (ZTT) at Worms University of Applied Sciences, Germany. However, the FeeServer core for the FMD was completely rewritten in C++ to simplify the implementation and ease the configuration for the FMD. The newly implemented FeeServer adheres to the protocol and behaviour of the original FeeServer from Worms, but is very different in the actual implementation.

RcuCE is an implementation of the control engine that accesses the RCU via the DCSC interface library. The interface library provides low-level access to the memory mapped register, memories, and commands of the RCU. The RcuCE provides the outside world with a number of general purpose and specialised commands for accessing the various registers, memories, and commands of the RCU. The result of these RcuCE commands are sent back to the caller over the FeeServer acknowledge channel. In this way, clients of the server, like the intercom layer and Rcu++, can have direct access to the hardware through the FeeServer.

The RcuCE monitors some registers and memories of the RCU:

- The *Error* and *Status* register (`errst`). This register bits of various error and status flags, including front-end card communication errors, ALTR0 bus busy, inconsistent read-out, and so on.
- The *Active front-end card list* (`af1`). This register is a bit-pattern of the currently active i.e., turned-on, front-end cards.
- The *Status* memory (`fec_status`). If a front-end card raises an interrupt, its current state is recorded in this memory. Up to 32 such lines can be present, corresponding to the total possible number of front-end cards that can be attached to a single RCU.

These values are read periodically from the front-end, and published to the outside world by using the core FeeServer functionality for this. The RcuCE also defines a simple finite state-machine to capture the state of the RCU while running. The current state is published as a service via the core FeeServer. The RcuCE will automatically change state when certain criteria are met, but can also be asked to change state explicitly. Exactly which states will be in the finite state machine and how they are reached is still an open question — one that we will answer as we accumulate more experience with the full system.

When ever the RcuCE detects that a new front-end card (FMDD) has been turned on, reflected in the corresponding bit in the RCU `af1` register, it will ask an abstract factory to create a representation for this card. Exactly how the front-end card is represented and which values, if any, should be published for the front-end card, is up to the user implementing the factory. The representation is likewise discarded, when the RcuCE detects that card has been turned off, ensuring consistency.

The original control engine for the RCU was designed and implemented by University of Bergen, Norway, but was revisited and reimplemented for the FMD. Again, this was done to extend its adaptability and simplify the code base.

The FMD implementation of the front-end card representation factory defines for each front-end card a number of services. All monitor ADC values (see Table 5.6) are published, as well as the configuration and status registers `csr0` and `csr1`. These registers are read periodically and

stored in the representation. When a value is changed, the new value is pushed to the clients of the server. At the time of writing, the exact frequency of updates and the number of services has not been finalised. As we gain more experience with the system, we will know better what is crucial and what is not. However, it is important to note that nothing is precluded at this time.

Combining these four parts⁵ — the core FeeServer, the DCSC interface library, the RcuCE, and the FMD front-end card representation factory — forms the FMD FeeServer, or simply the FeeServer. Note, that there are 3 such FeeServers running — one for each RCU i.e., one for each sub-detector.

The InterCom Layer To provide the higher levels of a single point of entry and to minimise the network load, all three FeeServers of the three sub-detectors FMD1, FMD2, and FMD3 are connected to a single *InterCom Layer* (ICL) process running on `alifmwn002`. Or rather, the ICL is a client of all three FeeServers. To the higher levels, the ICL implements a Front-end device server (FedServer). A FedServer is an abstract interface defined by the ALICE DCS group that all interfaces to the detector front-end electronics should implement (see Table 6.1). The interface is based on the DIM protocol and consist of two commands for configuring the front-end and the front-end chain it self, as well as two commands for accessing low-level commands of the front-end and the chain. The implementation should provide a special acknowledge channel for acknowledging low-level commands, and a message channel for logging. Optionally can the implementation define a number of services for monitoring the front-end electronics and the detector if so applicable.

Command	Parameters	Description
ConfigureFero	<code>char[20] target,</code> <code>int tag</code>	Ask the command coder to configure the electronics.
ConfigureFeeCom	<code>int,</code> <code>int,</code> <code>float,</code> <code>char[256]</code>	Configure the front-end chain (ICL and FeeServer).
ControlFero	<code>char[20] target,</code> <code>char* data</code>	Send low-level commands to electronics.
ControlFeeCom	<code>int,</code> <code>int,</code> <code>float,</code> <code>char[256]</code>	Send low-level commands to front-end chain (ICL and FeeServer)
Service	Parameters	Description
<code><icl_name_ACK></code>	<code>int status,</code> <code>char[256] msg</code>	Acknowledge of front-end chain commands.
<code><icl_name_MSG></code>	<code>int status,</code> <code>char[4] scope,</code> <code>char[256] where,</code> <code>char[256] msg,</code> <code>char[20] date</code>	Messages concerning the operation of the front-end chain (ICL and FeeServer) operations and state.

Table 6.1: Commands and services provided by the ICL. The `<icl_name>` is typically something like `ztt_dimfed_server`, but can be configured.

The FMD ICL is based on the ICL developed for the TPC, TRD, PHOS, and EMCAL by ZTT at Worms University of Applied Sciences, Germany, and adapted to the needs of the FMD.

⁵All parts can be downloaded from <http://fmd.nbi.dk/fmd/fee> where also more (technical) information is available.

Configuration and control of the front-end chain itself is entirely standard — meaning all levels can be configured for log-levels, timing, updates of frequencies and so on.

The Worms ICL simply forwards low-level commands to the front-end to the FeeServer, and amends the DCSs groups FedServer with specific acknowledge channels that are simply the FeeServer acknowledge channels forwarded. The FMD group takes advantage of this to implement a library Rcu++, that can communicate directly with the hardware through either the FeeServer or the FedServer, as indicated in Figure 6.4. In fact, this library can also communicate over the DAQ RORC interface as indicated on the figure. The Rcu++ library was used extensively in testing, and now forms the basis of an expert tool in the DCS for debugging and diagnosing problems in the FMD front-end.

The Worms ICL specifies that configuration of the front-end should be done based on look-ups in a database. It provides an abstract interface for a Command Coder that should be implemented by the detectors. The command coder is will get a *tag* that describes what we are configuring for (physics, pedestal calibrations, pulser calibrations, cosmic data, and so on), and is expected to return binary code for the FeeServer to execute. The ICL does not do anything to the binary code returned by the command coder, other than to pass it untouched to the addressed FeeServer.

Since the FMD group had already developed the Rcu++ tool for encoding binary data that the FeeServer can execute, it was natural to use that library for implementing our command coder. The RcuConf library provides the Rcu++ library with code to do look-ups in a relational database, and together they form the FMD command coder.

The ICL is, as outlined above, essentially a configurable switch-board. From the low (front-) end, it forwards services and messages to the higher levels. From the higher levels it forwards commands and configuration requests to the lower end.

PVSS The PVSS project running on `alifmdwn001` collects all the components described above into a single entity for controlling the FMD. It implements an OPC client for controlling the low and bias voltages, and a FedClient for controlling the front-end chain and electronics.

In the project there is a finite state machine for each component attached to the system. These are then controlled by a top-level finite state machine which constitutes the interface to the outside world — the human operator or the experimental control system (ECS). The state machine is sketched in Figure 6.5.

The diagram shows the possible states, and the action possible on each state. For example, in the *Standby* state the actions *Go Off* and *Configure* are possible. Actions are sent from the experimental control system when the human operator sees fit. Alternatively, the actions can be sent from a human operator of the PVSS system alone if the detector is run in stand-alone mode. The actions of the state machine implements the appropriate behaviour for the given target state.

Some actions requires an intermediate state change, like going from the *Standby* state to the *Standby configured* state (see Figure 6.5), since these actions can potentially take a longer time to complete than most other actions. The intermediate states (such as *Downloading*) are there to let the ECS know that something is going on, and the state machine has not dead-locked or crashed. The state machine will automatically leave these intermediate states, and change to the appropriate target state, unless of course, an error occurred, in which case the state machine will go to the *Error* state.

Go Standby: The following operations are performed in order. If any step fails, go into the *Error* state.

1. Check that racks holding the power supplies are on, and continue evaluating this.
2. Check that the cooling plant is supplying the FMD with cooling water, and continue evaluating this.
3. Turn on the low voltage easy crate, and turn on power for the RCU card and daughter SIU and DCSC.

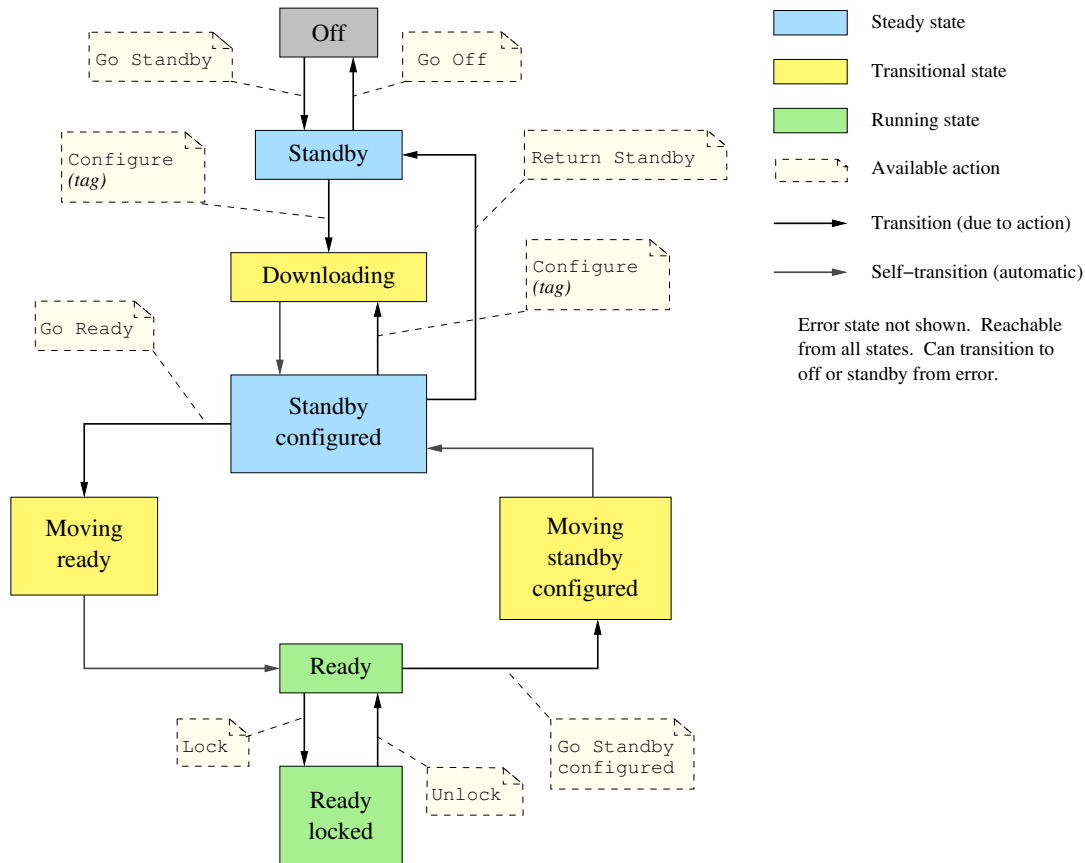


Figure 6.5: The FMD top-level DCS finite state machine.

4. Check that the FeeServer and RCU is up by checking the services published. Continue periodic checking of the RCU state.
5. Enable power output to the FMDD cards.

Go Off: Essentially this does the inverse of the above command.

1. Turn off power for the FMDD cards.
2. Turn off power for the RCU cards and their daughter cards.
3. Turn off the low voltage crate.
4. Ignore the state of racks and cooling.

Configure: This action has one associated parameter — the run type or *tag*. The tag encodes what kind of run we are configuring for e.g., physics, pedestals, pulser, and so on. The steps taken by the action itself does not depend on the value of the tag, but the configuration of the front-end electronics does (see also Section 6.2.4), so the tag value must be passed in here.

1. Send the command `ConfigureFero(tag)` to the ICL.
2. Go to the *Standby configured* state.

Return Standby: Send the command `ConfigureFero(standby)` to the ICL to let the FeeServer turn off the FMDD cards.

Go Ready: This action is executed before starting a run.

1. Enable output of bias voltage.

2. Enable state machine that monitors the FMDD. If an FMDD should certainly be turned off or on, then this sub-state machine will turn off or on respectively the bias voltage to the associated silicon sensors. Since the sub state machine will immediately detect that FMDDs has been turned on, it automatically ramp up the power to the associated sensors.
3. Go to the *Ready* state.

Go Standby configured: When exiting a run, this action is executed. We go back to the state *Standby configured* since the front-end is still configured at this time. One exception to this rule is after pulser runs. The front-end is left in an unsuitable state for further data taking, so the state machine may jump back to the *Standby* state.

1. Disable monitoring of the FMDD for the purpose of turning on or off bias voltage to the sensors.
2. Ramp down any bias voltages.
3. Disable the bias voltage outputs.
4. Go to the *Standby configured* state if the last run was not a pulser run — otherwise execute the action *Return Standby*.

Lock and Unlock: Both of these actions have no associated procedure. The state *Ready Locked* exists merely to provide rudimentary access control. When in the *Ready Locked* state, one *must* execute the action *Unlock* to do anything else. The *Unlock* action can in principle require authorisation.

In any given state, there is a number of services that are monitored, as defined by lower level state machines. The values monitored in each state of the top-level state machine are

Off: Nothing is monitored.

Standby: Rack power, cooling of the detector, and RCU and FeeServer state.

Standby configured: As *Standby* as well as FMDD status (temperatures, voltages, and currents, see Table 5.6).

Moving Ready: As *Standby configured*.

Ready: In addition to the values monitored in *Standby configured* will a sub state machine monitor the front-end cards and turn on/off the associated bias voltage supplies. The bias voltages are monitored, and if one should drop out without the corresponding FMDD card is turned off, the state machine goes into the error state.

Ready Locked: As *Ready*.

Moving Standby configured: As *Standby configured*.

If a monitored value is out of bounds, the lower level state machine may transition into an error state, which the top-level state machine will react to be changing into the *Error* state itself. In certain cases it may be possible to do automatic recovery of the error, but in all other cases the ECS will see the error and the human operator can take appropriate action e.g., take the FMD out of the run and have the expert take control to try to recover from the error.

Reprogramming of the FMDD BCs via remote JTAG Part of the FMD detector control system, though not part of the PVSS controls, are also the possibility to re-program the FMDD BCs via uploading new firmware to the flash memory on the FMDD. This is done from a so-called engineering node in the DCS network. The 10 JTAG connections, corresponding to the 10 BCs in the FMD are connected to a custom switch board near the engineering node. A program on the computer will load the firmware onto each BC in turn, switching through the channel using a USB interface. In this way, bugs fixes, and upgrades can easily be done to the BC firmware on the FMDD.

6.2.4 Configuration of the Front-end

Figure 6.6 illustrates the steps involved in going from the *Off* state to the *Ready* state. On the right hand side of the figure is a blow-up of the steps for the front-end chain.

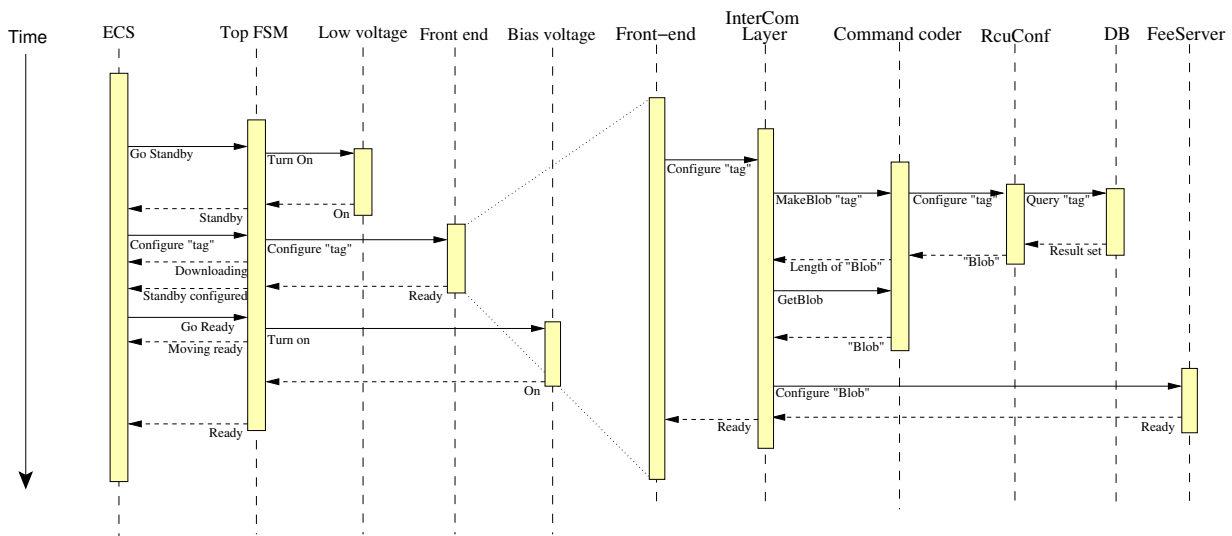


Figure 6.6: Steps executed in going from the *Off* state to the *Ready* state.

When the ICL receives the `ConfigureFero(tag)` command, it will ask the user defined command coder to make a *blob* (a chunk of binary data⁶) which it will then hand over to the FeeServers and let them process it.

The FMD implementation of the command coder is part of the RcuConf package, which in turn uses the Rcu++ to do the actual encoding. When the command coder receives the request to make a blob, it does several queries against a relational database⁷. The schema (or structure) of the front-end electronics configuration database is shown in Figure 6.7.

The command coder executes the following sequence of queries on the database

1. The command coder queries the database for the most recent version of an entry in the *Config* database with a tag matching the passed tag. If none is found, the command coder queries for the default configuration. If there is no default configuration, the command coder returns an error. The fields *x*, *y*, and *z* are for making sub-detector specific configurations, and are reserved for future use.

⁶First off, please notice that the blob described here is not a database blob. Secondly, some people think that blob is an acronym for Binary Large Object, but in fact it is not — Jim Starkey, the inventor of the term describes it as "... the thing that ate Cincinnati, Cleveland, or whatever." For the so-inclined the full, rather amusing story can be found at http://www.cvalde.net/misc/blob_true_history.htm

⁷The exact implementation of database manager is irrelevant, but for the DCS production environment, an Oracle implementation was chosen. For testing purposes a MySQL implementation was used.

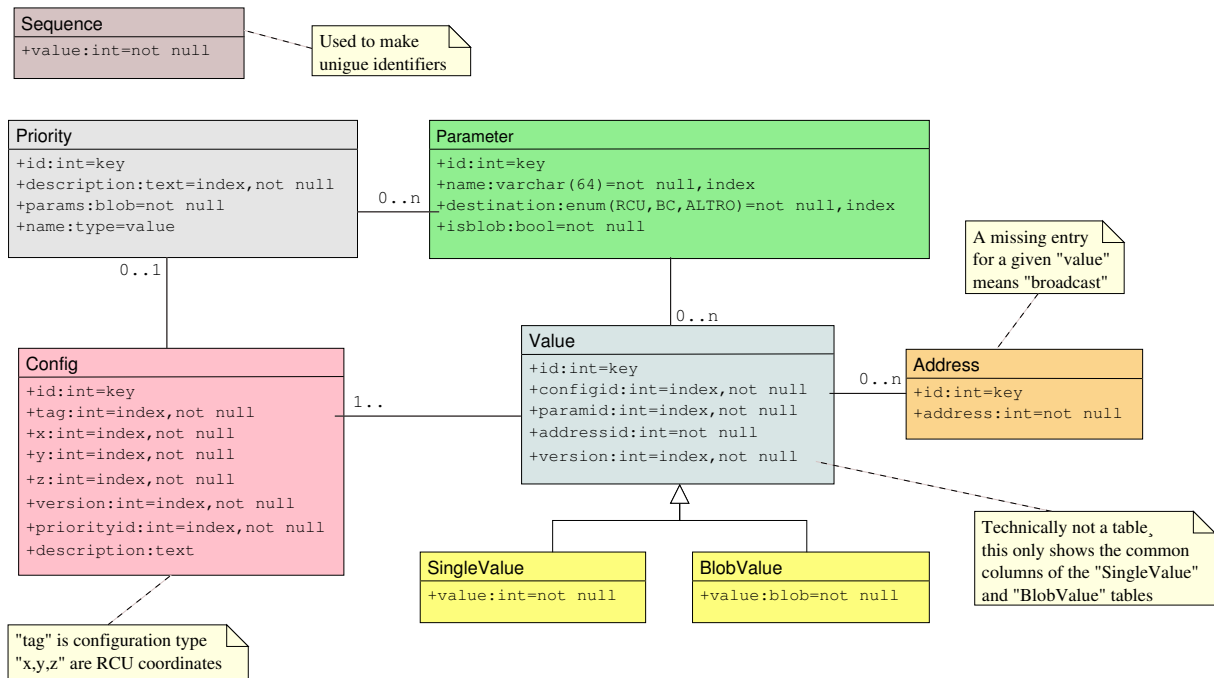


Figure 6.7: Schema of the DCS front-end electronics configuration database.

- Then, the command coder queries the database for an entry in the *Priority* table with a unique identifier (*id*) matching the priority field of the found configuration (*priorityid*). If no matching entry is found, the command coder exists with an error. The entries in the *Priority* table encode the sequence of what registers, memories, and commands should be written or executed when. That is, the list of parameters (*params*) specifies the order of the operations needed to configure the front-end. Many configuration entries can point to the same priority entry, minimising the table size for better performance.
- For each element of the *params* list in the found *Priority* entry, the entry in the *Parameter* table with the same unique identifier is retrieved from the database. The entries of the *Parameter* table gives information about the register, memory, or command that the FeeServer should set or execute. It has a name for identification purposes, a destination for the parameter in case of name-clashes, and a flag that tells the code whether the parameter has single-values (i.e., a register or command) or multi-valued (i.e., a memory).
- For each parameter entry, the command coder checks whether the parameter is single-valued or multi-valued. If the parameter is single-valued, a number of queries is made against the database table *SingleValue*, or if it is multi-valued, against the table *BlobValue*.

The queries has in both cases the same structure. First, the command coder checks if it can find the most recent version of an entry in one of the value tables with a parameter identifier (*paramid*) that matches the parameter in question, a configuration identifier (*configid*) of the found configuration from step 1 above, *and* no address reference (*addressid*=0). If no such entry is found, a new query is made, but this time with the default configuration type (*configid*=0). If an entry is found with either of these queries, the parameter and value data is passed to the Rcu++ library for encoding. The values found this way are considered to be broadcast to all front-end components of the relevant type.

The command coder then executes the same queries again, but this time with out specifying an address identification. The result is in general a list of values, each with their own unique corresponding entry in the *Address* table. The values found, plus their corresponding

address entry, as well as the parameter information is passed to the Rcu++ for encoding. The values found this way are intended to be written to the *specific* front-end components addressed by the corresponding address entry.

Thus the strategy is first to write all values in broadcast and then to override for values for specific cards if needed. Since the FMD front-end is very homogeneous, it is most like that overrides for a specific front-end component are rare. Thus broadcasts should constitute the bulk of the configuration.

Once the command coder has done the above queries and the FeeServer command blob has been encoded, it is handed off to the ICL which will then send it untouched to the FeeServers. The FeeServers then interpret the commands, and changes the internal state representation of the attached RCU to *configured*. Thus the front-end chain finite state machine can wait for all RCUs to be in this state before transitioning to *Standby Configured*.

Various tools have been developed to edit the front-end configuration database, including a graphical user interface allowing painless updates and reconfigurations. It is important to realise, that no entries in the database are ever overwritten or deleted, to ensure that one can reconstruct the circumstances for a given run if so needed. Also note, that all fields in the database, with the a few exceptions, are human read-able. That ensures that it is transparent to the FMD experts exactly what is done during the configuration of the front-end and makes the database easier to edit using standard tools.

Chapter 7

Test beam results

To gain experience and to verify the design of the FMD, some of the sensors and the read-out electronics were brought to University of Århus, Denmark, to be put in a 630 MeV electron test-beam. The beam is provided by the Aarhus *Storage Ring* in Denmark (ASTRID) [94], as a parasitic beam extracted by a septum magnet, as shown in Figure 7.1.

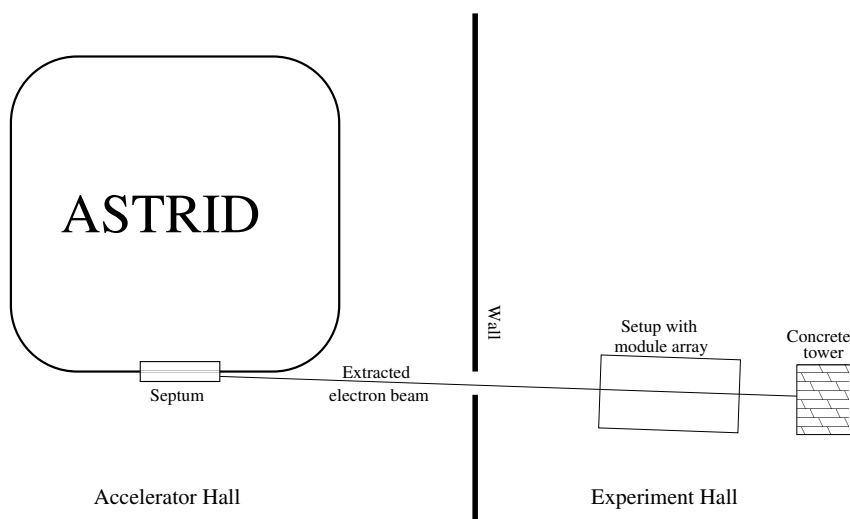


Figure 7.1: Extraction of parasitic beam from ASTRID.

The 630 MeV electron beam allows us to study the detector response to minimum ionising particles. In particular, the signal-to-noise ratio is of interest, as it gives an indication of how well we can reconstruct the over all multiplicity in A-A collisions. The requirement put forth in the technical design report was a signal-to-noise ratio of at least 10:1 for both inner and outer types of modules.

Initially, single sensor-hybrid modules were put in the beam to test the response of the silicon sensors and the VA13s. Later on, an array of sensors were put in the beam to allow studies of detector efficiency.

7.1 Initial Tests of the Modules

Setup A single module, either of the inner or the outer type was placed in the beam. Two small cubic $2 \times 2 \times 2$ cm scintillators were placed on either side of the module, and the coincidence of both these counters was used as the trigger. The read-out system used was the test system provided by IDEAS — the VATAc. Labview software running on a Windows computer was used to read-out the, store and analyse the data. The VATAc system has a single channel 13bit ADC,

making the intrinsic resolution of the ADC $2^3 = 8$ times better than the 10bit ALTRO resolution. Pedestal and pulser runs were taken to calibrate the system.

Collected Data More than 1 000 000 events were collected on a subset of the strips on a single module. The distribution of the calibrated signal for a single strip is shown in Figure 7.2.




Figure 7.2: Spectrum of calibrated single strip signal. Notice the many counts between the pedestal (tall peak on the left at 0) and the minimum ionising signal (at around 3.8).

The abundance of counts between the pedestal (the tall peak around 0 on the left in Figure 7.2) and the minimum ionising signal (around 3.8 in Figure 7.2) is characteristic of signal *sharing* between neighbouring strips. This can be seen by plotting the correlation of the signals of two neighbouring strips, as done in the left panel of Figure 7.4.




Figure 7.3: Hit sharing from particle traversing more than one strip

Sharing occurs when a particle traverses more than one strip, by impinging on the sensor at an angle $\neq 90^\circ$, as sketched in Figure 7.3. The energy loss suffered by the particle then goes to ionisation into two strips. It is important to note, that energy loss is a stochastic process, and may or may not happen in either strip. On average though, the amount of energy loosed in either strip is proportional to the path length through the strip.

The insert on the left of Figure 7.4 clearly shows that the sum correlated signal is Landau distributed, corresponding to the expected spectral shape of a single particle (see also Section 5.1.2).




Figure 7.4: Plotted on the left panel is the correlation of the the signals of two neighbouring strips. The insert shows the distribution of the sum of the signals from the two strips when they are both above the cuts indicated by the horizontal and vertical lines in the main panel. To the right is the full signal spectrum shown, as well as the spectrum after disregarding the shared signals.

On the right is again the full signal spectrum plot, as well as the spectrum after disregarding all shared signals. Clearly the summed distribution of the correlated signals, in the insert on the left, and the signal after subtraction of shared signals are Landau distributed with more or less the same with and most probable value.

Results After the shared signals are subtracted from the signal distribution, the Gaussian distribution is fitted to the pedestal and a Landau distribution to the signal data. The signal-to-noise S/N is then calculated as

$$S/N = \frac{\Delta_p}{\sigma}$$

where Δ_p is the most probable value of the fitted Landau distribution (see Section 5.1.2) and σ is the standard deviation of Gaussian distribution fitted to the pedestal. A signal-to-noise between 57:1 and 63:1 was found [95] for the inner modules, and around half that, 35:1 for the outer modules. For the inners, this is in excellent correspondence with the estimates of 58:1 to 64:1 done calculated in Table 5.1.

7.2 Full Test of the Modules and Electronics, and Efficiency

After the FMD digitiser cards had been developed and manufactured, a full setup of the front-end electronics and sensor modules were brought back to Århus for a full electronics system test. The purpose of this round of tests was to gain experience with the system, test the final configuration for characteristics such as signal-to-noise, and to determine the detection efficiency of the silicon sensors.

7.2.1 Setup

The test-beam setup consisted of

- 5 silicon modules.

- Two scintillator counters measuring roughly $12 \times 12 \times 1$ cm for triggering.
- A NIM crate for making triggers sequences.
- One or two FMDD cards (an inner or an inner and an outer).
- An RCU card with daughter DCSC and SIUcard.
- A standard computer with a RORC installed and running the ALICE DAQ program DATE. This machine was running the services of an LDC in stand-alone mode.

All of this was put in a custom made rack to make it easier to align the sensors relative to the beam¹.

Slow control was done using custom-made programs based on Rcu++, including a graphical user interface (RcuGui) for low-level access to the front-end, and an on-line monitor for monitoring data as it arrived at the DAQ machine.

The silicon modules were placed in a specially made box to keep light away from the light-sensitive VA1₃ chips, as shown in Figure 7.5. The box was put on elevation tables to be able to relatively easy move the silicon modules up and down in the beam. The two scintillator counters were placed in front and behind the box relative to the beam direction. Note, that the scintillator counters are large enough to fully cover the acceptance of the silicon modules.

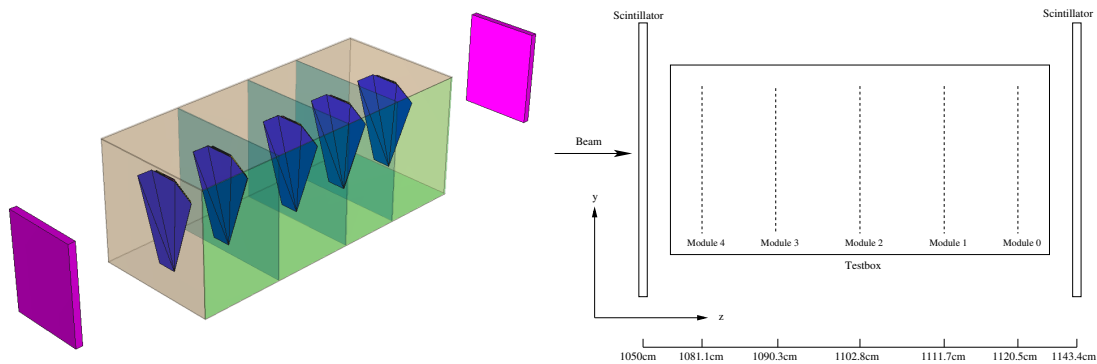


Figure 7.5: Arrangement of the silicon modules in specially manufactured box. The beam enters the box on the left, and exists on the right. The right-hand side of the figure shows a plane view of the arrangement of the detectors, and the distances of each silicon module to the origin of the beam (the septum magnet).

The beam profile was measured [96] by counting coincidences of a small cubic scintillator with the front planar scintillator on a grid of 2 cm lattice spacing. The beam was found to be rather flatly distributed with smooth fall of at the edges. The width of the beam was found to be a little smaller than the scintillating counters. The raw coincidence rate of the two scintillating counters was around 3 kHz at the start of a store in ASTRID, dropping to around 1 kHz near the end of a 24 hour store. In other words, ASTRID was able to provide us with an excellent beam throughout the running periods.

7.2.2 Overall performance

The trigger was setup to deliver a 10 triggers at rate of 100 Hz, with a derived $\overline{11}$ trigger arriving $5.3 \mu\text{s}$ after. At this rate, the system could be run for many hours, in some cases over-night, until the hard disk of the computer was filled up with data. The system maintained good performance throughout the running period and there were no long term effects observed on the recorded data. The system kept a stable operating temperature despite the lack of a cooling system.

¹See also photo in Figure F.3.

7.2.3 Signal-to-Noise

Since the beam profile was not broad enough to envelope the whole sensors, three separate runs were done, one with the sensors centred in the beam, and two with the sensors displaced roughly half a sensor with up and down. In this way, we could collect enough statistics to get a reasonable data set to do signal-to-noise analysis of the whole sensors. Figure 7.6 shows the extracted signal-to-noise ratio for an inner and an outer strip [97].




Figure 7.6: Measured signal-to-noise ratio of inner (left) and outer (right) silicon sensors using 630 MeV e^- from ASTRID

Note, that there is substantially less sharing in this data than in the initial test data (see Section 7.1). This is attributed to better alignment of the sensors to the beam than in the initial data. The signal-to-noise ratio observed with the full read-out system, as it will be in ALICE, was approximately 40:1 for the inner type sensors and 22:1 for the outer type. This is less than what was predicted by the estimates of Table 5.1, and less than the observed signal-to-noise seen with the VATAC system. The exact reason for the worse signal-to-noise ratio is not understood. One source of this additional noise, is the resolution of the ALTR0, which makes the pedestal relatively wider in number of ADC counts than the signal peak, but this does not account for all of the decrease in the signal-to-noise ratio. Figure 7.7 and 7.8 shows a summary of the measurements done for an inner and outer module.

Clearly the performance of the sensors are rather constant, with the exception of the two top-most VA1₃s on the inner cards (strips 384 to 640 on Figure 7.7). For these strips, the lines that connect the silicon bonds to the VA1₃ pitch adapter run over the hybrid card PCB for some distance, and these lines are believed to be the source of the additional noise seen for those strips.

7.2.4 Efficiency

The setup outlined in Section 7.2.1 was also used for estimating the detection efficiency of the silicon sensors. Carsten Sogaard did extensive studies of various ways of determining the efficiency in his master thesis [96]. The method decided upon meet certain criteria, quoting [96]

- [The method] must be insensitive to the acceptance of the detector array.




Figure 7.7: Count statistics, signal, noise, and signal-to-noise ratio of the 1024 strips in a single inner module. More information is available in [96].




Figure 7.8: Count statistics, signal, noise, and signal-to-noise ratio of the 512 strips in a single outer module. See also [96] for more information.

- Secondary particles must be omitted if they are [not] countable.
- The method must be robust i.e., it has to be insensitive to small changes in the parameters of the method.

To investigate the efficiency of a target module, for example module 3 (see Figure 7.5) one module on either side is chosen as the first reference module e.g., module 0. Another module on the other side of the investigated module e.g., module 4 is chosen as the second reference module.

For each event, all possible *reasonable* combinations of straight lines from the first to the second reference module is made. Reasonable in this case means, that the line *must* be relatively parallel to the beam direction. If there are more than one possible line that starts from the same strip in the first reference module but join to different hits in the second reference module, then the line with the lowest sum of square residuals is chosen and the other discarded. This ensures that lines from secondary particles (created in scattering or the like) are not counted.

The strip where the line intersects the target module is calculated. If there is a hit in the target module, within some cut, in the expected strip, then the hit is counted as being detected. If there are multiple lines that could join up with a hit in the target module, the line that has the best χ^2 , when counting the hit in target module, is chosen. The efficiency is then calculated as

$$\epsilon = \frac{N_d}{N_t},$$

where N_t is the total number of possible hits from particles, N_d is the number of detected hits. The uncertainty of the efficiency $\delta\epsilon$ is calculated as the 68.3% confidence interval around the value ϵ [98].

The setup was simulated, including the beam profile, using GEANT to better understand the results of the analysis [96]. The geometry and composition of the silicon modules in the simulation is exactly the same as the one used in general ALICE simulations (see Section 8.1). The data from the simulation was analysed in exactly the same way as the data from the test beam experiment. Since the simulation is in principle 100% efficient, except for particle decays and the fluctuations built in to the physics process of GEANT, it provides a good basis for comparing the analysis of the test-beam data.

Figure 7.9 shows the efficiency deduced from the test beam data (top panel) as well efficiency deduced from the simulated data (middle panel). The lower panel of the figure shows the test beam efficiency to the simulated efficiency. Clearly there is a good correspondence between the two, indicating that systematic errors are small. Notice the dip in both the top and middle panel near strip 350. The reason for the lower efficiency in this region is scatterings in the copper feet that hold the hybrid to its support plate. The amount of data available near the edges is severely limited and the use of the method in these regions is dubious.

The conclusion of [96], is that FMD detection efficiency is better than 99.5%. Coupled with the fact that less than 1% dead strips were found in the post-assembly testing, it is clear that the FMD will be a very good detector. Furthermore, the nice correspondence between the test beam data and the simulated data gives faith in the correctness of the simulation.




Figure 7.9: Efficiency of strips in a single inner module. The top panel shows the efficiency as determined from data, while the middle panel is the efficiency deduced from simulations. The bottom panel shows the ratio of the test beam efficiency to the simulated efficiency. Notice the dip in both the test beam and simulated data efficiency near 350. It is caused by scatterings in the copper feet that hold the hybrid to its support plate. The amount of data near the edges limits the applicability of the method in these regions. Adapted from [96].




Figure 7.10: Zoom of Figure 7.9 around the strips that has the most data.

Chapter 8

An Analysis Chain for the FMD Data

The previous chapters described how the signal induced by a charged particle traversing the silicon sensors is amplified, digitised, formatted into data blocks, sent to the DAQ and finally stored on permanent storage. Those steps constitute the on-line system, but are of course not the end of the story. The stored data needs to be processed and the events reconstructed off-line to form the basis of the physics analysis, which will eventually be used for making physics results published in scientific journals.

The ALICE off-line architecture is based on the ROOT¹ [99] toolkit, which provides hardware abstraction, statistical data analysis tools, network transparency, database functionality, to mention a few of the things ROOT can do. The choice of ROOT as the root of all off-line code in ALICE means that the individual developers are freed from thinking about many issues already handled by ROOT, and allows ALICE to run its off-line software in heterogeneous environments.

ALIROOT is the name of the ALICE off-line project. The flow of data through ALIROOT is shown in Figure 8.1. The left-hand side corresponds to simulated data (see Section 8.2 and 8.3), while the right hand-side corresponds to handling of data, whether it is from the experiment or from simulated data (see Section 8.4 and 8.5).

At the top-left of Figure 8.1 are event generators. They represent modules of the particle production in an event, whether it is pp or Pb-Pb collisions. The representations of the particles from the event generator are fed to a detector simulation, which propagates the particles through the ALICE geometry and simulates their interactions in matter, including the decay of short-lived particles. The output of the detector simulation is a list of hits in elements of the geometry defined to be active e.g., the FMD silicon modules or the TPC gas volume. These hits represent the *true*² knowledge of the how the particles interacted with the material, for example the energy deposition, and the space-time coordinates of the interaction. To make the simulated hits resemble data as it comes from the detectors, the hits are digitised to produce digits e.g., ADC counts in the detector. The final step in the simulation branch is to pack these digits into the on-line event format.

At the bottom of the parabola, is the input of raw data, whether it comes from the experiment or from the simulation branch. The first step is to decode the raw data into a format convenient for the reconstruction, namely the digit format. The reconstruction then takes over. In the reconstruction pass the digits are analysed and the space-time points with information about energy loss, interactions, and so on is reconstructed. Then, reconstruction of the tracks left in the detector by the particles is done, and we end up with a description of the event in terms of particle trajectories, and global event properties such as charged particle multiplicity, centrality, and so on.

The outcome of the reconstruction, the so-called *Event Summary Data* (ESD) is what is used for the physics analysis. Since the amount of data produced by ALICE is phenomenal (1 MB

¹The author is a frequent contributor to ROOT, and had the pleasure of working with the ROOT core team as a CERN summer student in the summer of 2000.

²*True* in this context means the truth according to the detector simulation.

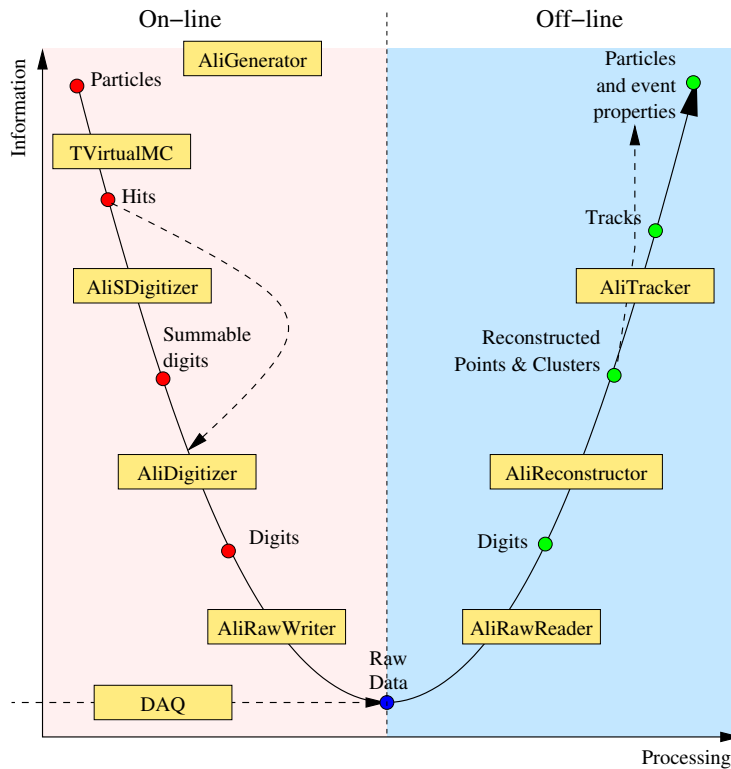


Figure 8.1: Flow of data through ALIROOT

per event in pp and 12.5 MB per event in A–A, corresponding to approximately 1.25 GB/s and 2 PB/year for 1 month of Pb–Pb and 5 months of pp collisions), only 3 reconstruction passes are planned: 1 shortly after the data has been recorded, and 2 some months after. It is therefore crucial that enough information is available in the ESD to extract the Physics results. The data is collected at the so-called Tier-0 system at CERN where it will be put in permanent storage. The later reconstruction passes must however run in a distributed fashion over the global Grid [100], and the resulting ESDs are stored on the various national and regional Tier-1 centres. From there, anyone with a valid ALICE credential can pick-up and analyse the data in the ESDs. The Physics Working Groups will in addition have pre-scheduled passes over the ESDs to produce the so-called *Analysis-Object-Data* (AOD) for specific physics analysis.

This chapter will outline the off-line code developed for the FMD, starting with the geometric description of the FMD and working our way along the parabola of Figure 8.1.

8.1 Geometry

Within the ALIROOT framework, the geometry FMD detector has been described in detail. The basic building block of the geometric description, as it is in the physical detector, are the silicon modules.

A sensor of the appropriate geometry is segmented into two sectors azimuthally and, 256 (for the outers) or 512 (for the inners) strips radially. Behind them are placed thin layers of PCB, copper, and gold to simulate the hybrid stack-up³. In effect, this description of the hybrid card smears the material of the card over the full area, which is reasonable, that is, no attempt has been made to describe the VA1₃ bulk volumes or the ceramic pitch-adapters. However, as shown in Figure 7.9, such a level of detail is not needed to make an appropriate description of the material in the FMD. What is included in the description, are the copper feet that hold a hybrid

³The stack of layers in a printed circuit board, and mounted components.

card to its support. Again, from Figure 7.9 it is clear that the description is quite adequate. Representations of sensors of the inner and outer types are then placed in logical volumes that represents a half ring.

For each sub-detector, instances of the appropriate half-rings are then grouped, along with representations of their supporting honeycomb plates and carbon support structures into half sub-detectors. This corresponds to how the detectors were assembled and installed in the ALICE experiment hall, and allows for intuitive manipulation of the FMD volumes for alignment purposes. The half rings are then placed at the appropriate locations inside the greater ALICE geometric description.

The elements, the sensor modules and half rings, thus placed, are put at their nominal position. Since the detectors may not sit exactly where they were intended to be, the ALIROOT framework allows one to create small translations and rotations and apply these to the geometric description. The FMD2 and 3 half-rings are fixed with a precision of $100\ \mu\text{m}$ or better, relative to the ITS, and similar for the FMD1 half-rings to V0-A. By construction, the location of the FMD elements are known to a precision of $100\ \mu\text{m}$ or better, relative to the half-rings. Since the η overlap between tracking detectors (like the ITS) and the FMD is very small, there is currently no ambition to refine this alignment from data.

Note, however, that FMD2 and 3 are physically attached to the ITS support structure, meaning that if the ITS is physically moved, so will FMD2 and 3. How to propagate the alignment by the ITS to FMD2 and 3 is currently under investigation of the ITS and FMD group. FMD1 is physically mounted to V0-A, and the precision of the installation of the V0-A and FMD1 compound is currently not known. An option for doing optical survey of this compound has been reserved in case it is found to be needed after the installation of FMD1 in January/February 2008.

The geometric description of the FMD sub-detectors is based on a parametric model — that is, the shapes, size, and positions are based on a hand-full of parameters, and any change of these parameters will be propagated to the description automatically. This allowed the FMD group to easily check various geometry options based on a few design parameters. The geometric description is used throughout the rest of the code, with a few exceptions. For example, the simulation uses the geometry to do the proper propagation of tracks and interactions in matter, while the reconstruction code uses the geometry to map from detector coordinates (detector,ring,sector,strip) to physical coordinates (η, φ) .

8.2 Simulations of events

The ALIROOT simulation framework is constructed around the ROOT abstract simulation interface `TVirtualMC`. Concrete implementations of this interface are GEANT 3, GEANT 4, and Fluka. From the point of view of the developer, the 3 available implementations are the same, since the details are hidden. That means, that one can easily switch simulation implementation without changing even a single line of code. The three available simulation packages differ in how they treat physical processes. For example, GEANT 3 treats most processes effectively, while FLUKA does more microscopic calculations.

The flow of the simulation is driven by the input data, no matter what implementation is used. That is, for a given event, the simulation takes one particle off the top of the stack of particles and propagates it through the detector. It then takes the next particle and propagate that, and so on, until the stack is empty. If a particle decays, or somehow produces new particles as a result of an interaction, then the new particles are pushed onto the stack, and will be removed when its their turn to be propagated.

Whenever something happens in the simulation e.g., a particles enters or exits a volume declared active, a particle losses energy in a volume, and so on, some user routines are called. For example, whenever a particle loses energy in the FMD sensors, some FMD code is called, and it

is up to that code to define a hit that will be output to disk. The FMD hit output contains the space–time coordinates (x, y, z, t) as well as the detector coordinates (detector,ring,sector,strip), the total energy lost by the particle, and the length of the particles trajectory through the strip.

Figure 8.2 shows the energy loss per unit length as recorded from 100 simulated HIJING events. The simulation implementation used here, and elsewhere, was is GEANT 3. In the simulation, the production of δ –electrons was turned on, allowing these to carry away some of the energy loss without detection. If the production of δ 's was not turned on, the correspondence with the known results (figure 27.6 in [81]) would be far worse (see also Figure E.1). Also notice the rather long tail towards higher Δ/x around $\beta\gamma = 20$. This will affect how well the FMD can reconstruct multiple hits in a single strip, since the long tail will tend to obscure multiple hit signals.




Figure 8.2: Energy loss in the FMD silicon sensor from a GEANT 3 simulation (points). Imposed are various calculations as explained in the legend. The *Stopping* curve is the full mean energy loss curve (see also Figure 5.9). Two variants of this, without radiative energy loss or δ –electron production is also shown. Δ_p/x refers a parameterisation of the observed most probable energy loss 320 μm in Silicon. The arrow indicates the $\beta\gamma$ for which $d(\Delta/x)/d\beta\gamma \sim 0$. *RPP* in the legend refers to [81]. Note, that in the simulation production of δ –electrons was turned on.

The output of the simulation — the hits, are useful for checking that the simulation works as expected. Since everything about the event, and how the particles propagated through the experiment, is known, the simulation can be used to estimate various things, such as primary to secondary corrections and the acceptance of the detector.

8.3 Digitisation of Simulated Data

To make the simulated data look more like real data, another step is needed. This is known as digitisation. The true (GEANT produced) hits are processed and converted into a simulated detector response. The information in the hits is sometimes also known as the GEANT *truth*. For the FMD this consists of taking the integrated energy loss Δ_i of all particles that hit a silicon strip in an event, and convert that into ADC counts, or so-called *digits*. To make the digits more

realistic, random noise is added to the ADC values. For a given strip, the digit ADC count c_i is given by

$$c_i = p_i + x + g_i \Delta_i C \quad (8.1)$$

where p_i is the pedestal of the strip, x is random Gaussian deviate with a standard deviation equal to n_i , the noise of the strip, g_i is the gain of the strip, and C is the fixed conversion factor from gain calibrated signal to ADC counts. The pedestal p_i and noise n_i is determined in pedestal runs, while the gain factor g_i is determined using special pulser gain runs as outlined in Section 5.3.2. The conversion factor C depends on the electronic response of the VA1₃ pre-amplifier, the receivers on the BC, and the ALTR0 digitisation. It is approximately equal to

$$C = \frac{1}{\Delta_p} \frac{\text{max ADC}}{\text{range}_{mip}},$$

where Δ_p is the most probably energy loss, $\text{max ADC} = 1023$ is the maximum value of the ALTR0 ADC, and $\text{range}_{mip} = 20$ is the dynamic range of the VA1₃ chip. However, a better value may be obtained by fitting a Landau distribution the experimental measured gain-corrected ADC spectra, and use the most-probable value scale factor. Thus, the parameters used to make simulated digits, are exactly the same as the parameters used in the reconstruction from raw data. Therefore, the parameters used for the reconstruction is used in the digitisation. This provides a consistency check, in that the reconstructed data should reproduce the true hits. Figure 8.3, shows the correlation of the digitised value (the ADC counts) and the true energy loss.




Figure 8.3: Conversion of energy deposition into ADC counts

Notice, that the ADC saturates at 1023 ($= 2^{10} - 1$) since the ALTR0 is a ten-bit ADC. Therefore, signals with very large energy deposition will be registered as 1023, and this limits the available range for reconstruction.

In case of an oversampling ratio (see Section 5.3.1) different from 1, more than one ADC count must be made for each strip, since in this configuration, the ALTR0 will sample the input several times. When the analogue input to the ALTR0 ADC shifts from one level (corresponding to the signal of strip i) to a new level (signal from strip $i + 1$), it takes a certain amount of time for the signal to reach its full value. The time it takes is the rise time of the electronics. The

digitisation code simulates the electronic rise time of the signal by the function

$$f(c_{i,s}, t, c_{i-1}) = \begin{cases} (c_{i,s} - c_{i-1}) \left(1 - e^{-bt}\right) + c_{i,s} & c_{i,s} > c_{i-1} \\ (c_{i-1} - c_{i,s}) e^{-bt} + c_{i,s} & \text{otherwise} \end{cases} \\ = c_{i,s} + (c_{i-1} - c_{i,s}) e^{-bt} , \quad (8.2)$$

where $t = s/\nu_{samp}$ — the sample number divided by the oversampling rate, and b is a constant. This is illustrated in Figure 8.4.




Figure 8.4: Illustration of ADC value generation in case of over sampling rate 1, 2, 3, and 4. The squares shows the fully converted ADC counts (corresponding to no oversampling), while up and down triangles, and circles corresponds to an oversampling rate of 2, 3, and 4. The rise time of the electronics is simulated according to (8.2).

If one is interested in mixing signal (e.g., rare particle production) and background events (events with no rare signal), one can use the so-called *sum-able* digits (or *s*-digits). They are exactly like digits, except that the noise term x is not added to the output. The strategy is to take *s*-digits from a background event, add *s*-digits from a signal event, and then apply smearing (or noise) to the combined signal-background event.

After the true hits have been digitised, the resulting digits are processed again to put them into the format as written by the DAQ (see Section 6.1.1). For the FMD, this consists of encoding the ADC values into the `ALTR0` channel event format, as outlined in Table 5.23. The rest of the event structure is encoded by the core off-line code.

The code that encodes the `ALTR0` channel data, is shared among all detectors that use the `ALTR0` format. However, each detector must provide a custom map from detector coordinates i.e., for the FMD(detector,ring,sector,strip) to hardware addresses i.e., (ddl,FMDD,ALTR0,channel), where `ddl` is the unique number assigned to the FMD RORCs. Note, that the ADC sample number is not part of the mapping — it is taken care of by the encoding by simply giving the ADC values in the proper order. For the FMD a static map is used i.e., it is always the same and the mapping is calculated. However, an option has been left open to have a dynamic mapping in case of hardware reconfigurations.

8.4 Reading Raw Data

Raw data, as produced either by the experiment or via simulations, as outlined above, is read in, and converted to digit objects. Core code takes care of the common parts of the event structure,

while custom code decodes the ALTR0 channel event format into digit objects. Essentially, the FMD code gets a block of bytes that it should interpret as a list of ALTR0 channel events as described in Table 5.23.

The code to decode the ALTR0 channel event is common among all detectors that use the ALTR0 chip. The particular detector only needs to provide a map from the hardware addresses (ddl,FMDD,ALTR0,channel) to detector coordinates e.g., for the FMD (detector,ring,sector,strip). Note, that the strip number returned is the first strip that the ALTR0 channel is reading data from. The rest of the strip index values are calculated from the oversampling rate.

8.5 Reconstructing Bare Multiplicity

The task of the reconstruction is to estimate the number of particles that hit a silicon strip, whether the particle came from the primary interaction or through secondary interactions in surrounding material or decays. The reconstruction starts from digits — either directly from simulation or from raw data, irrespective of whether the raw data was produced by the experiment or from simulations, and produce a near multiplicity for each strip. The process is illustrated in Figure 8.5

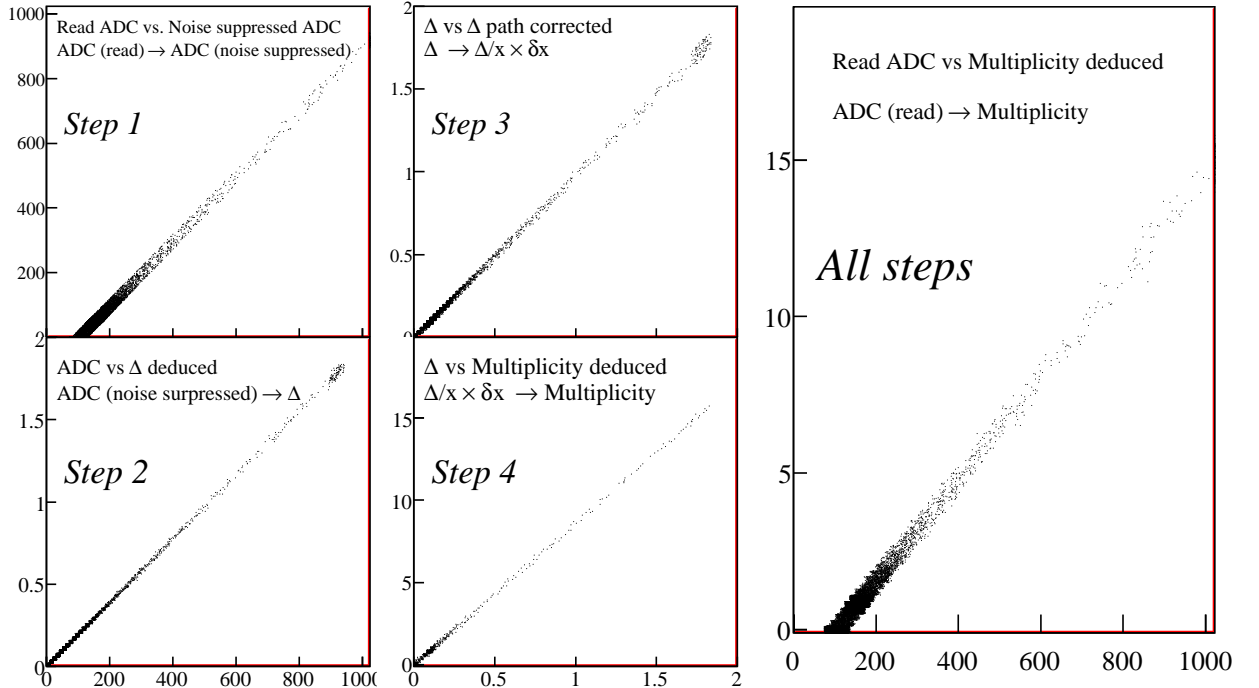


Figure 8.5: The process of reconstruction the bare multiplicity in a strip from the raw ADC counts. The top-left panel (Step 1) is the conversion from the raw read ADC counts to noise and pedestal suppressed ADC values, as given by (8.3) with $f = 3$. Continuing to the bottom-left panel (Step 2), the noise and pedestal suppressed ADC values are converted to energy loss Δ (see (8.4)), and in the top-middle panel (Step 3), the energy loss is corrected for the path-length through the silicon (see (8.5)). In the bottom-middle panel, the corrected energy loss is divided by the most probable energy loss Δ_p (see (8.6)) to give the bare multiplicity. The right most panel shows the full path from read ADC values to reconstructed bare multiplicity.

For each digit, the reconstruction calculates first the pedestal subtracted ADC counts, and if this is smaller than the noise of the strip by some factor f , then the strip is considered empty

and the digit is not processed any further (*Step 1* in Figure 8.5). That is

$$c'_i = \begin{cases} c_i - p_i & \text{for } c_i > p_i + fn_i \\ 0 & \text{otherwise} \end{cases} . \quad (8.3)$$

Appropriate values of f is in the range from 2 to 4, where 2 corresponds to suppressing 95.5% of the noise, and 3 corresponds to suppressing 99.4%, since the noise is approximately Gaussian distributed. In case of on-line pedestal subtraction on the front-end (as outlined in Section 5.3.1), then $p_i = 0$ in the above expression.

Since the input channels of the VA1₃ have different gain factors, and are therefore not directly comparable, the next step is to apply the gain factor g_i extracted from pulser gains. This will give us derived energy deposited Δ'_i when scaled by the inverse common scale C (*Step 2* of Figure 8.5)

$$\Delta'_i = c'_i \frac{1}{g_i} \frac{1}{C} , \quad (8.4)$$

where g_i is the gain for strip i .

Since particles more often than not traverse the silicon sensor at an angle $\neq 90^\circ$, as illustrated in Figure 8.6, and thus see more material, the Δ'_i calculated for each strip in (8.4) are not directly comparable. The signal Δ'_i needs to be corrected for the path length through the detector by a factor of $\delta x_i = \cos \vartheta_i$ where ϑ_i is given by

$$\vartheta_i = \tan^{-1} \frac{r_i}{z_i - v_z}$$

where r_i is the radial distance from the beam to the strip, z_i is the z -coordinate of the strip, and v_z is the location of the primary interaction along the beam axis. This correction essentially normalises the signal to the nominal sensor thickness by a projection onto the normal of the sensor. The path length corrected signal Δ_i (*Step 3* in Figure 8.5) is given by

$$\begin{aligned} \Delta_i &= \Delta'_i \delta x_i \\ &= \Delta'_i \cos \left[\tan^{-1} \left(\frac{r_i}{z_i - v_z} \right) \right] . \end{aligned} \quad (8.5)$$




Figure 8.6: Particles traverse the silicon sensors at some angle, meaning that it will see more or less material depending on strip. Hence the signal Δ' need to be path-length corrected by a factor of $\cos \vartheta$.

The parameters r_i and z_i are derived from the geometric description, after proper alignment, in the off-line code, and are known to a precision of roughly a hundred μm . v_z is determined by other detectors, such as the T0, or better, the ITS. The ITS quotes an uncertainty of [101]

$$\delta v_z = \frac{260 \mu\text{m}}{\frac{dM_{\text{ch}}}{d\eta}} + 7 \mu\text{m}$$

where the “additive constant represent the influence of the ITS residual misalignment.” Thus the uncertainty varies from roughly $90\ \mu\text{m}$ in pp to $10\ \mu\text{m}$ for Pb–Pb.

To get the final bare multiplicity per strip, the derived energy deposited Δ_i is divided by the most probable energy loss for a minimum ionising particle Δ_p (see Section 5.1.2). Thus we arrive at *Step 4* of Figure 8.5 with the expression

$$\begin{aligned}
 M_{\text{ch},i} &= \frac{\Delta_i}{\Delta_p} & (8.6) \\
 &= \frac{1}{\Delta_{mp}} \cos \left[\tan^{-1} \left(\frac{r_i}{z_i - v_z} \right) \right] \Delta'_i \\
 &= \frac{1}{\Delta_p} \frac{1}{g_i} \frac{1}{C} \cos \left[\tan^{-1} \left(\frac{r_i}{z_i - v_z} \right) \right] c'_i \\
 &= \frac{1}{\Delta_p} \frac{1}{g_i} \frac{1}{C} \cos \left[\tan^{-1} \left(\frac{r_i}{z_i - v_z} \right) \right] \begin{cases} c_i - p_i & \text{for } c_i > p_i + fn_i \\ 0 & \text{otherwise} \end{cases}, & (8.7)
 \end{aligned}$$

for the full conversion from ADC count c_i to bare charged particle multiplicity $M_{\text{ch},i}$ of a single strip. The total conversion is shown in the last panel of Figure 8.5.

The uncertainty in this number is given by

$$\begin{aligned}
 (\delta M_{\text{ch},i})^2 &= \left(\frac{\partial M_{\text{ch},i}}{\partial g_i} \right)^2 \delta g_i^2 + \left(\frac{\partial M_{\text{ch},i}}{\partial C} \right)^2 \delta C^2 + \left(\frac{\partial M_{\text{ch},i}}{\partial r_i} \right)^2 \delta r_i^2 + \left(\frac{\partial M_{\text{ch},i}}{\partial z_i} \right)^2 \delta z_i^2 \\
 &\quad + \left(\frac{\partial M_{\text{ch},i}}{\partial v_z} \right)^2 \delta v_z^2 + \left(\frac{\partial M_{\text{ch},i}}{\partial c_i} \right)^2 \delta c_i^2 + \left(\frac{\partial M_{\text{ch},i}}{\partial p_i} \right)^2 \delta p_i^2 \\
 &= \left[\frac{\delta c_i^2 + \delta p_i^2}{(c_i - p_i)^2} + \frac{\delta g_i^2}{g_i^2} + \frac{\delta C^2}{C^2} \right] M_{\text{ch},i}^2 \\
 &\quad + \frac{\delta v_z^2 r_i^2 + \delta z_i^2 r_i^2 + \delta r_i^2 (v_z - z_i)^2}{r_i^2 + (z_i - v_z)^2} \frac{1}{\Delta_p^2} \frac{1}{g_i^2} \frac{1}{C^2} \sin^2 \left[\tan^{-1} \left(\frac{r_i}{z_i - v_z} \right) \right]. & (8.8)
 \end{aligned}$$

The uncertainties δc_i and δp_i are typically 0.5 ADC counts, so that the first term of the first parenthesis is negligible. Similarly, the relative error on the gain g_i is rather small — around a 1% or so, and can be largely ignored. δr_i , δz_i , and δv_z are all of the order of a hundred μm (or $10^{-2}\ \text{cm}$), while r_i , z_i , and v_z are of the order of 10 cm. Thus the first factor in the last term is of order 10^{-4} . The second factor is likewise small, since the angle ϑ^2 is typically small ($< 20^\circ$). It is therefore obvious that the main source of the uncertainty in $M_{\text{ch},i}$ is the error on fit parameter C , and

$$\delta M_{\text{ch},i} \approx \frac{\delta C}{C} M_{\text{ch},i} \quad .$$

The bare multiplicity per strip $M_{\text{ch},i}$ is stored in *Event Summary Data* (ESD) structures, along with the unique pseudo-rapidity η_j of each strip. That is, for a given radial distance r_j and distance along the beam line z_j corresponding to a strip, only one η_j value is stored, since that value is the same for all strips with $r_i = r_j$ and $z_i = z_j$. The pseudo-rapidity is calculated as

$$\begin{aligned}
 \eta_j &= -\log \tan \left(\frac{\vartheta_j}{2} \right) \\
 &= -\log \tan \left(\frac{\tan^{-1} \frac{r_i}{z_i - v_z}}{2} \right) & (8.9)
 \end{aligned}$$

The output data also contains the v_z used in the calculations of η_j and the path-length correction. This allows one to reapply a possibly better known z coordinate of the primary interaction point to the FMD data. The data stored in the ESD structures is thus independent of the precise

geometry and calibration parameters, while still retaining enough information to use the FMD data in various ways.

It is important to understand, that the $M_{ch,i}$ for each strip, after the reconstruction, is *not* the final multiplicity. The output data still needs to be corrected for sharing and secondary particle corrections.

8.6 Hit Sharing Corrections and Particle Cuts

As mentioned earlier, a particle that traverses the silicon sensor at some incident angle ϑ' can deposit energy in more than one strip, leading to the effect of sharing. That is, one strip does not see the full energy deposited by the particle, but only a fraction thereof, while a neighbouring strip will see the remaining energy signal. The fraction of energy seen by a strip depends on the radial distance from the point-of-entry to the strip boundary Δr and the angle of the particle ϑ' relative to the beam, as shown in Figure 7.3. However, since the distance Δr is not known, there is no way to determine the amount of sharing *a priori*. Instead the sharing must be corrected for directly on the data.

Clearly, a tell-tale sign of a strip having a shared signal, is that the signal is below the expected Δ_p for a minimum ionising particle, and that one of the neighbouring strips has an equally small signal. The strategy is therefore to loop over all strips in a given event, looking for pairs of strips that have signals less than Δ_p for a MIP, and then combine these into a single signal. In principle, a signal could be shared over many more than two strips. However, for that to happen, the incident angle ϑ' must be larger than 45° , due to the geometry of the strips (see also Figure 7.3). Due to the position of geometry and the distance of the FMD to the primary interaction point, $|z_i - v_z|$ allowing for a realistic spread of ± 15 cm of v_z , there will never be primary particles originating from the interaction with such a large angle. Furthermore, since the point of origin and the direction of secondaries is not known, it is unreasonable to impose further corrections from sharing over 3 or more strips.




Figure 8.7: The distribution of strip multiplicity before and after the sharing correction, as well as the merged signals and the signals removed.

Figure 8.7 shows the process sharing correction. If a strips or its neighbours signal is less than

the sharing cut Ξ_1 (indicated on Figure 8.7 by the dashed line), then the strip and the neighbour are merged into a single strip, by adding their signals. The strips thus selected are shown in Figure 8.7 as the merged signal. The strip chosen to merge the two signals into, is the strip that before the correction had the larger signal. The other strip signal is forced to zero. The signals thus forced to zero are indicated on Figure 8.7 as the removed signals. Both strips are marked as processed, and are processed no further in the reconstruction of the event. The result is a drastic clean-up of the low-signal entries, as seen on Figure 8.7.

Note, that the choice of the sharing cut Ξ_1 has some influence on extracted multiplicity cuts. However, the effect has been estimated to be around 5% and can be further improved by more meticulous studies taking into account η dependencies, and similar effects [102].

Once the data has been corrected for sharing as above, one can define a set of cuts corresponding to 1 and 2 particle hits. This is done as a three step process.

1. A Landau distribution is fitted to the single strip multiplicity spectrum in a narrow range around the maximum peak of the spectrum. The range is defined as $[0, \bar{m} + s/2]$ where \bar{m} is the mean and s is the standard deviation of the spectrum.
2. Another Landau distribution is then fitted to the spectrum in the range $[\xi'_1 + 3\sigma'_1, \bar{m}' + s']$ where ξ'_1 is the most probable value and σ'_1 is the variance of the Landau distribution from the previous step, and \bar{m}' is the average and s' the variance of the spectrum above $\xi'_1 + 3\sigma'_1$.
3. The 6 parameters $(A'_1, \xi'_1, \sigma'_1)$ and $(A'_2, \xi'_2, \sigma'_2)$, where A'_i is the over-all scale of the two Landau distributions, and the second set is from the fit in second step above, is then used to seed a the fit of a double Landau distribution to the full spectrum.

From the combined fit, we can now extract (A_1, ξ_1, σ_1) and (A_2, ξ_2, σ_2) for the de-convoluted 1 particle Landau and 2 particle Landau distributions. We can use these to define 1 particle and 2 particle cuts e.g.,

$$\begin{cases} 0 & 0 \leq m < \xi_1 - 3\sigma_1 \\ 1 & \xi_1 - 3\sigma_1 \leq m < \xi_2 - 2\sigma_2 \\ 2 & \xi_2 - 2\sigma_2 \leq m < \xi_2 - 2\sigma_2 + \Delta\xi \end{cases}$$

where $\Delta\xi = \xi_2 - \xi_1$. This is shown in Figure 8.8. Since the sum of n Landau distributed random variates with $\xi = \xi'$ and $\sigma = \sigma'$ is itself Landau distributed with $\xi = n\xi'$ and $\sigma = n\sigma'$, we can use the above to define higher cuts, simply by adding $\Delta\xi$ to the last cut. However, since at least the simulations show a very long single and double particle hit Landau tail, it is not reasonable to simply divide up the spectrum according to $\Delta\xi$, as there is a relatively high probability that a single or two particles depositing a rather large energy signal, as compared to the probability of more than one or two particle hits.

Therefore, a cut value Ξ_2 is defined as

$$P(\Xi_2, 1 \text{ particle}) < P(\Xi_2, 2 \text{ particle})$$

where $P(x)$ are the probability distributions for x . That is, the value where the Landau probability distribution for 1 particle is smaller than the Landau distribution for 2 particles. The cut is indicated as the dashed line in Figure 8.8(b). Any signal above the value of Ξ_2 is taken to be a double particle hit. On the other end of the spectrum, any signal less than Ξ_1 is taken to be no hit. Figure 8.9 shows the correspondence of the result of this analysis to the GEANT 3 truth. Disregarding the edges of the FMD rings, we see a correspondence on the level of 5% to 10%. The result can be drastically improved increasing the statistics, which will allow for a 3 particle cut to be defined [102].




Figure 8.8: (a) Analysing the single strip multiplicity for single and double particle hits. First, a Landau distribution is fitted to the data around the maximum peak, then a Landau distribution is fitted to the rest of the spectrum, and the parameters of these two functions are used to seed a combined 1+2 particle Landau to the full spectrum. (b) The de-convoluted 1 and 2 particle Landau distributions, as well as a derived 3 particle Landau distribution.

8.7 Alternative for Low Multiplicity Events

The method for extracting the inclusive charged particle density $\frac{d^2 M_{ch}}{d\eta d\varphi}$ is suitable for all kinds of events. However, for events with a lower hit density, like peripheral Pb–Pb, lighter A–A, or pp, one can optionally use a method based on counting statistics. The distribution of number of hits (m) on a strip is the Poisson distribution

$$P\{m\} = \frac{\lambda^m e^{-\lambda}}{m!},$$

where λ is the mean number of hits. The probability for no hits on a strip is $P\{0\} = e^{-\lambda}$. Since

$$P\{0\} = \frac{N_{empty}}{N_{total}},$$

where N_{empty} and N_{total} are the number of empty strips and the total number of strips in the selected region of the detector, respectively, the average multiplicity, λ , can be determined as

$$\lambda = -\log \frac{N_{empty}}{N_{total}} \quad (8.10)$$

The method goes as follows.

- Select regions of the FMD full acceptance i.e., (η, φ) bins of reasonable size — $\Delta\varphi = 18^\circ$ for inner rings, and $\Delta\varphi = 9^\circ$ for outer rings, while $\Delta\eta$ typically will correspond to 64, 128, 256, or 512 strips. The requirement to the size of the region is, that the number of hit strips is relatively uniformly distributed across the region, and that there are more empty strips than full.
- Then for each event count how many strips did *not* have a signal larger than some threshold — say Ξ_1 . These strips are considered to have been hit by *no* particle i.e., empty strips.

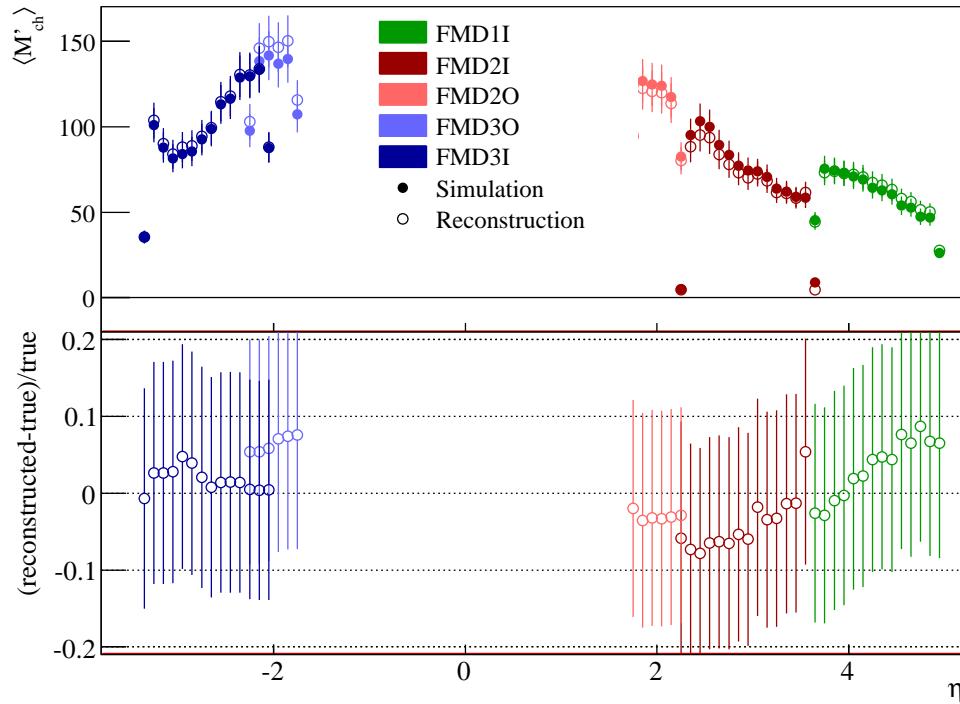


Figure 8.9: Reconstruction of the charged particle density. The top panel shows the reconstructed charged particle density (including secondaries) as well as the GEANT 3 true charged particle density. The bottom panel shows the scaled difference of the reconstructed charged particle density to the GEANT 3 truth. Note, that the error bars on lower plot are calculated that the error the reconstructed and true distributions are not correlated. However, as the reconstructed distribution is derived from the simulated one, as described in the text, this assumption is obviously not true. It is, however, not a straight forward matter to de-convolute the correlation, and too much emphasis should not be put on the error bars.

- Calculate the average multiplicity in the region by (8.10).

Figure 8.10 shows the inclusive charged particle density calculated using this Poisson method compared to the GEANT 3 truth. By combining the analysis based on the single strip multiplicity and the Poisson method one can estimate the systematic error of the analysis, which is typically on the order of a few percent [102].

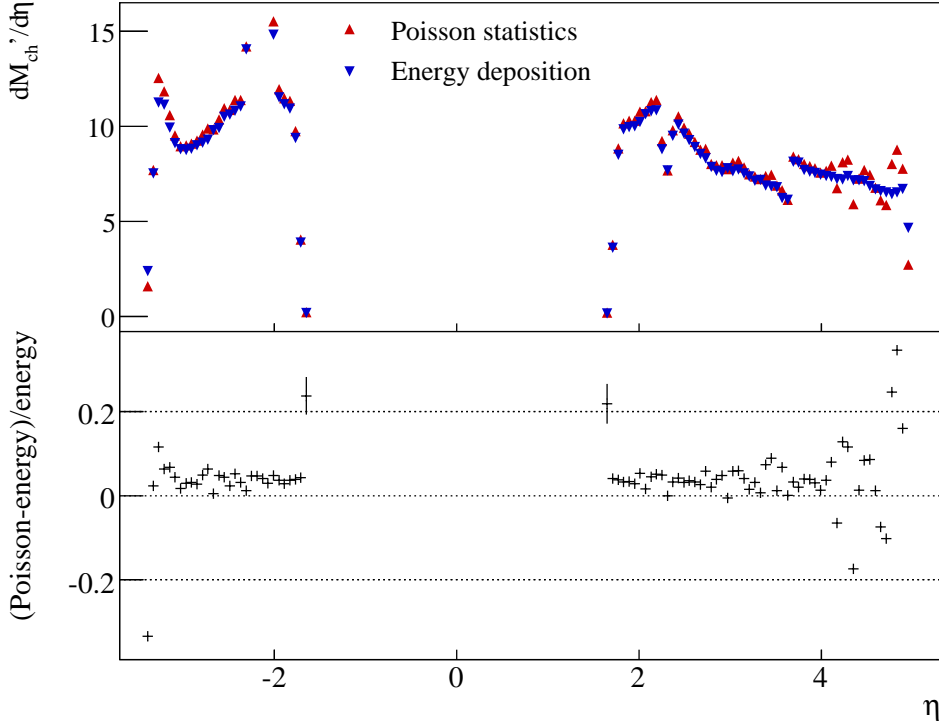


Figure 8.10: Comparison of methods to reproduce the bare charged particle multiplicity in pp. The input distribution came from 100 000 PYTHIA events at $\sqrt{s_{NN}} = 14000$ GeV [102].

8.8 Acceptance and Secondary Particle Corrections

To properly scale the data (η, φ) charged particle multiplicity to obtain $dM_{ch}/d\eta$, one must correct for the incomplete (η, φ) acceptance of the FMD, including the small overlaps between the various rings. This is done by running a simulation of the FMD detector *only* and throwing a flat distribution of charged particles at the FMD sub-detectors. Also, processes that can produce secondary particles are turned off to ensure that only the geometric acceptance is investigated. This procedure is repeated for different values of the z coordinate v_z of the primary interaction point. The ratio of detected hits in the FMD sector to the number of particles thrown in an (η, φ) bin is then the acceptance $A_{v_z}(\eta, \varphi)$ of the FMD in that bin for a given range of z coordinates of the interaction point.

The result shown in Figure 8.9 is *inclusive*, in the sense that it contains contributions from not only particles created in the primary interaction, but also from particles created in secondary interactions in surrounding material and through decays. To obtain the physical $dM_{ch}/d\eta$ we need to make another correction for these *secondary* particles.

The method to make this contribution is based around a detector simulation. A generator is run, and the particles are propagated through the geometric description of ALICE. Whenever a new particle is created, its point of creation is noted down and the particle is marked as a secondary. Then, when a particle hits the FMD sensors, the code asks whether the particle was

a secondary or not, and sum these contributions. This is done in (η, φ) bins, as the surround material is not necessarily evenly distributed in φ . Figure 8.11 shows the relative strength of secondaries as produced in various sources. Clearly, a large fraction of the secondaries seen in the FMD are from the ITS, especially the services of the ITS, and the beam pipe. As for the acceptance correction, this needs to be done in bins of the z coordinate of the primary interaction point, as the amount of material traversed is different for different values of v_z . We end up with a set of two dimensional maps $S_{v_z}(\eta, \varphi)$, one for each primary interaction v_z range and binned in (η, φ) .




Figure 8.11: Relative strength of sources of secondary particles seen by the FMD [102].

8.9 Charged Particle Multiplicity Density

Given the inclusive multiplicity, the acceptance map, and the secondary particle correction, all in (η, φ) we can construct the charged particle multiplicity density $dM_{\text{ch}}/d\eta$. The first step is to determine for each event with the primary interaction at v_x

$$\left. \frac{d^2 M_{\text{ch}}}{d\eta d\varphi} \right|_i = \left. \frac{d^2 M_{\text{ch}}'}{d\eta d\varphi} \right|_i \times A_{v_z}(\eta, \varphi) \times S_{v_z}(\eta, \varphi) \quad , \quad (8.11)$$

where $\left. \frac{d^2 M_{\text{ch}}'}{d\eta d\varphi} \right|_i$ is the inclusive charged particle density, $S_{v_z}(\eta, \varphi)$ is the secondary particle correction, and $A_{v_z}(\eta, \varphi)$ is the acceptance. These two dimensional distributions for each event are now directly comparable and can be average over many events. The final charged particle multiplicity density is then given by the integral

$$\frac{dM_{\text{ch}}}{d\eta} = \int_0^{2\pi} d\varphi \frac{1}{N} \sum_i^N \left. \frac{d^2 M_{\text{ch}}}{d\eta d\varphi} \right|_i \quad .$$

The result of this analysis, compared to the input HIJING simulation of 1000 events with $b \in [0, 5]$ fm, is shown in Figure 8.12.

8.10 Event Plane Angle and Azimuthal Anisotropy

The measurement of the azimuthal anisotropy v_2 (see Section 2.3) has attracted quite a lot of attention in recent years, with a corresponding large number of publications on the subject.




Figure 8.12: The primary charged particle density as reconstructed by the FMD, and compared to the input of 1000 events of HIJING model data with $b \in [0, 5]$ fm. The error bars indicate the statistical error, while the grey boxes indicate the systematic error estimated by comparing to the Poisson method [102].

The FMD aims at providing an independent measurement of v_2 in the forward regions inaccessible to other ALICE detectors. Furthermore, we aim at providing an independent measurement of the event plane angle for the central barrel detectors to use in determining the v_2 in that region.

8.10.1 Event Plane

For each event j the event plane is given by

$$\Psi_2 = \tan^{-1} \left(\frac{\sum_i w_i \sin(2\varphi_i)}{\sum_i w_i \cos(2\varphi_i)} \right) ,$$

with appropriate weight factors w_i , and where the sum runs over all observations of (ϕ_i, w_i) . In the FMD analysis, the event plane is calculated in η bins, and the weight is set to

$$w_i = \frac{d^2 M_{\text{ch}}}{d\eta_i d\varphi_i} \Big|_j ,$$

that is the acceptance and background corrected charged particle density in bin $(d\eta_i d\varphi_i)$. However, as outlined in Section 2.3 this is only half the story. Crucial for subsequent analysis of the azimuthal anisotropy v_2 is the estimation of the event plane resolution (see also Section 2.3.1)

$$R_k(\chi) = \frac{\pi}{\chi} e^{-\chi^2/4} \left(I_{\frac{k-1}{2}}(\chi^2/4) + I_{\frac{k+1}{2}}(\chi^2/4) \right) , \quad (8.12)$$

where $I_n(x)$ is the modified Bessel function of the first kind.

To find a value for $R_k(\chi)$ we need to evaluate χ . Following [61, 103–105], this is done by dividing the event into 2 random sub-events, A and B , of equal size in each η bin, and then

calculate the event plane angle in either sub-event, Ψ_A and Ψ_B . In [103] the following expression is given for χ

$$\frac{k}{N} = \frac{e^{-\chi^2/2}}{2} \quad ,$$

where k is the number of events where $|\Psi_A - \Psi_B| > \pi/2$, and N is the total number of events used. From this we isolate χ in terms of the measurable quantities k and N

$$\chi = \mp \sqrt{-2 \log(2k/N)} \quad (8.13)$$

To evaluate the uncertainty of R_k we identify

$$y = \chi^2/4 \quad C = \frac{\sqrt{\pi}e^{-y}}{2\sqrt{2}}$$

and introduce the short hands $f_2(y) = I_{\frac{k-1}{2}}$ and $f_3(y) = I_{\frac{k+1}{2}}(y)$, so we can write (8.12) more compact as $R_k(y) = Cy(f_2(y) + f_3(y))$. The derivative of the resolution function is

$$\frac{dR_k(y)}{dy} = \frac{C}{2} \left[4\sqrt{y} \left(\frac{f_2(y)}{dy} - \frac{f_3(y)}{dy} \right) - (4y - 2)(f_2(y) + f_3(y)) \right] \quad .$$

Since $\frac{I_\nu(x)}{dx} = I_{\nu-1}(x) - \frac{\nu}{x}I_\nu(x)$, and identifying $f_1(y) = I_{\frac{k-3}{2}}(y)$, we get

$$\frac{dR_k(y)}{dy} = \frac{C}{2} \left[4yf_1(y) + (4 - 2k)f_2(y) - (4y + 2k)f_3(y) \right] \quad .$$

As $r = k/N$ is obviously an efficiency-like ratio, the uncertainty δr is given by Binomial errors

$$\delta r = \sqrt{r \frac{1-r}{N}} \quad .$$

The total error $\delta\chi$ then becomes

$$\delta^2\chi = \left(\frac{d\chi}{dr} \right)^2 \delta^2r = \frac{r-1}{4k \log(2r)} \quad ,$$

and we find that the total error on R_k to be

$$\delta R_k = \frac{dR_k(\chi)}{d\chi} \delta\chi \quad .$$

100 parameterised HIJING events combined with a standard ALICE flow after-burner was simulated and sent through the detector simulation. The parameterised HIJING is a simple event generator that creates pions, kaons, protons, and neutrons according to simple parameterisations of the standard HIJING event generator. The after-burner reads the generated data and adds flow by modifying the transverse momentum of the particle, according to the prescription given in [61].

Within a bin, two sub-events of equal size were chosen randomly. That is, all non-zero signals after sharing and background correction, were shuffled at random and the first half was assigned to sub-event A and the last half to sub-event B . The reaction plane angles Ψ_A and Ψ_B were calculated independently in these two sub-events, and used to calculate R_k as outlined above.

Figure 8.13 shows the average distance from the simulated true event plane Ψ_R to the reconstructed second order event plane Ψ_2 for FMD1i, FMD2i, FMD2o, FMD3o, FMD3i, and all of the FMD. Also shown is the event plane resolution R_k for the same bins. Although $\langle \Psi_2 - \Psi_R \rangle$ is not exactly zero, the values are consistent, and within 2 variances consistent with zero. Further




Figure 8.13: Difference of reconstructed event plane Ψ_2 to true event plane Ψ_R . The input data was 100 events from a parameterised HIJING event generator with a standard flow after burner.

improvements is possible by better understanding of the secondary particle correction. The secondary particle correction is crucial in this analysis, as the secondary particles tend to be less correlated with the reaction plane than the primary particles are. Their contribution is therefore to suppress the asymmetry.

An reaction plane resolution of some 90 % is quite reasonable, and the FMD can consequently be used other detectors, like the ITS and TPC for determining the reaction plane of the collisions. Note, however, that proper determination of the reaction plane resolution, is conditional on the sub-event A and B samples are large enough and roughly of equal size. Preliminary studies show that a reasonable sub-event size is around 500 non-zero signals in each sub-event, or roughly a hit density of 10 % in a given FMD ring. With an expected average charged particle multiplicity of around 2000 at FMD1 in semi-peripheral to central events (Figure 8.12), these requirements should be fulfilled.

8.10.2 Azimuthal Anisotropy — v_2

As discussed in Section 2.3, the azimuthal anisotropy is measured by

$$v_2^{obs} = \langle \cos(2[\varphi - \Psi_2]) \rangle \quad , \quad (8.14)$$

where the average is taken over all observation of φ in all events, that is

$$v_2^{obs} = \frac{1}{\sum_i^{N_{event}} N_{obj,i}} \sum_i^{N_{event}} \sum_j^{N_{obs,i}} \cos(2[\varphi_j - \Psi_{2,i}]) \quad ,$$

where N_{event} is the number of events, and $N_{obs,i}$ is the number of observations of φ in event i . However, this observed value needs to be corrected for the event plane angle resolution by

$$v_2 = \frac{v_2^{obs}}{R_k} \quad ,$$

with R_k estimated as in Section 8.10.1 [61]. The uncertainty of the corrected observed value of v_2 is then given by

$$\begin{aligned} \delta^2 v_n &= \left(\frac{dv_2}{dv_2^{obs}} \right)^2 \delta^2 v_2^{obs} + \left(\frac{dv_2}{dR_k} \right)^2 \delta^2 R_k \\ &= \frac{\delta^2 v_2^{obs} R_k^2 + \delta^2 R_k (v_2^{obs})^2}{R_k^4} \quad , \end{aligned} \quad (8.15)$$

where $\delta^2 v_2^{obs}$ is given the observed standard deviation.

In analysing the flow signal from the FMD, it is important to note, that a 2, 3, or higher particle signal — after background correction — in a single strip must be treated as a corresponding number of observations of ϕ . That is, the sums in (8.10.2) are replaced by the weighted mean

$$v_2^{obs} = \frac{1}{\sum_i^{N_{event}} \sum_j^{N_{obj,i}} w_j} \sum_i^{N_{event}} \sum_j^{N_{obs,i}} w_j \cos(2[\varphi_j - \Psi_{2,i}]) \quad , \quad (8.16)$$

where, again, the charged particle density after background subtraction is chosen as the weights w_j . Figure 8.14 shows the result of analysing 100 events generated by a parameterised HIJING generator with a 6 % flow after-burner, and propagated through the detector simulation.

Clearly, the input flow of 6 % is not reproduced in all bins, though most are consistent with the input within 2 standard deviations. The reconstruction fails near the edges, and where the secondary particle correction is sketchy at best due to lack of statistics. Also, a better determination of the 1, 2, and especially, 3 particle cuts will certainly improve measurement




Figure 8.14: Reconstructed v_2 as a function of η .

considerably. The contribution from secondary particles is not easy to extract from the data, since it must necessarily be done in a purely statistical fashion. However, in a single event, the secondary contribution will tend to wash out the asymmetry, this lowering the observed v_2 . The seemingly complete break down at FMD1 is not fully understood at this time, but is most likely due to low counting statistics and an incomplete understanding of the background from secondaries.

Review

Various components of computer software has been developed to investigate all aspects of the FMD, including simulations, reconstruction of events from data, and analysis of events to produce results on the charged particle multiplicity $dM_{\text{ch}}/d\eta$, the event plane angle Ψ_R , and the azimuthal anisotropy v_2 . Simulations show that the FMD data is dominated by background from secondary particle production in external material, and that it is of the utmost importance that this background is understood. However, even with a coarse background correction, as used in the current chapter, it is still possible to extract somewhat reliable results, giving hope that future analysis can excel.

Chapter 9

Conclusion

From the publication of the technical design report [2] in the summer of 2004, to the installation of the three FMD sub-detectors almost 4 years have passed. The first sub-detector, FMD3, was installed and successfully commissioned in April 2007, and FMD2 followed in August 2007¹. The last sub-detector, FMD1 will be installed in late January 2008. In the time in between, from the technical design report to the physical installation of the detector, much development, construction, and testing has taken place.

After assembly of the silicon sensors with the hybrid boards, every single module has been tested, catalogued, and characterised. The result is that less than 1‰ of the strips are faulty or broken, and that we have a baseline to compare remaining tests to. Based on technologies developed for the ALICE TPC, a special digitisation card was developed. The mixed design of an extensible FPGA and hardware components ensures flexibility and adaptability while retaining a high degree of stability. The FMDD board was designed in close collaboration between engineers and physicists with the Physics requirements of the FMD foremost in mind. The chain of sensors and front-end electronics was tested thoroughly in the laboratory, and components, such as the ALTR0 bus cables, were modified and customised to meet the requirements of the FMD.

The full sensor and front-end electronics chains was tested on several occasions at the ASTRID storage ring in Århus, Denmark. The tests showed, that the full chain performs rather better than initially required by the Physics goals set forth in the technical design report. A signal-to-noise ratio of $\approx 40 : 1$ for the inner type sensors and $\approx 23 : 1$ for the outer type sensors was found. The test beam also showed that the system is overall stable enough to run uninterrupted for more than 8 hours. Independent studies showed that detection efficiency to be better than 99.5%. After installation of FMD3 in the ALICE pit at CERN, the full chain was once again tested, and showed no performance loss relative to previous tests in the laboratory and test beam at ASTRID. Due to the overall installation schedule it will not be possible to perform similar tests on FMD1 and 2 until February 2008.

During the first dress rehearsal of December 2007, the FMD3 sub-detector was tested with all the final hardware, including trigger system, cooling, data acquisition, and power supplies. The sub-detector was also partially integrated into the overall ALICE control system, and data collected successfully. The mechanism for automatic online calibrations of pedestals and gains, including writing the calibration parameters to globally available storage was tested and found successful. The cooling scheme was found to be performing well, with some 30°C in the area around the sensors and front-end electronics. The power supply lines was found to be limiting in that there is too large a potential drop over the some 80 m from the power supply to the detector. Plans are being discussed for increasing the cross section of these cables to lower the resistance and thereby the potential drop.

What remains for the full integration of the FMD systems in the overall ALICE system, is the

¹See photos of the installed detectors in Figure F.4, and a picture of some very happy FMD team members in Figure F.5.

finalisation of the detector control system at the supervisory level — in particular the various state machines that controls the behaviour the various hardware components is missing. Active development is taking place to have this ready for the next dress rehearsal in February 2008.

The full software chain, as sketched in Figure 8.1, was developed and tested, as wells as tools for analysing the FMD data of reconstructed events for the charged particle multiplicity, event plane angle, and asymmetric anisotropy. The core of the software has been tested and verified to reproduce the signals from the detector satisfactory. However, as simulation indicates that a rather large fraction of the signals recorded in the FMD sub-detectors come from the secondary particle background, it is crucial to understand this background better. The tools for these studies have been developed, although the studies themselves have not been carried out to a high enough detail.

The forth-coming pp run will provide an excellent opportunity to study background effects in the FMD. The low multiplicity, and corresponding low hit density in the FMD, provides a relatively clean data sample which can be studied and understood more easily than a high multiplicity A–A data sample. The pp run will also provide insight into the exact energy loss distribution, and thus provide data for a refinement of 1 particle cuts. Coupled with an analysis based on a Poisson distribution, this could bring the uncertainty of these cuts down considerably, and give a handle on the systematic errors of the analysis.

Clearly, the density of charged particles from pp is not high enough to reconstruct any kind of event plane angle nor any azimuthal anisotropy. The Physics results is therefore restricted to $dM_{\text{ch}}/d\eta$ for the pp run. It is, however, in and of itself an important analysis for understanding the A–A data, as it provides a baseline to compare to.

In conclusion, the FMD can be expected to deliver high-quality data in a unique η range for the ALICE experiment. It has a high efficiency, high signal-to-noise, and overall good and stable performance. The up-coming dress rehearsals should see a fully operational and integrated FMD.

Christian Holm Christensen
Copenhagen, December 2007

Bibliography

- [1] ALICE Collaboration, A. Ahmad *et al.*, *ALICE: Technical proposal for A Large Ion Collider Experiment at the CERN LHC*. No. 3 in LHC Tech. Proposal. CERN, Geneva, 1995. <http://cdsweb.cern.ch/record/293391>.
- [2] ALICE Collaboration, I. G. Bearden, H. Bøggild, C. H. Christensen, J. J. Gaardhøje, J.-Y. Grossiord, T. Malkeiwicz, B. S. Nielsen, G. Paic, and T. Wladyslaw, eds., *ALICE technical design report on forward detectors: FMD, T0 and V0*. CERN, 2004. CERN-LHCC-2004-025.
- [3] B. B. Back *et al.*, ‘The significance of the fragmentation region in ultrarelativistic heavy ion collisions,’ *Phys. Rev. Lett.*, **91** (2003) 052303, [nucl-ex/0210015](#).
- [4] PHOBOS Collaboration, B. B. Back *et al.*, ‘Charged-particle pseudorapidity distributions in Au + Au collisions at $s(\text{NN})^{1/2} = 62.4\text{-GeV}$,’ *Phys. Rev.*, **C74** (2006) 021901, [nucl-ex/0509034](#).
- [5] G. Gabrielse, D. Hanneke, T. Kinoshita, M. Nio, and B. Odom, ‘New Determination of the Fine Structure Constant from the Electron g Value and QED,’ *Phys. Rev. Lett.*, **97** (2006) 030802.
- [6] G. Gabrielse, D. Hanneke, T. Kinoshita, M. Nio, and B. Odom, ‘Erratum: New Determination of the Fine Structure Constant from the Electron g Value and QED,’ *Phys. Rev. Lett.*, **99** (2007) 039902.
- [7] G. S. Bali, K. Schilling, C. Schlichter, and A. Wachter, ‘Confining Forces and String Formation from the Lattice,’ [hep-lat/9410002](#).
- [8] M. E. Peskin and D. V. Schroeder, *Introduction to Quantum Field Theory*. Addison-Wesley, Reading, USA, 1997.
- [9] B. Muller, ‘Physics of the Quark-Gluon Plasma,’ [nucl-th/9211010](#).
- [10] ALEPH Collaboration, R. Barate *et al.*, ‘Studies of quantum chromodynamics with the ALEPH detector,’ *Phys. Rept.*, **294** (1998) 1–165.
- [11] Z. Fodor and S. D. Katz, ‘Lattice QCD results at finite T and μ ,’ *Heavy Ion Phys.*, **18** (2003) 41–48, [hep-lat/0204029](#).
- [12] T. Blum, L. Karkkainen, D. Toussaint, and S. Gottlieb, ‘The Beta Function and Equation of State for QCD with Two Flavors of Quarks,’ *Phys. Rev.*, **D51** (1995) 5153–5164, [hep-lat/9410014](#).
- [13] G. Boyd *et al.*, ‘Equation of State for the SU(3) Gauge Theory,’ *Phys. Rev. Lett.*, **75** (1995) 4169–4172, [hep-lat/9506025](#).
- [14] J. W. Harris and B. Muller, ‘The Search for the Quark-Gluon Plasma,’ *Ann. Rev. Nucl. Part. Sci.*, **46** (1996) 71–107, [hep-ph/9602235](#).

- [15] H. Satz, ‘QCD and QGP: A summary,’ [hep-ph/9706342](#).
- [16] P. Braun-Munzinger and J. Stachel, ‘Particle ratios, equilibration, and the QCD phase boundary,’ *J. Phys.*, **G28** (2002) 1971–1976, [nucl-th/0112051](#).
- [17] J. J. Brehm and W. J. Mullin, *Introduction to the Structure of Matter*. John Wiley & Sons, Inc., 1989.
- [18] X.-N. Wang and M. Gyulassy, ‘HIJING: A Monte Carlo Model for Multiple Jet Production in p+p, p+A and A+A Collisions,’ *Phys. Rev.*, **D44** (1991) 3501–3516.
- [19] M. Gyulassy and X.-N. Wang, ‘HIJING 1.0: A Monte Carlo Program for Parton and Particle Production in High-Energy Hadronic and Nuclear Collisions,’ *Comput. Phys. Commun.*, **83** (1994) 307, [nucl-th/9502021](#).
- [20] BRAHMS Collaboration, I. G. Bearden *et al.*, ‘Rapidity dependence of antiproton to proton ratios in Au+Au collisions at $\sqrt{s_{NN}} = 130$ GeV,’ *Phys. Rev. Lett.*, **87** (2001) 112305, [nucl-ex/0106011](#).
- [21] BRAHMS Collaboration, I. G. Bearden *et al.*, ‘Rapidity Dependence of Charged Antiparticle-to-Particle Ratios in Au+Au Collisions at $\sqrt{s_{NN}} = 200$ GeV,’ [nucl-ex/0207006](#).
- [22] Z.-w. Lin, S. Pal, C. M. Ko, B.-A. Li, and B. Zhang, ‘Multiphase Transport Model for Heavy Ion Collisions at RHIC,’ *Nucl. Phys.*, **A698** (2002) 375–378, [nucl-th/0105044](#).
- [23] BRAHMS Collaboration, I. G. Bearden *et al.*, ‘Rapidity dependence of charged antihadron to hadron ratios in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV,’ *Phys. Rev. Lett.*, **90** (2003) 102301.
- [24] BRAHMS Collaboration, I. G. Bearden *et al.*, ‘Nuclear stopping in Au + Au collisions at $\sqrt{s_{NN}} = 200$ GeV,’ *Phys. Rev. Lett.*, **93** (2004) 102301, [nucl-ex/0312023](#).
- [25] BRAHMS Collaboration, I. Arsene *et al.*, *Nuclear stopping in Au + Au collisions at $\sqrt{s_{NN}} = 63$ GeV* Submitted for publication, 2007.
- [26] BRAHMS Collaboration, I. G. Bearden *et al.*, ‘Pseudorapidity distributions of charged particles from Au+Au collisions at the maximum RHIC energy,’ *Phys. Rev. Lett.*, **88** (2002) 202301, [nucl-ex/0112001](#).
- [27] J. Bondorf, T. Døssing, J. J. Gaardhøje, H. Heiselberg, A. Jackson, B. Lauritzen, I. Mishustin, and H. Schulz, *Topics in Nuclear Physics*. Niels Bohr Institute, University of Copenhagen, 1999.
- [28] STAR Collaboration, R. Snellings, ‘Single Particle Inclusive Spectra, HBT and Elliptic Flow: A Consistent Picture at RHIC?,’ [hep-ph/0111437](#).
- [29] BRAHMS Collaboration, I. Arsene *et al.*, ‘Transverse momentum spectra in Au+Au and d+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and the pseudorapidity dependence of high p_T suppression,’ *Phys. Rev. Lett.*, **91** (2003) 072305, [nucl-ex/0307003](#).
- [30] BRAHMS Collaboration, I. G. Bearden *et al.*, ‘Charged meson rapidity distributions in central Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV,’ *Phys. Rev. Lett.*, **94** (2005) 162301, [nucl-ex/0403050](#).
- [31] BRAHMS Collaboration, I. Arsene *et al.*, ‘Centrality dependent particle production at $y = 0$ and $y \approx .1$ in Au–Au collisions at $\sqrt{s_{NN}} = 200$ GeV,’ *Phys. Rev.*, **C72** (2005) 014908, [nucl-ex/0503010](#).

- [32] STAR Collaboration, H. Z. Huang, ‘Anti-baryon to Baryon Ratios in Au+Au Collisions at $\sqrt{s_{NN}} = 130$ GeV,’ *Nucl. Phys.*, **A698** (2002) 663–666.
- [33] PHENIX Collaboration, H. Ohnishi, ‘Particle Ratios in PHENIX at RHIC,’ *Nucl. Phys.*, **A698** (2002) 659–662.
- [34] PHOBOS Collaboration, B. B. Back *et al.*, ‘Ratios of Charged Antiparticles to Particles Near Mid- Rapidity in Au+Au collisions at $\sqrt{s_{NN}} = 130$ GeV,’ *Phys. Rev. Lett.*, **87** (2001) 102301, [hep-ex/0104032](#).
- [35] NA50 Collaboration, M. C. Abreu *et al.*, ‘Evidence for Deconfinement of Quarks and Gluons from the J/psi Suppression Pattern Measured in Pb+Pb Collisions at the CERN-SPS,’ *Phys. Lett.*, **B477** (2000) 28–36.
- [36] STAR Collaboration, C. Adler *et al.*, ‘Centrality dependence of high p(T) hadron suppression in Au + Au collisions at $s(NN)^{*(1/2)} = 130$ -GeV,’ *Phys. Rev. Lett.*, **89** (2002) 202301, [nucl-ex/0206011](#).
- [37] STAR Collaboration, J. Adams *et al.*, ‘Evidence from d + Au measurements for final-state suppression of high p(T) hadrons in Au + Au collisions at RHIC,’ *Phys. Rev. Lett.*, **91** (2003) 072304, [nucl-ex/0306024](#).
- [38] STAR Collaboration, J. Adams *et al.*, ‘Azimuthal anisotropy and correlations at large transverse momenta in p + p and Au + Au collisions at $s(NN)^{*(1/2)} = 200$ -GeV,’ *Phys. Rev. Lett.*, **93** (2004) 252301, [nucl-ex/0407007](#).
- [39] STAR Collaboration, J. Adams *et al.*, ‘Direct observation of dijets in central Au + Au collisions at $s(NN)^{*(1/2)} = 200$ -GeV,’ *Phys. Rev. Lett.*, **97** (2006) 162301, [nucl-ex/0604018](#).
- [40] STAR Collaboration, J. Adams *et al.*, ‘Experimental and theoretical challenges in the search for the quark gluon plasma: The STAR collaboration’s critical assessment of the evidence from RHIC collisions,’ *Nucl. Phys.*, **A757** (2005) 102–183, [nucl-ex/0501009](#).
- [41] C. Ristean, ‘The Rapidity Dependence of High p_{\perp} Suppression in Au–Au Collisions at BRAHMS Experiment,’. PhD thesis, Niels Bohr Institute, 2007. http://www.nbi.dk/~ristea/PhD_CR_Thesis.pdf.
- [42] T. M. Larsen, ‘High Momentum Jet Quenching in $\sqrt{s_{NN}} = 62.4$ GeV Cu–Cu and Au–Au Collisions,’. PhD thesis, Niels Bohr Institute, 2007. <http://www.nbi.dk/~trulsm1/phdthesis.pdf>.
- [43] P. F. Kolb and U. W. Heinz, ‘Hydrodynamic description of ultrarelativistic heavy-ion collisions,’ [nucl-th/0305084](#).
- [44] V. Koch, ‘Aspects of Chiral Symmetry,’ *Int. J. Mod. Phys.*, **E6** (1997) 203–250, [nucl-th/9706075](#).
- [45] PHENIX Collaboration, S. Huang, ‘Measurements of phi mesons reconstructed from hadronic and leptonic decays by the PHENIX experiment at RHIC,’.
- [46] M. Bleicher, J. Randrup, R. Snellings, and X. N. Wang, ‘Enhanced event-by-event fluctuations in pion multiplicity as a signal of disoriented chiral condensates at RHIC,’ *Phys. Rev.*, **C62** (2000) 041901, [nucl-th/0006047](#).
- [47] STAR Collaboration, J. Adams *et al.*, ‘Production of charged pions and hadrons in Au + Au collisions at $s(NN)^{*(1/2)} = 130$ -GeV,’ [nucl-ex/0311017](#).

- [48] K. Yagi, T. Hatsuda, and Y. Miake, *Quark-gluon plasma: From big bang to little bang*. No. 23 in Cambridge monographs on particle physics, nuclear physics and cosmology. Cambridge University Press, Cambridge, 2005.
- [49] B. B. Back *et al.*, ‘The PHOBOS perspective on discoveries at RHIC,’ *Nucl. Phys.*, **A757** (2005) 28–101, [nucl-ex/0410022](#).
- [50] BRAHMS Collaboration, I. Arsene *et al.*, ‘Quark gluon plasma and color glass condensate at RHIC? The perspective from the BRAHMS experiment,’ *Nucl. Phys.*, **A757** (2005) 1–27, [nucl-ex/0410020](#).
- [51] J. D. Bjorken, ‘Highly Relativistic Nucleus-Nucleus Collisions: The Central Rapidity Region,’ *Phys. Rev.*, **D27** (1983) 140–151.
- [52] D. Kharzeev, E. Levin, and M. Nardi, ‘Color glass condensate at the LHC: Hadron multiplicities in p p, p A and A A collisions,’ *Nucl. Phys.*, **A747** (2005) 609–629, [hep-ph/0408050](#).
- [53] J. C. Collins, D. E. Soper, and G. Sterman, ‘Factorization for Short Distance Hadron - Hadron Scattering,’ *Nucl. Phys.*, **B261** (1985) 104.
- [54] D. Kharzeev, E. Levin, and L. McLerran, ‘Parton saturation and N_{part} scaling of semi-hard processes in QCD,’ *Phys. Lett.*, **B561** (2003) 93–101, [hep-ph/0210332](#).
- [55] L. McLerran, ‘The color glass condensate to the strongly interacting quark gluon plasma: Theoretical developments,’ *Nucl. Phys.*, **A787** (2007) 1–8.
- [56] L. McLerran, ‘The color glass condensate: An intuitive description,’ *Int. J. Mod. Phys.*, **A21** (2006) 694–698.
- [57] L. McLerran, , ch. The Color Glass Condensate and Small-x Physics, pp. 291–334. No. 583 in *Lecture Notes in Physics*. Springer, Berlin / Heidelberg, 2002.
- [58] L. McLerran, ‘The color glass condensate and RHIC,’ *Nucl. Phys.*, **A752** (2005) 355–371.
- [59] D. Kharzeev and E. Levin, ‘Manifestations of high density QCD in the first RHIC data,’ *Phys. Lett.*, **B523** (2001) 79–87, [nucl-th/0108006](#).
- [60] S. Voloshin and Y. Zhang, ‘Flow study in relativistic nuclear collisions by Fourier expansion of Azimuthal particle distributions,’ *Z. Phys.*, **C70** (1996) 665–672, [hep-ph/9407282](#).
- [61] A. M. Poskanzer and S. A. Voloshin, ‘Methods for analyzing anisotropic flow in relativistic nuclear collisions,’ *Phys. Rev.*, **C58** (1998) 1671–1678, [nucl-ex/9805001](#).
- [62] P. Danielewicz, ‘Effects of compression and collective expansion on particle emission from central heavy ion reactions,’ *Phys. Rev.*, **C51** (1995) 716–750, [nucl-th/9408018](#).
- [63] e. . d’Enterria, D. *et al.*, ‘CMS physics technical design report: Addendum on high density QCD with heavy ions,’ *J. Phys.*, **G34** (2007) 2307–2455.
- [64] PHENIX Collaboration, K. Adcox *et al.*, ‘Formation of dense partonic matter in relativistic nucleus nucleus collisions at RHIC: Experimental evaluation by the PHENIX collaboration,’ *Nucl. Phys.*, **A757** (2005) 184–283, [nucl-ex/0410003](#).
- [65] BRAHMS Collaboration, H. Ito, ‘Rapidity dependence of pion elliptic flow,’ *Nucl. Phys.*, **A774** (2006) 519–522.

- [66] PHOBOS Collaboration, B. B. Back *et al.*, ‘Centrality and pseudorapidity dependence of elliptic flow for charged hadrons in Au + Au collisions at $\sqrt{s_{NN}} = 200$ -GeV,’ *Phys. Rev.*, **C72** (2005) 051901, nucl-ex/0407012.
- [67] PHOBOS Collaboration, B. Alver *et al.*, ‘System size, energy, pseudorapidity, and centrality dependence of elliptic flow,’ *Phys. Rev. Lett.*, **98** (2007) 242302, nucl-ex/0610037.
- [68] PHENIX Collaboration, A. Adare *et al.*, ‘Scaling properties of azimuthal anisotropy in Au–Au and Cu + Cu collisions at $\sqrt{s_{NN}} = 200$ GeV,’ *Phys. Rev. Lett.*, **98** (2007) 162301, nucl-ex/0608033.
- [69] C. Vanoli, ‘Cern accelerator complex,’ 2006. <http://documents.cern.ch/cgi-bin/setlink?base=PH0&categ=photo-di&id=0606052>.
- [70] O. S. Brüning, P. Collier, P. Lebrun, S. Myers, R. Ostojic, J. Poole, and P. Proudlock, *LHC Design Report*, vol. I: The LHC main ring. CERN, Geneva, 2004. <http://cdsweb.cern.ch/record/782076>.
- [71] M. Benedikt, P. Collier, V. Mertens, J. Poole, and K. Schindl, *LHC Design Report*, vol. III: The LHC Injector Chain. CERN, Geneva, 2004. <http://cdsweb.cern.ch/record/823808>.
- [72] ALICE Collaboration, A. Andronic *et al.*, *ALICE transition-radiation detector Technical Design Report*. Technical Design Report ALICE. CERN, Geneva, 2001. <http://cdsweb.cern.ch/record/519145>.
- [73] P. Fonte, ‘Applications and new developments in resistive plate chambers,’ *IEEE Trans. Nucl. Sci.*, **49** (2002) 881–887.
- [74] R. Santonico and R. Cardarelli, ‘DEVELOPMENT OF RESISTIVE PLATE COUNTERS,’ *Nucl. Instrum. Meth.*, **187** (1981) 377–380.
- [75] E. Cerron Zeballos *et al.*, ‘A New type of resistive plate chamber: The Multigap RPC,’ *Nucl. Instrum. Meth.*, **A374** (1996) 132–136.
- [76] ALICE Collaboration, A. Alici *et al.*, *ALICE Time-Of Flight system (TOF) addendum to the Technical Design Report*. Technical Design Report ALICE. CERN, Geneva, 2002. <http://cdsweb.cern.ch/record/545834>.
- [77] W. R. Leo, *Techniques for Nuclear and Particle Physics Experiments: A How-to Approach*. Springer, Berlin, Germany, 1987.
- [78] T. Cormier, C. W. Fabjan, L. Riccati, and H. de Groot, *The Electromagnetic Calorimeter Addendum to the ALICE Technical Proposal* Tech. Rep. CERN-LHCC-2006-014. CERN-LHCC-96-32-ADD-3, CERN, Geneva, Mar, 2006.
- [79] V. I. Man’ko, W. Klempt, L. Leistam, J. De Groot, and J. Schkraft, *ALICE Photon Spectrometer (PHOS) Technical Design Report*. Technical Design Report ALICE. CERN, Geneva, 1999.
- [80] P. Azzurri, ‘The CMS silicon strip tracker,’ *J. Phys. Conf. Ser.*, **41** (2006) 127–134, physics/0512097.
- [81] Particle Data Group Collaboration, W.-M. Yao, C. Amsler, D. Asner, R. Barnett, J. Beringer, P. Burchat, C. Carone, C. Caso, O. Dahl, G. D’Ambrosio, A. DeGouvea, M. Doser, S. Eidelman, J. Feng, T. Gherghetta, M. Goodman, C. Grab, D. Groom,

- A. Gurtu, K. Hagiwara, K. Hayes, J. Hernández-Rey, K. Hikasa, H. Jawahery, C. Kolda, K. Y., M. Mangano, A. Manohar, A. Masoni, R. Miquel, K. Mönig, H. Murayama, K. Nakamura, S. Navas, K. Olive, L. Pape, C. Patrignani, A. Piepke, G. Punzi, G. Raffelt, J. Smith, M. Tanabashi, J. Terning, N. Törnqvist, T. Trippe, P. Vogel, T. Watari, C. Wohl, R. Workman, P. Zyla, B. Armstrong, G. Harper, V. Lgovsky, P. Schaffner, M. Artuso, K. Babu, H. Band, E. Barberio, M. Battaglia, H. Bichsel, O. Biebel, P. Bloch, E. Blucher, R. Cahn, D. Casper, A. Cattai, A. Ceccucci, D. Chakraborty, R. Chivukula, G. Cowan, T. Damour, T. DeGrand, K. Desler, M. Dobbs, M. Drees, A. Edwards, D. Edwards, V. Elvira, J. Erler, V. Ezhela, W. Fetscher, B. Fields, B. Foster, D. Froidevaux, T. Gaisser, L. Garren, H.-J. Gerber, G. Gerbier, L. Gibbons, F. Gilman, G. Giudice, A. Gribsan, M. Grünewald, H. Haber, C. Hagmann, I. Hinchliffe, A. Höcker, P. Igo-Kemenes, J. Jackson, K. Johnson, D. Karlen, B. Kayser, D. Kirkby, S. Klein, K. Kleinknecht, I. Knowles, R. Kowalewski, P. Kreitz, B. Krusche, Y. Kuyanov, O. Lahav, P. Langacker, A. Liddle, Z. Ligeti, T. Liss, L. Littenberg, L. Liu, K. Lgovsky, S. Lgovsky, T. Mannel, D. Manley, W. Marciano, A. Martin, D. Milstead, M. Narain, P. Nason, Y. Nir, J. Peacock, S. Prell, A. Quadt, S. Raby, B. Ratcliff, E. Razuvaev, B. Renk, P. Richardson, S. Roesler, G. Rolandi, M. Ronan, L. Rosenberg, C. Sachrajda, S. Sarkar, M. Schmitt, O. Schneider, D. Scott, T. Sjöstrand, G. Smoot, P. Sokolsky, S. Spanier, H. Spieler, A. Stahl, T. Stanev, R. Streitmatter, T. Sumiyoshi, N. Tkachenko, G. Trilling, G. Valencia, K. van Bibber, M. Vincet, D. Ward, B. Webber, J. Wells, M. Whalley, L. Wolfenstein, J. Womersley, C. Woody, A. Yamamoto, O. Zenin, J. Zhang, and R.-Y. Zhu, ‘Review of Particle Physics,’ *J. Phys.*, **G33** (2006) 1–1232.
- [82] O. Toker, S. Masciocchi, E. Nygard, A. Rudge, and P. Weilhammer, ‘VIKING: A CMOS low noise monolithic 128-channel front end for Si strip detector readout,’ *Nucl. Instrum. Meth.*, **A340** (1994) 572–579.
- [83] E. Nygard, P. Aspell, P. Jarron, P. Weilhammer, and K. Yoshioka, ‘CMOS low noise amplifier for microstrip readout: Design and results,’ *Nucl. Instrum. Meth.*, **A301** (1991) 506–516.
- [84] I. ASA, *VA1₃* tech. rep., IDEAS ASA, Høvik, Norway, September, 2004. Not published, available on request.
- [85] R. Esteve Bosch, A. Jimenez de Parga, B. Mota, and L. Musa, ‘The ALTRO chip: A 16-channel A/D converter and digital processor for gas detectors,’ *IEEE Trans. Nucl. Sci.*, **50** (2003) 2460–2469.
- [86] J. Baechler *et al.*, ‘Front-end electronics for the ALICE TPC-detector,’.
- [87] S. Brown and J. Rose, ‘Architecture of FPGAs and CPLDs: A Tutorial,’ *IEEE Design and Test of Computers*, **13** (1996), no. 2, 42–57.
- [88] B. Mota *et al.*, ‘Performance of the ALTRO chip on data acquired on an ALICE TPC prototype,’ *Nucl. Instrum. Meth.*, **A535** (2004) 500–505.
- [89] ALICE Collaboration, R. Esteve Bosch *et al.*, ‘Readout control unit of the front end electronics for the ALICE time projection chamber,’.
- [90] C. Gaspar and M. Donszelmann, ‘DIM: A distributed information management system for the DELPHI experiment at CERN,’.
- [91] M. Donszelmann and C. Gaspar, ‘The DELPHI distributed information system for exchanging LEP machine related information,’ *Nucl. Instrum. Meth.*, **A352** (1994) 280–282.

- [92] C. Gaspar and B. Franek, ‘Tools for the automation of large distributed control systems,’ *IEEE Trans. Nucl. Sci.*, **53** (2006) 974–979.
- [93] B. Franek and C. Gaspar, ‘SMI++: Object oriented framework for designing and implementing distributed control systems,’.
- [94] R. Stensgaard, ‘ASTRID - THE AARHUS STORAGE RING,’ *Phys. Scripta*, **T22** (1988) 315–317.
- [95] ALICE Collaboration, K. Gulbrandsen, I. Bearden, P. H. Bertelsen, C. H. Christensen, J. J. Gaardhøje, and B. S. Nielsen, ‘The ALICE forward multiplicity detector,’ *Nucl. Phys.*, **A774** (2006) 919–922.
- [96] C. Søgaard, ‘The ALICE Forward Multiplicity Detector — Test Beam Results,’ Master’s thesis, Niels Bohr Institute, University of Copenhagen, 2007.
<http://www.nbi.dk/~soegaard/thesis.pdf>.
- [97] ALICE Collaboration, C. H. Christensen, J. J. Gaardhøje, K. Gulbrandsen, B. S. Nielsen, and S. C., ‘The ALICE forward multiplicity detector,’ *Int. J. Mod. Phys.*, **E16** (2008) 2432–2437, [arXiv:0712.1117](https://arxiv.org/abs/0712.1117) [nucl-ex].
- [98] M. Paterno, *Calculating Efficiencies and Their Uncertainties* D0 notes 2861, FNAL, 2003. FERMILAB-TM-2286-CD, <http://home.fnal.gov/~paterno/images/effic.pdf>.
- [99] R. Brun and F. Rademakers, ‘ROOT - An Object Oriented Data Analysis Framework,’ in *Proceedings AIHENP’96 Workshop, Lausanne*, vol. A389, pp. 81–86. 1997. See also <http://root.cern.ch/>.
- [100] F. Carminati, C. W. Fabjan, L. Riccati, and H. de Groot, *ALICE computing Technical Design Report*. Technical Design Report ALICE. CERN, Geneva, 2005. Submitted on 15 Jun 2005.
- [101] ALICE Collaboration, , ‘ALICE Technical Design Report of the Inner Tracking System (ITS),’.
- [102] H. H. Dalsgaard, ‘The Forward Multiplicity Detector for ALICE,’ Master’s thesis, Niels Bohr Institute, University of Copenhagen, 2007.
<http://www.nbi.dk/~canute/thesis.pdf>.
- [103] J.-Y. Ollitrault, ‘On the measurement of azimuthal anisotropies in nucleus– nucleus collisions,’ [nucl-ex/9711003](https://arxiv.org/abs/nuc1-ex/9711003).
- [104] J.-Y. Ollitrault, ‘Collective flow from azimuthal correlations,’ *Nucl. Phys.*, **A590** (1995) 561c–564c.
- [105] J.-Y. Ollitrault, ‘Determination of the reaction plane in ultrarelativistic nuclear collisions,’ *Phys. Rev.*, **D48** (1993) 1132–1139, [hep-ph/9303247](https://arxiv.org/abs/hep-ph/9303247).

Appendix A

Variables

A.1 Physical Variables

For a particle with three-momentum $\mathbf{p} = (p_x, p_y, p_z)$ and mass m we have

$$\begin{aligned} p &\equiv \sqrt{p_x^2 + p_y^2 + p_z^2} \\ p_\perp &\equiv \sqrt{p_x^2 + p_y^2} \\ p_\parallel &\equiv p_z \\ m_\perp &\equiv \sqrt{m^2 + p_\perp^2} \\ E &= \sqrt{m^2 + p^2} \\ \beta &= \frac{v}{c} = \frac{p/m}{c} \\ \gamma &= \frac{1}{\sqrt{1 - \beta^2}} \end{aligned}$$

We define the rapidity y as

$$y \equiv \frac{1}{2} \log \frac{E + p_\parallel}{E - p_\parallel} . \quad (\text{A.1})$$

y has the nice property that it is additive, so that seen a particle seen moving with y in the frame S , is seen in the frame S' moving with y_β relative to S is simply given by

$$\begin{aligned} y' &= \frac{1}{2} \log \frac{p' + p'_\parallel}{p' - p'_\parallel} \\ &= \frac{1}{2} \log \frac{\gamma(1 - \beta)(p' + p'_\parallel)}{\gamma(1 + \beta)(p' - p'_\parallel)} \\ &= \frac{1}{2} \log \frac{p + p_\parallel}{p - p_\parallel} + \frac{1}{2} \log \frac{1 + \beta}{1 - \beta} \\ &= y - y_\beta \end{aligned}$$

When the energy of a particle is not known it is convenient to use the pseudorapidity η defined as:

$$\eta \equiv -\log \left(\tan \frac{\vartheta}{2} \right) , \quad (\text{A.2})$$

where $\vartheta \equiv \tan^{-1}(p_\perp/p_\parallel)$. For $m \ll p$, or $E \sim p$, it follows that

$$y = \frac{1}{2} \log \frac{E + p_\parallel}{E - p_\parallel} \sim \frac{1}{2} \log \frac{p + p_\parallel}{p - p_\parallel} = \frac{1}{2} \log \frac{1 + \cos \vartheta}{1 - \cos \vartheta} = \frac{1}{2} \log \tan^2 \frac{\vartheta}{2} = \eta$$

One often characterises the available energy in a collision by the centre-of-mass \sqrt{s} energy of the system. However, for nucleus-nucleus collision a more measure convenient measure is the centre-of-mass energy per colliding nucleon pair $\sqrt{s_{NN}}$. Sometimes one also sees the notation $\sqrt{s} = xA\text{GeV}$ which is the same as $\sqrt{s_{NN}} = x\text{ GeV}$.

Appendix B

Images of the ALICE Detectors

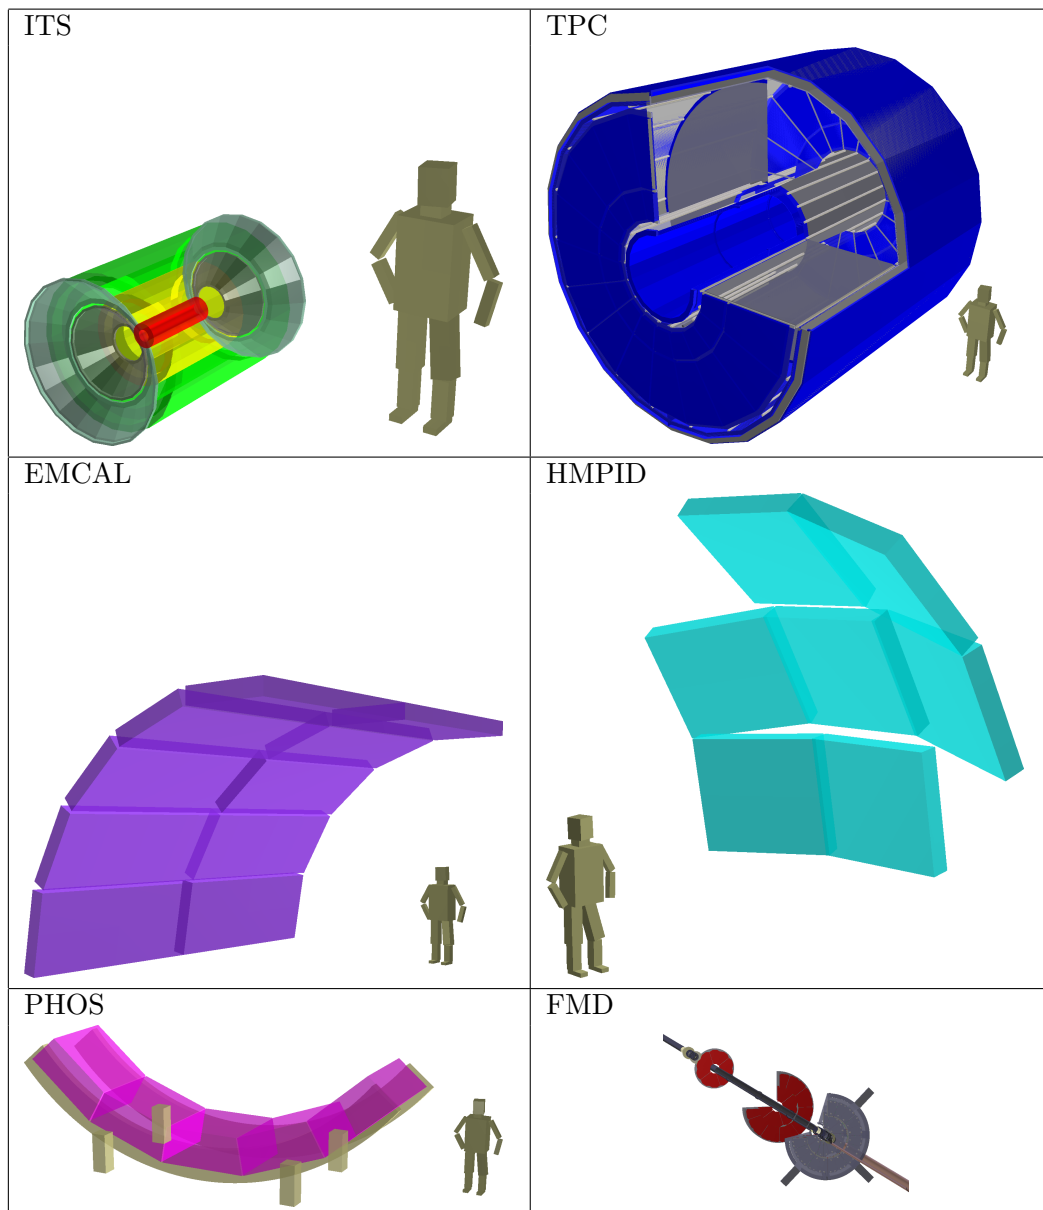


Figure B.1: ALICE detectors.




Figure B.2: The ALICE event display, showing a central Pb–Pb. Top panel shows the hits generated in the simulation, and the bottom panel shows the particle tracks through the detector simulation.

Appendix C

Additional Figures for Chapter 2




Figure C.1: Fragmentation region in Au–Au at $\sqrt{s_{NN}} = 19.4$ GeV, $\sqrt{s_{NN}} = 130$ GeV, and $\sqrt{s_{NN}} = 200$ GeV at various centralities, as measured by the PHOBOS experiment. Adapted from [3, 4].

Appendix D

Additional Figures for Chapter 5

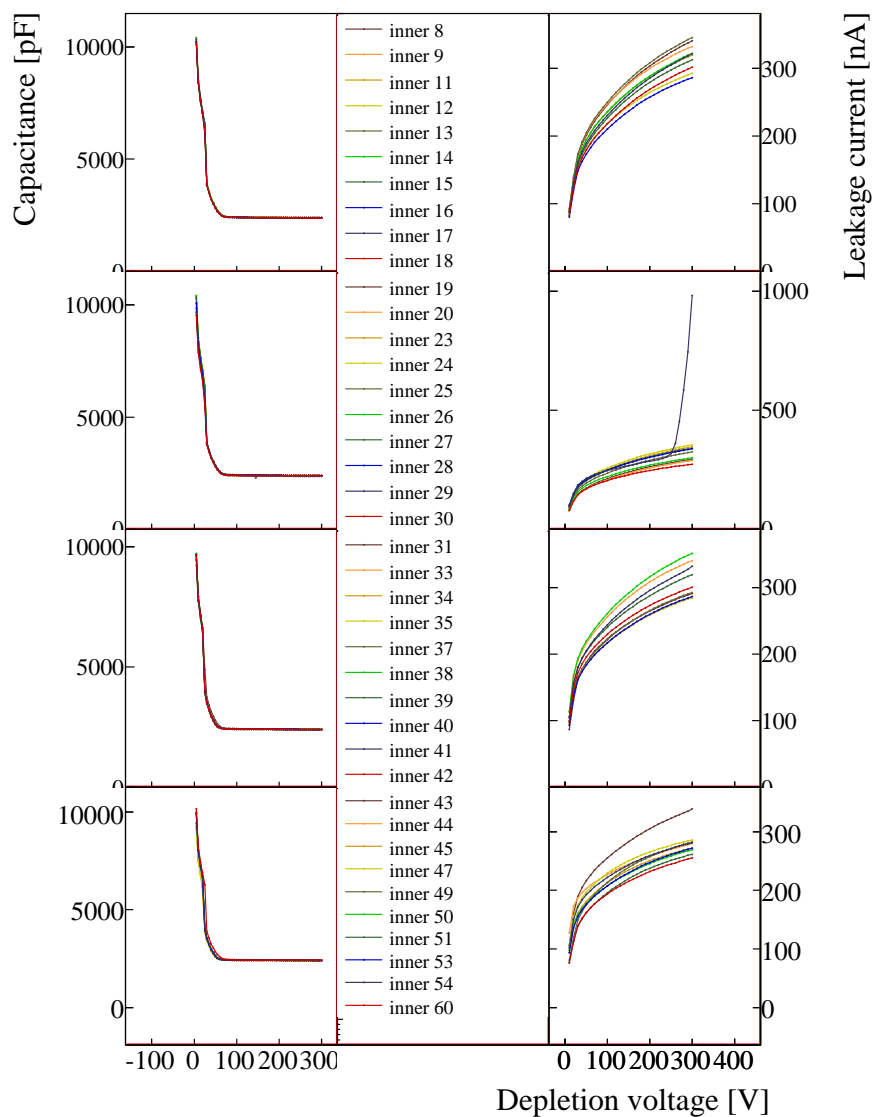


Figure D.1: Characteristics of all inner sensors, as measured by Hamamatsu. Capacitance and leakage current as a function of bias voltage for the inner sensors. In the 2nd panel from the top, on the right, one can clearly see the breakdown of the semi-conductor at a bias voltage of around 300 V.

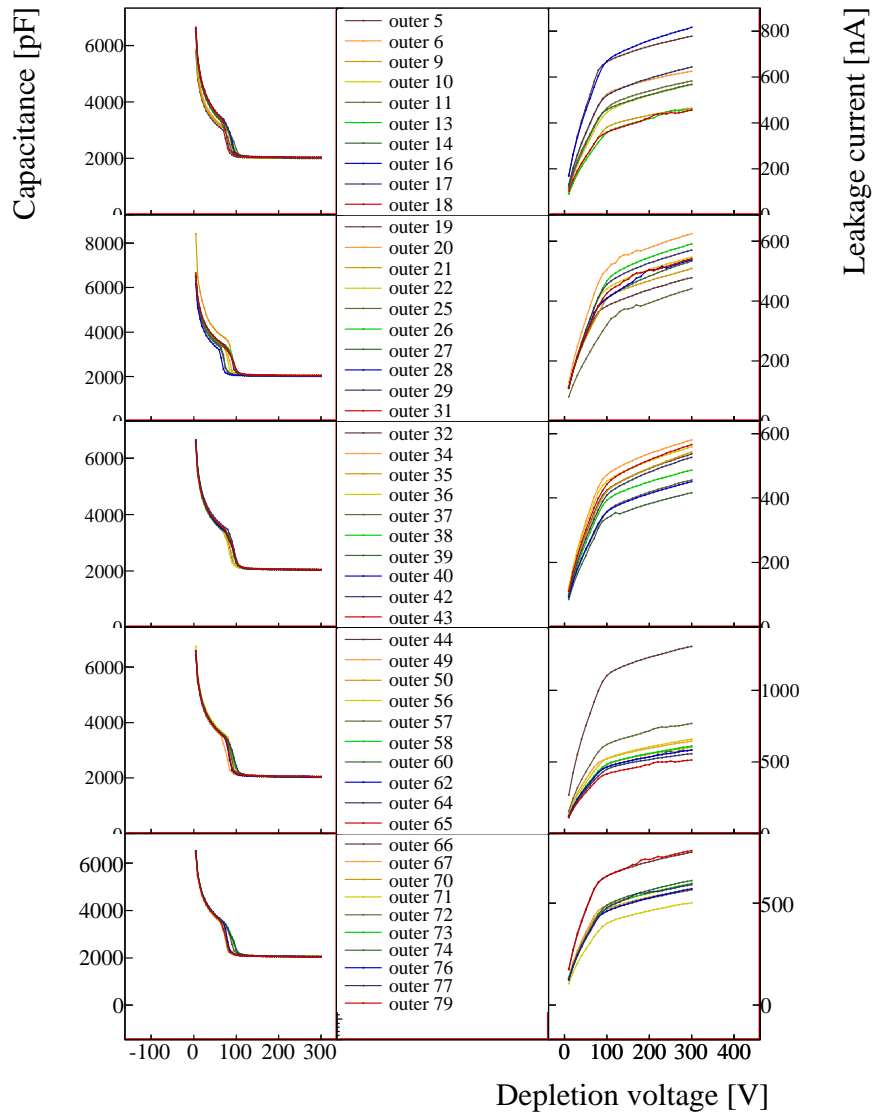


Figure D.2: Characteristics of all outer sensors, as measured by Hamamatsu. Capacitance and leakage current as a function of bias voltage for the outer sensors. In the 2nd panel from the top, on the right, one can clearly see the breakdown of the semi-conductor at a bias voltage of around 300 V.

External parameters	
Radiation dose (10 years)	5 000 Gy
Hadron flux (10 years)	$1 \times 10^{13} \text{ cm}^{-2}$
Neutron flux (10 years)	$2 \times 10^{12} \text{ cm}^{-2}$
1 MeV n eq. flux (10 years)	$3 \times 10^{13} \text{ cm}^{-2}$
Operational temperature	30 °C
Geometrical parameters	
Wafer diameter	150 mm
Effective sensor diameter	134 mm
Silicon thickness	$310 \pm 10 \mu\text{m}$
Number of radial strips	512 (inner) or 256 (outer)
Strip pitch	$250 \mu\text{m}$ (inner) or $500 \mu\text{m}$ (outer)
Strip length	13–50 mm (inner) or 24–42 mm (outer)
Guard and biasing ring width	$\sim 1 \text{ mm}$
Dimension of bonding pads	$\sim 100 \mu\text{m} \times 300 \mu\text{m}$
Silicon bulk parameters	
Silicon bulk type	n-type
Silicon lattice orientation ¹	$\langle 100 \rangle$
Silicon resistivity	$\sim 5 \text{ k}\Omega\text{cm}$
Silicon mask parameters	
Metal strip width	slightly larger than the p+ implant width
Metal strip thickness	$1 \mu\text{m}$ Al
p+ strip width/pitch ratio	0.20–0.25
Metal back side	$1 \mu\text{m}$ Al
Passivation on implant side	$1 \mu\text{m}$ PECVD
Alignment reference	reference mark on implant side mask
Sensor electrical parameters	
Full Depletion voltage	60–120 V
Operational voltage	70–130 V
Breakdown voltage	$> 200 \text{ V}$
Total leakage current	$< 1 \mu\text{A}$
Strip leakage current	$< 5 \text{ nA}$
Strip coupling capacitance	5–25 pF
Polysilicon bias resistors	$\sim 20 \text{ M}\Omega$
Bad strips	$< 1 \%$

Table D.1: Silicon sensor design parameters.

¹This is solid-state notation for the lattice orientation relative to the surface $\langle n_x, n_y, n_z \rangle$.

Appendix E

Additional Figures for Chapter 8

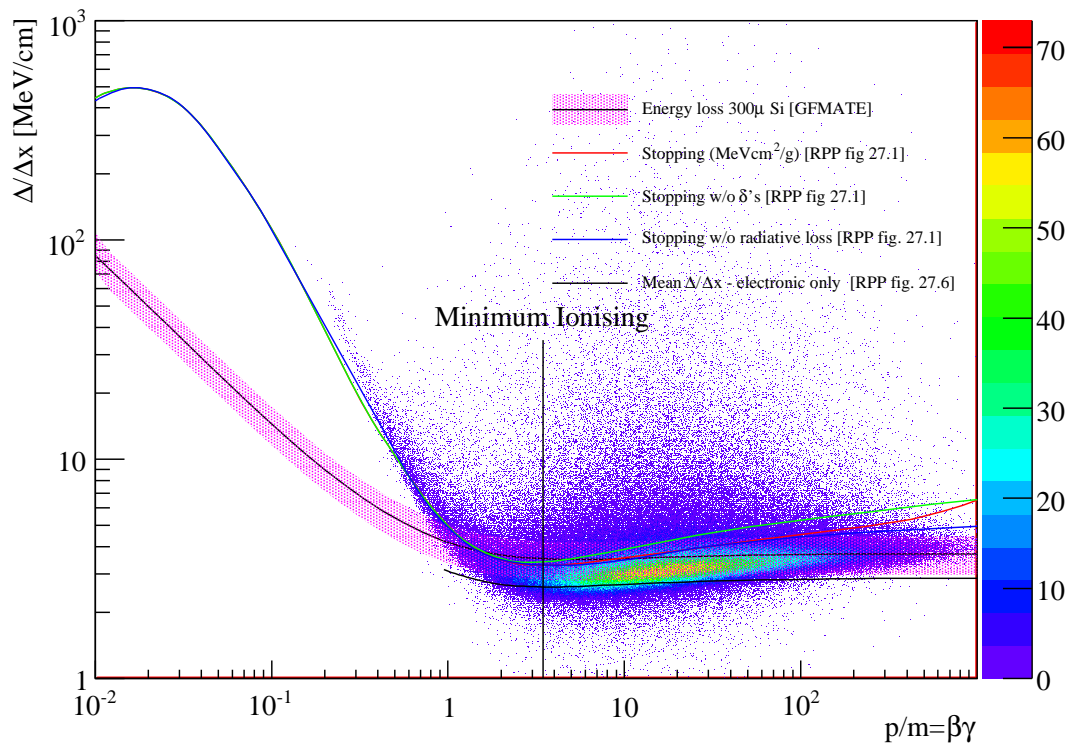


Figure E.1: Simulated energy loss in the FMD with no δ ray production.

Appendix F

Miscellaneous Photos




Figure F.1: An inner type sensor




Figure F.2: The FMD front-end electronics. From the top-left: An inner-type module of a sensor and hybrid card, inner type FMDD, outer type FMDD, and at the bottom, the RCU mother board with DCSC and SIU daughter cards.




Figure F.3: Test-beam setup at ASTRID. The e^- beam comes from ASTRID through the hole in the wall on the far left, and is stopped in the concrete blocks on the far right. In between are two scintillators used for triggering, and the devices-under-test: 6 silicon modules inside the grey box in the middle. Mounted on either side is an inner and an outer type FMDD.




Figure F.4: FMD2 (left) and 3 (right) installed in ALICE. Both pictures are taking in the direction of the MUON arm. On the left FMD2 is installed onto the ITS, while on the right, FMD3 is temporarily mounted on the MUON absorber nose. Later on, the support for FMD3 was moved from the absorber nose to the ITS.




Figure F.5: Some members of the FMD underneath the installed FMD3 sub-detector. Back row, from the left: Børge Svane Nielsen, Erik Kaimer Olsen, Jørn Westergaard, and Jimmy Cali Hansen. Front row, from the left: Gaël Renault, Kristján Herlache Gulbrandsen, Christian Holm Christensen, and Jens Jørgen Gaardhøje.

Glossary

Glossary of common terms and abbreviations. The numbers on the far right refers to the page where the term is used.

A+A	7
Collisions between nuclei.	
AGS	5
The Alternating Gradient Synchrotron facility at BNL.	
ALIRO	61
<i>ALice Tpc Readout</i> . The analogue-to-digital converter (ADC) developed for the ALICE TPC. Beside the TPC, the following detectors in ALICE also make use of the ALIRO: FMD, EMCAL, and PHOS. It is a 16-channel 10bit ADC, with parallel read-out and high radiation tolerance. The ALIRO is a rather advanced chip, and has many features and filters to improve the read-out.	
Asymptotic freedom	2
The freedom of quarks in the limit of high quark density, or short distances.	
b	7
(Impact parameter) The vector between the centres of two colliding nuclei. The length b is measured in fm.	
Baryon	3
Particles made of 3 (anti)quarks, such as the proton and neutron.	
Bethe–Bloch equation	46
The equation that describes energy loss (or stopping power) of charged particle traversing matter.	
BNL	5
Brookhaven National Laboratories on Long Island, New York, USA.	
CERN	5
European Organisation for Nuclear Research in Geneva, Switzerland. It used to be <i>Conseil Européen pour la Recherche Nucléaire</i> , hence the odd abbreviation.	
CGC	19
<i>Color Glass Condensate</i> . A description of the initial state of the colliding nuclei.	
$dM_{\text{ch}}/d\eta$	16
The charged particle density $dM_{\text{ch}}/d\eta$ is the number of charged particles (p , \bar{p} , π^\pm , K^\pm , and so on) per unit of pseudo-rapidity η .	

- Chemical Freeze-out** 11
The time after the collision, from which the chemical potential of various species does not any more.
- Chirality** 3
The handedness of particles. Particles with their spin *aligned* to their momentum vector are said to be *right-handed*; particles with their spin *anti-aligned* are said to be *left-handed*.
- Confinement** 2
The property of the QCD interaction, that confines the quarks inside the hadrons.
- DCSC** 82
Detector Control System Card. Daughter card of the RCU card that provides the interface to the detector control system. It has a small ARM processor running GNU/Linux on it.
- Dilepton** 13
A pair of a lepton and the corresponding anti-lepton, as for example μ^- and μ^+ .
- Disoriented Chiral Condensates** 13
Also DCC. The behaviour of the chiral condensate $\langle\psi\bar{\psi}\rangle$ when the chiral symmetry is explicitly broken, may lead to DCC.
- ENC** 53
Equivalent Noise Charge is a measure of the noise of a pre-amplifier. It is the root-mean-square of the distribution of charges that are equivalent to the electronic noise of the system.
- Energy loss straggling** 48
Since energy loss is a statistical process, the observed spectra of energy loss for a given $\beta\gamma$ is distributed in some way around some mean as described by the Bethe-Bloch equation.
- Event plane** 22
The plane spanned by the collision axis \mathbf{z} and the impact parameter \mathbf{b} . See also Figure 2.8.
- Gauge boson** 2
Force mediators, or interaction particles. ‘Gauge’ because they measure the force of the interaction.
- Hadron** 3
Particles made up of 2 or more quarks.
- Initial State** 19
The state of the colliding nuclei *before* the collision.
- Kinetic Freeze-out** 11
The time after the collision, from the particles essentially no longer transfer momentum between each other.
- Lattice QCD** 3
Due to the large coupling constant (strong interaction) of QCD in the limit of low energy and large distances, it is not possible to analytically calculate quantities in QCD. Therefore a method of numerical approximation where space-time is considered discrete (points in a lattice) is used; the method is called Lattice QCD.

- Lepton** 2
Fundamental particle in the standard model, that determines the chemical properties of atoms and molecules.
- LHC 5
The future *Large Hadron Collider* facility at CERN.
- Minimum Ionising Particles** 48
A minimum ionising particle (or MIP) is a particle with a $\beta\gamma$ that corresponds to the minimum of the Bethe–Bloch equation.
- N+N** 7
Collisions between nucleons.
- Nucleon** 3
The constituents of the atomic nucleus i.e., neutrons and protons.
- Occupancy** 31
The occupancy for a detector element, is the probability that at least one particle hit the element in an event. Some people use “Occupancy” in the meaning of “mean number of hits”, however, the above definition emphasis the probabilistic nature rather than the statistical.
- Participants** 8
Nucleons that interact during the collision.
- Quantum Chromodynamics** 2
The theory of the strong interactions within the standard model.
- Quantum Electrodynamics** 2
The theory of electromagnetic interactions within the standard model.
- Quantum Electroweakdynamics** 2
The unified theory of electromagnetic and weak interactions. The weak interaction is the one responsible for the radioactive decay of particles.
- Quark** 2
Fundamental particle in the standard model from which nucleons, and hence matter, are build
- Quark–Gluon Plasma** 5
The state of nuclear matter, predicted by QCD, where chiral symmetry is restored, and the temperature and/or density is so high, that the hadrons overlap, and the quarks become asymptotically free, or deconfined.
- RHIC 5
The *Relativistic Heavy Ion Collider* facility at BNL.
- Signature** 11
An observation that lends creditability to a given theory of Nature, but does not rule out other contradictory theories. This is in contrast to real evidence, where the observation can not be easily explained without the assumption of the theory.

- $\sqrt{s_{NN}}$ 132
 Centre-of-mass energy per colliding nucleon pair. $\sqrt{s} = x$ AGeV is the same as $\sqrt{s_{NN}} = x$ GeV.
- Spectators** 8
 Nucleons that does not interact during the collision.
- SPS 5
 The Super Proton Synchrotron at facility CERN.
- T 11
 The temperature is the inverse slope parameter of $\frac{d^3N}{dp^3} = \frac{Y}{4\pi T m^2 K_2(\frac{m}{T})} e^{-\frac{E}{T}}$
- Tracklets** 31
 Combinations of hits to form 3-vectors. Tracklets are different from tracks, in that no assumptions are made with regard to the true trajectory of the particle, and tracklets are usually fairly short. Tracklets are often used as seeds for more complex tracking algorithms.
- VA1₃** 50
 A pre-amplifier of the VIKING family. It amplifies and shapes the low signal from the silicon sensor, and multiplexes the signal from 128 strip into one channel. VA stands for VIKING ASIC, and ASIC in turn stands for Application Specific Integrated Circuit.
- Verilog** 65
 A moderate level *Hardware Description Language* (HDL) for describing hardware, implementing firmware, etc. in human-readable code.
- VHDL** 65
VHSIC Hardware Description Language is a higher level programming language for implementing firmware, describing hardware, etc. in human-readable code. VHSIC stands for *Very-High-Speed Integrated Circuit*.
- x 19
 The fraction of the total nucleon momentum carried by a parton — i.e., a quark, anti-quark, or gluon.

Index

- A+A, [7](#)
- AGS, [5](#)
- ALICE, [7](#), [30](#)
- ALTR0, [61](#)
 - Data formatting, [64](#)
 - Pedestal subtraction, [62](#)
 - Zero suppression, [63](#)
- ALTR0 bus
 - `cstb`, [70](#)
- ALTR0 bus
 - `ackn`, [70](#)
- ALTR0 bus, [59](#), [70](#)
 - `ackn`, [60](#)
 - `cstb`, [60](#)
 - `dstb`, [60](#)
 - `error`, [60](#)
 - `trsf`, [60](#)
 - `write`, [60](#)
 - bd, [60](#)
 - Protocol of, [70](#)
- Asymptotic freedom, [2](#)
- ATLAS, [7](#)
- Azimuthal anisotropy, *see* v_2 , *see* v_2
- b**, [7](#), [15](#)
- Barrel detectors, [31](#)
- Baryon, [3](#)
- BC
 - VA1₃ interface, [66](#)
 - Commands, [67](#)
 - Implementation, [70](#)
 - Monitor ADC interface, [69](#)
 - Performance, [70](#)
 - Registers, [66](#), [67](#)
 - Testing, [70](#)
 - Trigger handling, [66](#)
- BC, x, [xiii](#), [58–62](#), [65–67](#), [69–71](#), [74–76](#), [88](#), [105](#), [151](#)
- β , [131](#)
- Bethe–Bloch equation
 - δ -electron, [48](#)
- Bethe–Bloch equation, [46](#)
 - Density correction to, [47](#)
 - Shell correction to, [47](#)
- Bethe–Bloch equation
 - Radiative loss, [48](#)
- BNL, [5](#)
- BRAHMS, [7](#)
- c, [15](#)
- c, [8](#)
- CAEN, [81](#), [151](#)
- Centrality, *see* c
- CERN, [5](#)
- CGC, [19](#)
- Charge carries, [43](#)
- $dM_{\text{ch}}/d\eta$, [16](#)
- Chemical Freeze-out, [11](#)
- Chirality, [3](#)
- clk, [61](#)
- CMS, [7](#)
- Collider, [7](#)
- Collision
 - central, [7](#), [15](#)
 - peripheral, [7](#), [15](#)
 - Stopping in a, [8](#)
 - Transparency of, [8](#)
- Confinement, [2](#)
- CTP, [61](#), [151](#)
- DA, [77](#)
 - FEE parameters, [79](#)
 - Gains, [79](#)
 - Pedestals, [77](#)
- DCC, *see* Disoriented Chiral Condensates
- DCS, x, [74](#), [80–82](#), [84–86](#), [88](#), [89](#), [151](#)
- DCSC, [82](#)
 - Embedded GNU/Linux box, [82](#)
- DCSC, [72–76](#), [82–85](#), [94](#), [144](#), [148](#), [151](#)
- DDL, [77](#), [79](#), [80](#), [151](#)
- δ -electron, [48](#)
- Detector algorithm, *see* DA
- Digitiser card, *see* FMDD
- Dilepton, [13](#), [34](#)
- DIM, [82](#), [84](#), [151](#)
- Disoriented Chiral Condensates, [13](#)
- DIU, [77](#), [151](#)
- Donor, [43](#)

- Electron–hole pairs, 43
- EMCAL, 33
- EMCAL, 30–33, 61, 84, 133, 147, 152
- ENC, 53
- Energy loss straggling, 48
 - thick absorbers, in, 48
- η , 131
- Event plane, 22
- Event plane angle, *see* Ψ

- FedServer, 84, 85, 152
- FeeServer, 82
- FeeServer, 82–90, 152
- FMD, 34
- FMD, iii, iv, ix–xi, xiii, 1, 2, 13, 15, 16, 23, 24, 26, 30–34, 39–41, 46, 50, 52, 54–56, 58, 60–63, 65–67, 69, 70, 72, 74–77, 79–88, 90, 91, 93, 98, 101–104, 106, 107, 109–112, 114–117, 119–122, 133, 141, 144, 145, 147, 152
- FMDD, 56
- FMDD, x, xiii, 41, 56, 58–62, 64–67, 69–76, 79, 82, 83, 86–88, 94, 106, 107, 121, 144, 151, 152
- Forward detectors, 33
- Freeze–out
 - Chemical, 11
 - Kinetic, 11
- FWD, iii, 152

- γ , 131
- Gauge boson, 2
- GDC, 77, 79, 152

- Hadron, 3
- HMPID, 33
- HMPID, 30–33, 133, 152
- Hybrid card, 50

- ItwoC bus, 59
- ICL, 84
- ICL, xiii, 82, 84–86, 88, 90, 152
- IDEAS, 91, 152
- Impact Parameter, *see* **b**
- Initial State, 19
- InterCom Layer, *see* ICL
- ITS, 31
- ITS, ix, 11, 13, 16, 30–34, 39, 103, 108, 115, 119, 133, 145, 152

- Jet quenching, 12

- Kinetic Freeze–out, 11
- L3, 30, 31, 34, 152
- Lattice QCD, 3
- LDC, 77, 79, 94, 152
- Leakage current, 44
- Lepton, 2
- LHCb, 7
- LHC, 5
- Longitudinal momentum, *see* p_{\parallel}
- Lorentz factor, *see* γ

- Minimum Ionising Particles, 48
- Module, *see* Hybrid card
- m_{\perp} , 131
- MUON, 34
- MUON, 31, 34, 35, 145, 152
- Muon spectrometer, *see* MUON

- N+N, 7
- Nucleon, 3

- Occupancy, 31

- Participants, 8
- Particle ratios, 11
- PDC, 77, 152
- Pedestal, 62
- PHENIX, 7
- PHOBOS, 7
- PHOS, 33
- PHOS, 30–32, 61, 84, 133, 147, 152
- PMD, 30, 152
- Pseudorapidity, *see* η
- Ψ , 116
- $\langle\psi\bar{\psi}\rangle$, 13
 - In Lattice QCD, 3
- p_{\perp} , 131
- p_{\parallel} , 131

- Quantum Chromodynamics, 2
- Quantum Electrodynamics, 2
- Quantum Electroweakdynamics, 2
- Quark, 2
- Quark–Gluon Plasma, 5

- Rapidity, *see* y
- RCU
 - I²C bus, 74
 - status**, 74
 - Architecture, 73
 - Branches, 73
 - Busy box, 73
 - Data transfer, 73
 - Instruction interpreter, 74

- Monitoring, 74
- Slow control, 74
- Trigger handling, 73
- RCU, x, 32, 41, 59–62, 64, 69–76, 79, 80, 82–87, 90, 94, 144, 148, 152, 153
- RcuCE, 82–84, 153
- Read-out clock, *see* `clk`
- RHIC, 5
- RORC, 77, 94, 106, 153
- `sample_clk`, 61
- Sample clock, *see* `sample_clk`
- SDD, 31, 153
- Semi-conductor, 42
- Signature, 11
- Silicon, 42
- Silicon crystal
 - n-type, 43
 - p-type, 43
- SIU, 72–74, 77, 85, 94, 144, 153
- $\sqrt{s_{NN}}$, 132
- SOD, 77
- SOD event, 73
- SPD, 31, 153
- Spectators, 8
- Sps, 5
- SSD, 31, 153
- Standard model, 2
- STAR, 7
- Star of data event, *see* SOD event
- T0, 34
- T0, iii, 30, 33, 34, 39, 108, 153
- T , 11
- TOF, 32
- TOF, 11, 30–32, 35, 153
- TPC, 32
- TPC, iii, iv, 11, 13, 30–33, 58, 61, 63, 65–67, 70, 72, 73, 82, 84, 101, 119, 121, 133, 147, 153
- Tracking, 33
- Tracklets, 31
- Transverse mass, *see* m_{\perp}
- Transverse momentum, *see* p_{\perp}
- TRD, 32
- TRD, 11, 30–33, 84, 153
- TTC, 73, 75, 153
- V0, 34
- V0, iii, 30, 33, 34, 39, 103, 153
- v_2 , 13, 119
- VA13, 50, 50
- Calibration pulser, 55
- Noise, 53
- Read-out, 55
- Verilog, 65
- VHDL, 65
- x , 19
- y , 131