

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

Development of a High-Level Trigger for the Dimuon Spectrometer of the ALICE Experiment at the Large Hadron Collider

Bruce Becker

Thesis Presented for the Degree of
DOCTOR of PHILOSOPHY
to the Department of Physics, Faculty of Science
University of Cape Town

Cape Town, September 2006

UT 530 BECK
814136

To those who made it happen.

University of Cape Town

Abstract

The ALICE experiment at CERN's Large Hadronic Collider will mark the beginning of a new phase in the study of ultra-relativistic heavy ion collisions. It will be possible to explore in great detail phenomena discovered or hinted at in the course of experiments at the Relativistic Heavy Ion Collider, in particular the signals of the quark-gluon plasma. One of the most promising signals of the creation of this new state of matter is the anomalous suppression of the Υ ($b\bar{b}$) and J/ψ ($c\bar{c}$) families. One of the main decay channels of these mesons is into dimuons and ALICE has a dedicated dimuon spectrometer in order to study the spectra of these interesting particles. The signal is, however, swamped by a large background from several other muonic sources. Due to the large data rate expected for ALICE and the limited bandwidth, a highly efficient and selective trigger is required for the experiment - the dimuon high-level trigger (dHLT).

This thesis concerns the context, development and implementation of the ALICE dimuon high-level trigger. The physics context of the experiment is described, as well as the technical requirements of the system. The performance of the prototype is investigated with the use of Monte-Carlo simulations. An investigation into the performance estimates of system in the expected physics environment was performed, which validated the benchmarks, as well as investigation of the effect of possible modifications of quarkonia yields due to QGP formation on the response of the dHLT.

Acknowledgements

This work was started at a time when the UCT-CERN Research Centre did not yet exist as a formal entity, and collaboration with ALICE was only barely starting. Joining an existing collaboration of hundreds of members based in a far-away lab is a daunting and exciting task, which would have been impossible to accomplish were it not for the great support we have received from them. During the course of this work, I have had the opportunity to work with many physicists and engineers from around the world. Innumerable contributions were made from too many people to mention, but I especially grateful to the following:

- To my supervisors, Dr. Zebon Vilakazi and Prof. Jean Cleymans, for giving me the freedom to explore various avenues, even stray very far from the path set out for me in order to find what I am good at. For allowing the group to grow in its own way.
- Dr. Florent Staley, for having foresight and patience and generally making life a little easier.
- To the members of my second family, the UCT-CERN Research Centre : Rory, Spencer, Sarah, Nawahl, Mike, Maciej, Prof. Fearick, Heather. Especially to:
 - Mark, for being there in the good times and the bad, and most importantly, behind me on the last rep so that I didn't crush myself. For giving me some perspective on work and rugby and being around to soak up all the negative attitude I was able to vent from time to time.
 - Gareth, for being able to take over from where I left off and building a machine better than I was able to.
 - Artur, for having the clear head to know what needed to be done and doing it without fuss.
- To the other members of the HLT and dHLT collaboration within ALICE : Dr. Volker Lindenstruth, Dr. Corrado Cicalo, Dr. Dieter Rörich, as well as Dr. Timm Steinbeck, Matthias Richter, Arne Wiebalck, Lord Hess and co, for their patient help and organisation, with the PubSub framework and all the other technical details of the implementation of the dHLT.
- Dr. Peter Steinberg and the PHOBOS team, for all of the great chats we've had, for giving some impression of what it's like to be a professional and for the 'PHOBOS experience'.

- Dr. Bernard Li, Dr. Stephen L. Scott and the rest of the OSCAR team, for their great support and lightning speed turnaround on issues we had when building CARMEN-I and CARMEN-II
- Dr. Philippe Crochet for being on hand to discuss at any time the physics with me and pointing out interesting alleys down which to explore new ideas.

During the time taken to finish this thesis, I have been very warmly received by several colleagues, old and new. I would like to acknowledge the support I've had from

- Schonland Research Centre for the Nuclear Sciences, from Dr. Simon Connell
- INFN, Cagliari, from Dr. Corrado Cicalò
- INFN, Milano, from Dr. Angela Bracco

Finally, to the two women in my life: Tiziana, for waiting and for believing. And my old mom, Wren, for supporting me no matter what. I promise you'll get to read my thesis, this time...

University of Cape Town

Contents

1	Introduction	2
1.1	Heavy ion experiments in context	2
1.2	The search for the QGP	3
1.2.1	Early experimental searches for the QGP	3
1.2.2	The Relativistic Heavy Ion Collider	4
1.3	The Large Hadron Collider	5
1.4	Thesis outline	5
2	Theoretical Overview and Background	7
2.1	The MIT Bag Model	7
2.2	QCD and lattice QCD	8
2.3	Ultra-relativistic heavy ion collisions	10
2.3.1	Kinematic variables	12
2.3.2	Collision dynamics	12
2.3.3	The evolution of the QGP	13
2.3.4	Experimental heavy ion collision programs	13
2.4	Observing the QGP	15
2.4.1	Global observables	15
2.4.2	Strangeness production	20
2.4.3	Elliptic flow observables	23
2.4.4	Specific signatures	24
2.4.5	Heavy quarkonia suppression	33
2.4.6	Non-plasma suppression mechanisms	37
2.4.7	J/ψ suppression from SPS and RHIC	40
2.5	The challenge for ALICE	41
3	The ALICE experiment at CERN's Large Hadronic Collider (LHC)	43
3.1	The Large Hadron Collider	43
3.2	The ALICE experiment overview	45

3.3	The central (barrel) ALICE systems	45
3.3.1	The L3 magnet	46
3.3.2	The Inner Tracking Chamber (ITS)	46
3.3.3	The Time Projection Chamber (TPC)	47
3.3.4	The Transition Radiation Detector (TRD)	47
3.3.5	The Time of Flight (TOF) detector	48
3.3.6	The High Momentum Particle IDentification (HMPID) detector	48
3.3.7	The PHOton Spectrometer (PHOS)	49
3.4	The forward-rapidity ALICE detectors	49
3.4.1	The Photon Multiplicity Detector (PMD)	50
3.4.2	The Forward Multiplicity Detector	51
3.4.3	The T0 and V0 detectors	51
3.5	The dimuon spectrometer	52
3.5.1	The absorbers	54
3.5.2	The spectrometer tracking system	54
3.5.3	The dimuon trigger subsystem	55
3.5.4	Dimuon spectroeter read-out	56
3.5.5	The dimuon spectrometer raw data format	59
3.5.6	Hit reconstruction (offline cluster finder)	60
3.6	The ALICE trigger and DAQ system	61
3.6.1	General trigger design considerations	61
3.6.2	The Central Trigger Processor	62
3.6.3	ALICE data acquisition	63
3.7	The ALICE offline computing framework	64
3.7.1	AliRoot functionality - simulated and real data	65
3.8	Summary	66
4	The ALICE High Level Trigger	67
4.1	Introduction	68
4.1.1	Physics observables and specific triggers	69
4.1.2	HLT design requirements	70
4.2	The online HLT	72
4.2.1	HLT data flow	72
4.2.2	Software integration into AliRoot	73
4.3	HLT computer network architecture	74
4.3.1	Connectivity and data flow	75
4.3.2	The HLT compute nodes	76
4.4	HLT communication framework	76

4.4.1	Publisher-Subscriber paradigm	76
4.4.2	The HLT analysis object	77
4.4.3	Load balancing	78
4.4.4	Publisher-Subscriber bridges	78
4.4.5	Publisher-Subscriber performance and fault tolerance	78
5	The design of the dimuon High Level Trigger (dHLT)	81
5.1	Introduction	81
5.2	Data flow of the dHLT	82
5.3	Raw data generation and unpacking	84
5.3.1	Raw data generation from ROOT trees	84
5.4	The dHLT hit reconstruction	85
5.4.1	Datastream model and choice of reconstruction algorithm	85
5.4.2	PC-CPU optimised hit reconstruction algorithm.	87
5.4.3	FPGA optimised hit reconstruction algorithm.	88
5.4.4	Hit reconstruction resolution performance	89
5.4.5	Hit reconstruction timing performances	90
5.5	The dHLT tracking algorithm	91
5.5.1	The “Manso” tracking algorithm	92
5.6	Benchmarks of the Manso tracking algorithm	95
5.6.1	Cut width and momentum reconstruction accuracy	96
5.6.2	Background rejection factor	98
5.6.3	Acceptance	98
5.7	Summary	98
6	Implementation of the dHLT	99
6.1	The online and offline cases	99
6.2	Interface to the HLT data transport framework	100
6.3	Offline implementation of the ALICE dHLT software	100
6.3.1	Raw data and reconstructed hits	100
6.3.2	Tracking component and Decision component	101
6.4	Online implementation	102
6.4.1	The “callback” option	102
6.4.2	The “monolithic” option	102
6.5	The dimuon High Level Trigger integration test	103
6.5.1	Test-bench technical specifications	103
6.5.2	Integration and stress test goals	103
6.5.3	Dummy load stress test	104

6.5.4	Monolithic dHLT tracker integration	108
6.6	Global HLT test over public Internet network	108
6.6.1	Global test setup	108
6.6.2	HLT <i>TaskManager</i> setup	111
6.6.3	Global test results	112
6.7	Next steps	112
7	Physics performance of the dHLT	114
7.1	Introduction	114
7.1.1	Selection of data sets	115
7.2	Baseline data production (no QGP)	116
7.2.1	Rapidity and transverse momentum spectra	118
7.2.2	Simulation and analysis frameworks	119
7.3	Background sources discussion	120
7.4	Detailed simulation of the background muon sources	122
7.4.1	Result of dHLT processing	122
7.4.2	Result of dHLT processing on signal	124
7.4.3	Mixed events	127
7.5	The dHLT with a quarkonia suppression model	129
7.5.1	Model description of the nature of the produced QGP	130
7.5.2	Quarkonia survival probability	131
7.6	Suppressed quarkonia spectra	134
7.6.1	Yields including feed-down decays	135
7.6.2	Modifications to the ALICE simulation framework	137
7.7	Analysis of Monte-Carlo suppression data	137
7.7.1	Results of dHLT processing	138
7.8	Discussion	138
7.8.1	Interpretation of dHLT results	138
7.8.2	Description and choice of QGP model	140
8	Summary and Outlook	141
8.1	Summary	141
8.2	Outlook	142
8.2.1	The future of the dHLT	142
8.2.2	Evolution of quarkonia models and implications for the dHLT	143
8.2.3	Conclusion	143

A The UCT-CERN Computing Facilities.	145
A.1 Introduction	145
A.2 Motivation and goals	145
A.3 Design considerations	147
A.4 Implementation	147
A.4.1 Physical setup	148
A.4.2 Data storage capacity	149
A.4.3 Operating system and resource management	150
A.4.4 Monitoring	151
A.5 Evolution	152
A.5.1 Node software and hardware upgrade	152
A.5.2 Network upgrade	152
A.5.3 Separation of roles	153
A.6 San Diego - 64 bit extension to UCT-CERN computing facilities.	155
A.6.1 Challenges set by the 64 bit platform	155
A.6.2 San Diego implementation	157
A.7 Outlook	158
Glossary	161
List of Acronyms	164

List of Figures

2.1	The energy density and pressure of different systems of quarks calculated on the lattice with QCD [29]. Note the large increase in the energy density around the critical temperature $T = T_c$	9
2.2	A phase diagram of QCD matter, showing the dependence of the phase transition on the mass of the constituent quarks.[29].	10
2.3	Schematic phase diagram showing the boundary between confinement and deconfinement and exotic states at high μ_B . This phase diagram indicates that a tri-critical point may exist around $\mu_B = 700$ MeV and indicates the regions (blue points) explored by current and future heavy ion experiments	11
2.4	Diagram of the space-time evolution of a heavy ion collision, in the presence (right) and absence (left) of a QGP [21].	14
2.5	An example of the centrality determination in the PHOBOS experiment at RHIC [32]. The correlation between the zero-degree calorimeters (at forward rapidities) and paddle counters (nearer mid-rapidity) is shown, with the corresponding calculation of the number of participants from a Glauber model	17
2.6	Transverse energy density as a function of the number of participating pairs in Au-Au collisions (left) and corresponding energy density calculated in the Bjorken picture, as measured by PHENIX [61].	19
2.7	Model predictions for charged particle multiplicities per participant in A-A and p-p collisions as a function of \sqrt{s} . The dashed line is the fitted function $0.68 \ln \frac{\sqrt{s}}{0.68}$, on the nuclear data, while the dotted line is a fit to $0.7 + 0.028 \ln^2 \sqrt{s}$. The long-dashed line is a prediction based on a saturation model [74]. The estimate for ALICE in these models is given by the intersection of the curves with the grey arrow on the right.	20
2.8	Rapidity density of net protons measured at the AGS, SPS and RHIC energies, for central collisions. The beam rapidity of each collider is shown. Note the rapid fall-off of net proton density with beam energy.	21

- 2.9 Comparison between the particle ratios from Au-Au collisions measured by the four RHIC experiments, and statistical model calculations [102], with $T = 174$ MeV and $\mu_B = 46$ MeV. 21
- 2.10 The yield of Λ , Ξ and Ω (containing one, two and three strange quarks respectively) in Pb-Pb collisions relative to the yield in p-Be collisions, as measured by the NA57 experiment [33], as a function of the number of wounded nucleons $\langle N_{\text{wound}} \rangle$. The yields are relative to the number of participating nucleons. Shown in the left panel are the yields for particles which may have had at least one quark in common with the projectile or target; the right panel shows the yield for particles whose constituent quarks would have had to be created during the collision. 23
- 2.11 Ratio of $\frac{K^+}{\pi^+}$ yield at mid-rapidity as measured by the AGS RHIC and SPS (see references 23 and 24 in [39]). 24
- 2.12 Measured ratio of $\frac{K^+}{\pi^+}$ yields as a function of $\sqrt{s_{NN}}$ (left) and strangeness-to-pion ratio $E_s = \frac{\langle \Lambda \rangle + \langle K + \bar{K} \rangle}{\langle \pi \rangle}$ as a function of $F = \frac{(\sqrt{s_{NN}} - 2m_N)^{\frac{3}{4}}}{\sqrt{s_{NN}}^{\frac{1}{4}}}$ (right). 25
- 2.13 Centrality dependence of v_2 , integrated over p_T . The triangles are the NA49 measurements for pions at $\sqrt{s_{NN}} = 17$ GeV [44], while the circles and crosses are measured by STAR at $\sqrt{s} = 130$ GeV [45] and 200 GeV [46], respectively (figure from [47]). 26
- 2.14 Hydrodynamics calculations for a hadron gas (EOS H) and QGP (EOS Q) equation of state [48], compared to STAR v_2 measurements for pions and protons in minimum bias 130 GeV Au-Au collisions [49] (figure from [47]). 26
- 2.15 Minimum-bias v_2 for identified particles as a function of p_T [68] and hydrodynamic calculations [69]. 27
- 2.16 Identified particle v_2 , scaled by the number of constituent quarks, fitted to a polynomial. The ratio of data/fit is shown in the lower panel [68]. 28
- 2.17 Direct γ excess double ratio (see text) and cross-section in 200 GeV p-p collisions at RHIC [50]. The photon cross-section is compared with pQCD calculations for various scale factors from [51] 28
- 2.18 Direct γ excess double ratio (see text) in 200 GeV Au-Au collisions at RHIC [50] as a function of p_T for various centralities. The solid line represents a Next-to-leading order (NLO) Perturbative QCD (pQCD) calculation [51] 29

2.19	Nuclear modification factor R_{AA} for charged hadrons and neutral pions in the most central 10% Au-Au collisions at RHIC (solid points). Error bars indicate statistical errors while surrounding bands indicate quadrature sum of systematic errors. Also shown are the modification factors for $\alpha + \alpha$ at $\sqrt{s} = 31$ GeV [54] and central Pb-Pb collisions at the SPS [55], indicated by the dotted uncertainty band.	30
2.20	Evidence of a deconfined medium created in Au-Au collisions at $\sqrt{s} = 200$ GeV, from the RHIC experiments PHENIX [58] and Solenoidal Tracker At RHIC (STAR) [59].	31
2.21	Correlation function of pairs of charged hadrons as measured by the PHENIX Collaboration in Au-Au collisions at $\sqrt{s} = 200$ GeV for different centralities. The solid bands indicate the estimated background contribution.	33
2.22	Correlation function with the underlying event subtracted for central (red), mid-central (black) and peripheral events (blue), showing the splitting of the away-side peak in central Au-Au collisions at $\sqrt{s} = 200$ GeV.	34
2.23	Feynmann diagrams of leading-order (LO) processes for $c\bar{c}$ ($b\bar{b}$) production.	36
2.24	The J/ψ (left) and Υ (right) suppression patterns as a function of temperature relative to the critical temperature of the plasma from [111].	37
2.25	J/ψ cross-sections per nucleon, times dimuon branching ratio, as a function of nuclear size taken from [85]. The line corresponds to the best fit to the function $B_{\mu\mu} = \sigma_0(A \times B)^{\alpha-1}$	38
2.26	Ratio of J/ψ to Drell-Yan production as a function of E_T measured by NA50. The dashed line is the normal nuclear absorption model [112]. The solid line is a co-mover absorption calculation [92]	39
2.27	Calculation of the nuclear modification of the gluon density due to the effect of “shadowing”, taken from [95]	40
2.28	Comparison of measured suppression of J/ψ to that expected from normal nuclear suppression, showing a threshold of “anomalous” suppression near $2.2 \text{ GeV}/\text{fm}^3$	41
2.29	R_{AA} of J/ψ mesons measured in Au-Au collisions at $\sqrt{s} = 200$ GeV as a function of centrality (N_{Npart}) [116]. The top pane shows R_{AA} at mid- and forward rapidity, while the bottom frame shows the ratio between R_{AA} measured at forward rapidity, to that measured at mid-rapidity.	42
3.1	The LHC accelerator complex at CERN, showing the location of the four experimental halls, including the ALICE experiment (Point 2).	44
3.2	Schematic diagram of the ALICE detector, to scale, showing various sub-systems.	45

3.3	The Inner Tracking System, showing the various component layers of silicon	47
3.4	Schematic scale diagram of the ALICE Time Projection Chamber, showing support spokes, and inner and outer sections.	48
3.5	Schematic diagram of the ALICE Transition Radiation Detector (to scale) .	49
3.6	Schematic diagram of the TOF (central barrel), PHOS (lower) and HMPID (top) detectors in the ALICE experiment, to scale.	50
3.7	The ALICE Zero degree calorimeter (ZDC), next to the beam pipe.	51
3.8	A schematic diagram of the ALICE Photon Multiplicity Detector (PMD) (to scale)	52
3.9	Schematic diagram of the ALICE dimuon arm front absorber, to scale, showing the distribution of materials used.	54
3.10	Schematic view of the ALICE dimuon spectrometer dipole magnet. The front (left) and side (right) views, as generated from aliroot, both show the third tracking station of the spectrometer within the volume of the dipole magnet.	55
3.11	Occupancy of the tracking chambers as a function of radial distance in the case of central Pb-Pb collisions [128]	56
3.12	MUON spectrometer tracking and trigger chambers. Dimensions are in cm.	57
3.13	Raw data format used in offline simulations.	59
3.14	Schematic diagram of the ALICE CTP, showing the flow of various data types and the subsystems with which the CTP interacts [143]	62
3.15	Data Acquisition (DAQ) data flow overview	64
3.16	Schematic overview of AliRoot functionality and interdependencies.	65
3.17	The ALICE data processing framework, showing the connection between simulated data and real data, with the relationship between the amount of information available and the amount of processing required.	66
4.1	Schematic diagram of the functionality of the HLT	73
4.2	A general overview of the HLT data flow, showing the logical functions performed at each stage [156]	74
4.3	Functionality of the HLT analysis object. Data is subscribed to, processed by the analysis object and then published back into the event stream . . .	77
4.4	Schematic diagram of the network bridge of the HLT communication framework, showing the abstraction layer.	79

4.5	Test of the HLT failure tolerance with one node failing and one spare [150]. The events marked by “1” and “2” are the network interruption on the original node and the re-distribution of the task to the spare node respectively. Note that for clarity, the curves have been scaled vertically in order see them better.	80
5.1	Schematic diagram of the data flow of the dHLT. The thick connections in black show the normal flow of data processed by the dHLT. The dashed blue connections show the input of calibration and control information from DCS and the conditions database. The thin connections in red show the extraction of monitoring data from the dHLT processing chain for quality assurance purposes.	83
5.2	The average number of digits per reconstructed hit for the different spectrometer stations (left) for muons of 20 GeV. The distribution of digits per reconstructed hit for muons of energy 5 GeV - 10 GeV and 20 GeV is shown for comparison on the right, for chamber 8.	86
5.3	Schematic diagram of the PC CPU-optimised online hit reconstruction algorithm.	88
5.4	Measured hit reconstruction residuals ($x_{\text{hitrec}} - x_{\text{GEANT}}$) in the bending (left) and non-bending (right) planes, fitted to Gaussian distributions. . .	90
5.5	Timing and efficiency of the online hit reconstruction algorithm on tracking chambers 4 and 5, as a function of DC cut, for a nominal number of 150 muons in the spectrometer.	91
5.6	Schematic diagram of the “Manso” tracking algorithm (not to scale). . . .	93
5.7	The last tracking chamber of the dimuon spectrometer (chamber 10) is shown with a sample (central) event. The digitised hits are shown, along with an example of a possible search radius from the L0 information. The search circle has its centre at R_P and a radius of R_S given as a function of R_P (see text).	94
5.8	ALICE dimuon tracking algorithm p_T calculation principle, adapted from the L0 trigger principle [128]	95
5.9	Fits of the difference between the p_T calculated by the L0 (left) and dHLT (right) algorithms and the input p_T of muons. Note that the width of the dHLT distribution is significantly smaller than that of the L0 distribution.	96
5.10	Fits of the L0 (left) and dHLT (right) efficiency functions. Note that the fitted skin width (γ) of the dHLT function is smaller than that of L0. . . .	97

6.1	Schematic diagram of the reconstructed hit (left) and tracking (right) processing components.	101
6.2	Schematic diagram of the physical data flow of the dummy stress test. Connections and framework processes are shown, but not individual dHLT framework components (see figure 6.3)	104
6.3	The HLT framework components implemented in the dummy stress test, along with the data flow.	105
6.4	File publisher node resource utilisation during the dHLT stress test.	106
6.5	Tracking node resource utilisation data during the dHLT stress test.	106
6.6	Resource utilisation for a <i>DecisionComponent</i> node during the dHLT stress test.	107
6.7	Global HLT setup with all involved sites and nodes. The components on each node are not shown.	109
6.8	Setup of the HLT processing modules at the Bergen cluster. Black ovals denote physical nodes, while grey boxes denote HLT processing components. The network bridges to the relay/forwarding node are also shown.	110
6.9	Cluster-finding subscription and tracking processes running on the Heidelberg cluster. Black ovals denote physical nodes, with grey boxes representing HLT processes.	110
6.10	The hierarchy of the global test <i>TaskManagers</i> at the various sites. Each L3 <i>TaskManager</i> slave is controlled by the closest L2 <i>TaskManager</i> servants. In the case of Dubna, which there was not a need for an intermediate relay/access node, the L3 slave <i>TaskManager</i> is controlled by the L1 Master in Heidelberg.	111
6.11	Network input rates to the CARMEN head node during the global framework test, for the full duration of the run, and during one hour.	112
7.1	Transverse momentum and rapidity parametrisations of for the J/ψ (blue) and Υ (red) mesons, for Pb-Pb collisions at 5.5 TeV. (Curves scaled to match CDF data and are relative to the J/ψ production.)	118
7.2	Transverse momentum spectra of the muons from charm (blue), beauty (pink) and light hadronic (π , K mesons - red) decays in Pb-Pb collisions at 5.5 TeV. The sum of all contributions is also shown in black.	121
7.3	Transverse momentum (left) and invariant mass spectra of background muon sources.	122

7.4	Left: Transverse momentum spectra for the pure background data sample, as well as the corresponding dHLT and L0 output data samples (charged tracks only). Right Reconstructed transverse momentum resolution, for the muon cocktail background sample (see text).	123
7.5	Raw input (black), L0 (red) and dHLT (blue) p_T distributions and corresponding efficiency functions for the background data sample with a 1 GeV cut.	123
7.6	Fitted efficiency functions for a 1 GeV p_T cut on the background event sample. Note that the fitted width of the dHLT function is roughly half that of the L0 width.	124
7.7	Summary plot for the J/ψ data set.	125
7.8	The L0 and dHLT efficiency spectra as a function of p_T , for a 1 GeV p_T cut, fitted to an error function for the J/ψ data set.	125
7.9	Left: The invariant mass spectrum of the Υ family data sample generated. Right: The transverse momentum spectrum of the Υ data sample.	127
7.10	Left: The L0 efficiency function as a function of p_T with a 2 GeV cut for the Upsilon data sample. Right: The dHLT efficiency function for the same cut on the Upsilon data sample. Both have been fitted to the error function (see text). Note that the fitted width of the L0 histogram is larger than that of the dHLT width.	127
7.11	Invariant mass (left) and transverse momentum spectra (right) of the full muon cocktail data sample.	128
7.12	Left: Efficiency as a function of transverse momentum of the 1 GeV L0 cut for the mixed data sample. Right: Efficiency as a function of transverse momentum of the 1 GeV dHLT cut for the mixed data sample. Efficiency functions have been fitted to the error function described in chapter 5. . .	129
7.13	Survival probability of $b\bar{b}$ quarkonia family for a SU(3) plasma radius of 1 fm (7.13(a)) and $R = R_{Pb}$ fm (7.13(b)). The p_T dependence of the suppression of Υ (solid blue), Υ' (dashed red) and χ_b (dot-dashed magenta) is shown. Figure adapted from [171]	133
7.14	Survival probability of $b\bar{b}$ quarkonia family for a 3 flavours plasma radius of different plasma radii. The p_T dependence of the suppression of Υ (solid blue), Υ' (dashed red) and χ_b (dot-dashed magenta) is shown. Note the extended scale on the right. (Figure adapted from [171]).	134

7.15	Convolution of the SU(N)-predicted QGP suppression scenario compared to the standard production distribution (black) of Υ (solid blue), Υ' (dashed-red) and χ_b (dot-dashed magenta), for two extreme cases of the plasma radius – 1 fm and 7.1 fm.	135
7.16	Convolution of the SU(3)-predicted QGP suppression scenario, of Υ (solid blue), Υ' (dashed red) and χ_b (dot-dashed magenta), for two extreme cases of the plasma radius – 1 fm and 7.1 fm (log scale shown for the $R = 7.1$ fm case in order to make the differences clearer).	135
7.17	Convolution of indirect production spectra of Υ (solid blue) and Υ' (dashed red) and the SU(N)-predicted QGP suppression distribution, with the standard production distribution (black), for two extreme cases – plasma radius of 1 fm and of 7.1 fm.	136
7.18	Convolution indirect production spectra of Υ (solid blue) and Υ' (dashed red) and the SU(3)-predicted QGP suppression distribution, with the standard production distribution (black).	137
7.19	Transverse momentum distributions for the SUN and SU3 QGP cases, for an event sample of pure signal - $1.85 \cdot 10^5$ and $4.09 \cdot 10^5$ events respectively.	138
7.20	L0 and dHLT efficiency spectra for the SU(N) Υ data sample.	139
A.1	A view of a CARMEN compute node, showing the PC components. Each node was equipped with 80 GB hard drive and 512 MB of RAM	148
A.2	Schematic diagram of the first implementation of the UCT-CERN computing facility CARMEN, showing network connections amongst nodes and data storage.	149
A.3	Photo of the current CARMEN cluster at UCT-CERN, showing 3 levels of the 4-level rack, compute nodes, RAID array and head node	150
A.4	Schematic diagram of the upgraded network and hardware configuration of CARMEN. Note the separation of file server and firewall roles from the head node to dedicated nodes.	154
A.5	Photo showing the dedicated file server and RAID array with bay-doors. Red tags denote spare disk and hot-swap disk, for fail-over capability.	155
A.6	Photo of the Sandiego cluster rack from Sun Microsystems	159
A.7	Diagram showing how the IPMItool application manages the Sandiego node platforms through out-of-band communications with the BMC and in-band keyboard controller style (KCS) communications with the server platform.	160

List of Tables

2.1	Zero-temperature properties of the lowest-lying heavy quarkonia spectrum ($b\bar{b}$ and $c\bar{c}$) [80].	35
3.1	Principle parameters of the Large Hadronic Collider at CERN [120], [141].	43
3.2	Resolution parameters for the Inner Tracking System for single and double tracks	46
3.3	Tracking chamber physical characteristics. The pad size of the tracking chambers is given for the two segmentation regions (inner and outer). . .	58
3.4	Table of latencies for the forward subsystems (including the ITS Pixels and PMD, which are included in the MUON event) and central detectors [150].	62
4.1	Foreseen ALICE HLT / DAQ modes of operation	72
4.2	DDL connections and data event payloads for the detectors in the HLT system. The rates are limited in TPC co-incidence mode to 200 Hz in Pb-Pb running and 1 kHz in p-p running (see text). The dimuon spectrometer can be triggered at 1kHz if no co-incidence with the TPC is required. . . .	75
5.1	Search radius parameters for the dHLT tracking algorithm for the region of interest (RoI) search.	93
7.1	Cross-sections in proton-proton collisions (σ_{pp}) and nucleus-nucleus collisions (σ_{AA}) for the lowest-lying $b\bar{b}$ family resonances [142] (including feed-down factors). Also shown is the resonance yield and number of dimuons from resonance decay expected in the data taking year, as well as the number of these dimuons expected in the spectrometer acceptance.	117
7.2	Cross-sections in proton-proton collisions (σ_{pp}) and nucleus-nucleus collisions (σ_{AA}) for the lowest-lying $c\bar{c}$ family resonances [142] (including feed-down factors). Also shown is the resonance yield and number of dimuons expected from resonance decay expected in the data taking year, as well as the number of these dimuons expected in the spectrometer acceptance. . .	117

7.3	Statistics of the various data sets used in the dHLT simulation	118
7.4	Percentage of tracks fulfilling various trigger selections for the J/ψ . The baseline criterion is the case where the track results in a hit on stations 4 and 5 and a hit on each of the trigger stations. In the case of a dHLT trigger with now L0 trigger, the percentage is of total L0 triggers.	126
7.5	Percentage of tracks fulfilling various trigger selections for the Υ data sample. The baseline criterion is the case where the track results in a hit on stations 4 and 5 and a hit on each of the trigger stations. In the case of a dHLT trigger with now L0 trigger, the percentage is of total L0 triggers.	126
7.6	Percentage of tracks fulfilling various trigger selections for the “minimum- bias” data sample. The baseline criterion is the case where the track results in a hit on stations 4 and 5 and a hit on each of the trigger stations. In the case of a dHLT trigger with now L0 trigger, the per- centage is of total L0 triggers.	128
7.7	Formation and dissociation times for the SU(3) [171] and SU(N) plasma [111] for the $b\bar{b}$ system.	131
7.8	Formation time and dissociation temperatures for the $c\bar{c}$ system in the 3- flavours plasma [171] and SU(N) [111] plasma case.	132
7.9	The percentage of L0 high p_T and asymmetric (L0 low p_T and dHLT high p_T triggers passed on average for the scenarios of no QGP production and the SU(N) and SU(3) cases.	138

University of Cape Town

Chapter 1

Introduction

1.1 Heavy ion experiments in context

The subset of physics today, which may be broadly described as “particle physics”, has been developed into a reasonably coherent, but somewhat incomplete description of nature on the very small scales relevant to the fundamental particles of the so-called “Standard Model”. The aim of this field of science is to find a description of the basic constituents of matter and the interactions between them. While there has been significant success in uniting these interaction theories at various scales, there are many areas which are still largely exploratory, with little more than phenomenology to guide the progress of the theory. This thesis deals with just such an area – the field of nuclear interactions at extreme conditions of energy density, temperature and pressure, known as ultra-relativistic heavy ion physics. Although progress in what is generally referred to as particle physics has been aimed at trying to determine with greater precision the microscopic properties of the theory with the aim of finding a unifying theory, heavy ion physics has attempted rather to investigate the large-scale aspects of the theory.

The experimental advance of this field has mostly been through the collision of nuclei with accelerators, or investigation of such collisions induced by high-energy cosmic radiation. The fundamental interactions between the constituents of nucleons - quarks and gluons, collectively termed “partons” - are described by the theory of Quantum Chromodynamics (QCD) [1]. This theory is non-Abelian and is described in terms of a scale, which is related to the distances of the interaction. At short¹ interaction distances, the coupling becomes weak and the theory can be treated perturbatively [2], a phenomenon known as “asymptotic freedom”². One of the consequences of asymptotic freedom is that there can exist a state, under conditions of high temperature and pressure, such that there is a very large parton density (and correspondingly a very small separation on average). Partons are thus free to ex-

¹In this case, “short” is compared to the typical hadronic size, which is of the order of 10^{-15} m.

²In 2004, the Nobel Prize Committee recognised the significance of the discovery of asymptotic freedom by awarding the Nobel Prize for Physics to work in this field. The prize was awarded to David J. Gross, H. David Politzer and Frank Wilczek “for the discovery of asymptotic freedom in the theory of the strong interaction”

plore larger volumes than in single hadronic systems. This state, termed the “quark soup” [3] or more commonly “quark-gluon plasma” [4], is generally believed to have been created in the very earliest moments of the universe’s evolution and may be reproducible in the laboratory.

While there has been remarkable agreement between experiments and perturbative aspects of this theory, finding solutions other than in the range of small coupling constant has so far proven to be elusive. The latter domain excludes all of the known stable configurations of nuclear matter, including the familiar neutron and proton. It is for this reason and others that this field has been driven in recent years by experimental advances and phenomenological models.

1.2 The search for the QGP

Since the creation of this state of matter requires a huge energy density, there are few experimental techniques to study the Quark-Gluon Plasma (QGP) in the laboratory. The most useful method – if not the only – so far for gaining experimental insight into the problem has been the study the collisions of heavy nuclei at high energy. By colliding heavy nuclei, a large, extended volume of pressurised nuclear matter can be created and the fundamental interactions can be investigated under these extreme conditions.

On the other hand, several purely theoretical advances have followed an intensively numerical approach, based on lattice gauge theory, solving QCD on a discretised lattice. These calculations, as well as estimates based on thermodynamic arguments, suggest that at sufficiently low baryonic densities and high temperatures, there is a phase transition from normal hadronic matter to a QGP, above a critical temperature T_c of about³ 160-200 MeV, or an energy density of about 1 GeV/fm³ [5].

1.2.1 Early experimental searches for the QGP

Experimental attempts to produce the QGP in the laboratory started more than twenty years ago, with exploratory research done at a few facilities at leading laboratories. These included Brookhaven National Laboratory (BNL)’s Alternating Gradient Synchrotron (AGS), European Centre for Nuclear Research (CERN)’s Super Proton Synchrotron (SPS) and Lawrence Berkeley National Laboratory (LBNL)’s Bevelac. These accelerators collided protons and light ions up to ³²S; subsequent upgrades allowed to accelerate even heavier species - Au and Pb. The centre of mass (cms) energies of these accelerators grew from $\sqrt{s_{NN}} \sim 1$ GeV at the Bevelac, $\sqrt{s_{NN}} \sim 5$ GeV at the AGS to $\sqrt{s_{NN}} \sim 17$ GeV and ~ 19 GeV (for Pb-Pb and S-S respectively) collisions and at the SPS.

These early experiments established the validity of the field, by demonstrating that the high energy densities and abundant particle production necessary for a statistical QCD description to apply were actually produced. In fact, the thermodynamic description of a strongly interacting system heated to a temperature of around 160-200

³A temperature is inferred, assuming a Boltzmann distribution of particles

MeV seems to be very efficient in describing the bulk properties [7, 8, 9] of the system created in such collisions.

Although there have been several exciting discoveries at CERN's SPS, culminating with the statement [11, 12] in the year 2000 that a "very dense state of matter" had been formed, decisive proof that the QGP was created there is still not claimed. The SPS experimental evidence remains somewhat circumstantial; with no direct evidence of the presence of a QGP, the data does not rule out purely hadronic descriptions altogether, although the only consistent interpretation seems to be that of a QGP. The data sets collected at the SPS are being progressively re-analysed and enhanced, however, and this situation may change in the near future.

Along with the creation of the dense and hot state by the collision of heavy ions, it is predicted that the chiral symmetry which is broken at normal temperatures and pressures would be restored temporarily. The restoration of chiral symmetry could provoke the creation of a metastable state called a Disoriented Chiral Condensate (DCC) (see for example reference [13] or more recently [14]). The formation of a DCC would result in an excess of low-momentum pions in a preferred direction in isospin space, thus giving rise to large fluctuations in the ratio of charged to neutral pions in a collision. The observation of so-called "Centrauro events" reported in cosmic ray literature, where clusters of mostly charged pions and very few π^0 were detected [15] was strong motivation for the existence of Disoriented Chiral Condensate (DCC)s. Experiments, such as the dedicated MiniMax experiment at Fermilab [16] or WA98 [17] at CERN have so far not been able to find any evidence for the creation of a DCC in heavy ion collisions, however.

1.2.2 The Relativistic Heavy Ion Collider

The experimental program in this field was further extended with the construction of the RHIC at BNL in the late 1990's. This facility became operational in July 2000 and aimed to explore a wider set of observables than previously possible at the SPS, with a set of four complementary experiments: Solenoidal Tracker At RHIC (STAR), Pioneering High Energy Nuclear Interaction eXperiment (PHENIX), PHOBOS and Broad Range Hadron Magnetic Spectrometers (BRAHMS). Unlike the previous experiments at the Bevelac, AGS and SPS accelerators, which were fixed-target experiments, RHIC collides two counter-rotating beams of nuclei. This translates into a higher $\sqrt{s_{NN}}$ for a given beam energy compared to fixed-target experiments, which means there is more energy available for particle production.

To date, RHIC has collided a wide variety of nuclear systems, including Au+Au, d+Au, p+p and most recently Cu+Cu, at energies of $\sqrt{s_{NN}} = 62.4 - 200$ GeV. There have been many surprises and vindications from the analysis of the data produced at RHIC, including some decisive statements by leading theoreticians that the QGP has definitely been created [18]. However, even if at the highest RHIC energies there is some direct evidence [19] for the *existence* of the QGP, there is not much room to move when it comes to the study of the deeper *nature* and *characteristics* of the plasma.

1.3 The Large Hadron Collider

The gap between proving the existence of the QGP by observing its effects, and probing its nature and characteristics implies an analogous gap in collider energies required. This gap should be closed, to a large extent, by the Large Hadronic Collider (LHC) currently being commissioned at CERN. Colliding heavy ions up to such large systems as Pb+Pb at energies of $\sqrt{s} = 5.5$ TeV per nucleon, the LHC will be the only machine on earth that will reach cosmologically relevant energies, such as existed in the very early stages of the evolution of the universe.

Extrapolations of observables from current data to such high energies vary significantly, but generally predict a favourable environment for all aspects of the study of the QGP. The energy density, size and lifetime of the system will increase by large factors and, with energy densities well above the deconfinement threshold, it should be possible to probe the QGP in its asymptotically free state. In contrast to current experiments, the central rapidity region in the centre of mass frame will have essentially zero net baryon density, dominated by an initial pre-equilibrium phase including semi-hard partonic interactions. This would significantly increase the initial temperature and ensure rapid thermalisation.

The LHC will open up a qualitatively different regime of particle physics to scrutiny. Four experiments – A Toroidal Large Acceptance Spectrometer (ATLAS), LHC b factory (LHCb), Compact Muon Solenoid (CMS) and A Large Ion Collider Experiment (ALICE) – will be present at the LHC. Of the four, ALICE will be the only one dedicated⁴ to the study of heavy ion collisions and the search for the QGP and as such, it will form the backdrop to this thesis. The experimental apparatus at the LHC will have to mirror this advance to take advantage of the physics environment. This thesis concerns the study and development of a high level trigger for a sub-detector of ALICE, the dimuon spectrometer. This subsystem will focus on the detection of muons resulting from the collisions and the identification of muons from the decay of heavy quarkonia such as the J/Ψ and Υ mesons, since one of the most widely accepted direct signs that a QGP is formed is an “anomalous” suppression in the yield of these mesons. Due to the high levels of muonic background and small branching ratio of the decay of these mesons into muons, in the absence of an online filter the proportion of useful data written to tape would be very small. Consequently a level of filtering will be used to reduce the level of background and reject unwanted events.

1.4 Thesis outline

This thesis concerns the design, development and testing of this trigger and a simulation of its effects on the foreseen data set. Chapter 2 concerns an introduction to the theoretical background of the field and introduces some of the relevant concepts and history. A description of the ALICE experiment can be found in chapter 3, with

⁴It should be noted that the CMS and ATLAS experiments have developed heavy-ion programs in addition to their dedicated p-p physics program. Similarly, ALICE also plans to study p-p data.

emphasis on the backward dimuon spectrometer subsystem. An outline of the design and functionality of the ALICE HLT is given in chapter 4, but the focus of the work will be in chapters 5 and 6, where the design and implementation of the dHLT will be described in detail. Chapter 7 is reserved for a study of the performance of the dHLT with respect to realistic spectra, while a discussion and outlook will be found in the 8th and final chapter.

University of Cape Town

Chapter 2

Theoretical Overview and Background

In the age of so-called “modern physics”, there has been a recurring theme – the study of matter under more and more extreme conditions, outside of our everyday experience. This theme, applied to the nucleus, has been the investigation of the behaviour of nuclear matter under extreme conditions of temperature, density and pressure.

Once the fundamental theory of nuclear matter, QCD, was discovered, it was also realised that the confinement of quarks in nuclei was due to the non-perturbative structure of the QCD vacuum. Along with this came the realisation that changes in the QCD vacuum at high temperature or pressure could change the confining behaviour of QCD and allow for the creation of deconfined states which are accessible to observation in the laboratory.

In practice, gaining a quantitative understanding of the QCD phase transition has proven to be a daunting task, due to the non-perturbative nature of the theory. Early progress was made in the direction of phenomenological models, such as the well-known MIT “Bag” model described in the next section. However, a more satisfying description of the QCD phase transition relied on solving the QCD field equations. It was found that this was possible by invoking a lattice and proceeding numerically. The rapid development of computers has helped to solve QCD on a discretised lattice and custom-built machines now exist expressly for this purpose.

2.1 The MIT Bag Model

The MIT Bag Model [20] is a phenomenological model based on classical statistical thermodynamics and was one of the first and most successful approaches at describing deconfinement phenomenologically. This model is particularly applicable when the baryonic chemical potential (μ_B), which describes the relative abundance of quarks to anti-quarks, is close to zero. Here, one deals with massless quarks and, in this case, the model predicts a first-order phase transition at a temperature known as the ‘critical temperature’, T_c .

Assuming, in a simple example, that the system consists only of massless pions,

the general expression of the pressure of such a system with n degrees of freedom is given by the Stefan-Boltzmann law :

$$P = n \frac{\pi^2}{90} T^4, \quad (2.1)$$

where T is the temperature of the system and n is the number of degrees of freedom. In the case of pions, the number of degrees of freedom is $n = 3$ (for π^\pm and π^0), but in the case of a QGP, describing a gas of light quarks and gluons, the degeneracy factor is $37 : 1 \times 2 \times 8$ for gluons (for spin and colour degeneracy respectively) and $7/8 \times 2 \times 2 \times 2 \times 3$ for quarks (for spin, flavour, isospin and flavour respectively). This describes a gas of free quarks and gluons, but in order to get the equation of state of such matter as we know it, we have to add a confining term. This is the so-called ‘‘Bag Constant’’ B , which represents in effect the pressure of the vacuum on the system and can be evaluated by considering the characteristic properties of the proton.

The energy in the bag, in equilibrium with its surroundings, can be written as a function of the radial dimension, R :

$$E(R) = B \times V_0 + \frac{C}{R}, \quad (2.2)$$

where V_0 is the volume of the bag and C/R is the energy of the quarks in the bag. Minimising this quantity with respect to R gives

$$E = 4BV_0. \quad (2.3)$$

In the case of the proton, where $E \sim 1\text{GeV}$ and $V_0 = \frac{4}{3}\pi R^3$, where $R \sim 0.7\text{fm}$, one finds a value for the bag constant of

$$B \sim 175\text{MeV}\cdot\text{fm}^{-3}. \quad (2.4)$$

Solving for T from the equation of state $P = 37\pi^2/90 \times T^4 - B$ gives a critical temperature of

$$T_{crit} \sim 200\text{MeV}. \quad (2.5)$$

2.2 QCD and lattice QCD

While perturbative QCD had many successes in describing high four-momentum transfer¹ phenomena in particle physics since its discovery, the real application of QCD to nuclear effects remained for a time impractical. It was realised that numerical analysis of lattice-regularised quantum fields could yield quantitative information on the non-perturbative properties of QCD [24, 25] and this quickly led to the application of discretised lattice calculations of the non-perturbative aspects of QCD. This approach allowed also to study the characteristics of the phase transition [26] as well as the equation of state of the plasma [27], in a thermal framework.

¹Also called ‘‘high- q^2 ’’, where q^2 is the four-momentum transferred during the process

The starting point of this investigation into the behaviour of QCD matter is the QCD Lagrangian \mathcal{L} , from which the partition function can be derived:

$$Z(T, V) = \text{Tr}(e^{-\beta \times H}) = \int dA \cdot d\psi \cdot d\bar{\psi} e^{-\int d^4x \int d\tau \mathcal{L}}. \quad (2.6)$$

Using statistical thermodynamics, one can then calculate many thermodynamic observables from the partition function, such as the energy density,

$$\epsilon = \left(\frac{T^2}{V}\right) \left(\frac{\partial \ln Z}{\partial T}\right)_V \quad (2.7)$$

or pressure,

$$P = T \left(\frac{\partial \ln Z}{\partial V}\right)_T. \quad (2.8)$$




Figure 2.1: The energy density and pressure of different systems of quarks calculated on the lattice with QCD [29]. Note the large increase in the energy density around the critical temperature $T = T_c$.

Many such simulations on the lattice have been done [28] and they show most notably that there is a phase transition associated with deconfinement around a critical temperature of $T_c = 150 - 200$ MeV. This is accompanied by a very large increase in the energy density (often associated with the latent heat of the transition). Lattice QCD, which encompasses all of the features of QCD, both perturbative at high q^2 and non-perturbative at low q^2 , thus also shows that a phase transition is predicted. A schematic 3-flavour phase diagram of lattice QCD is shown in figure 2.2, where the regions of first- and second-order phase transitions are shown, as well as the intermediate region, for which a rapid crossover is predicted. The second-order phase transitions lie on the boundary between the first-order transitions and the crossover region. The region where it is thought that the actual physical point (which represents our universe) lies is also shown, at two very light flavours (u and d) and one heavier flavour (s).

It is encouraging to see that both a phenomenological model like the Bag Model and the full theory predict the same phenomenon: a phase transition to deconfined matter at around $150 - 200$ MeV.

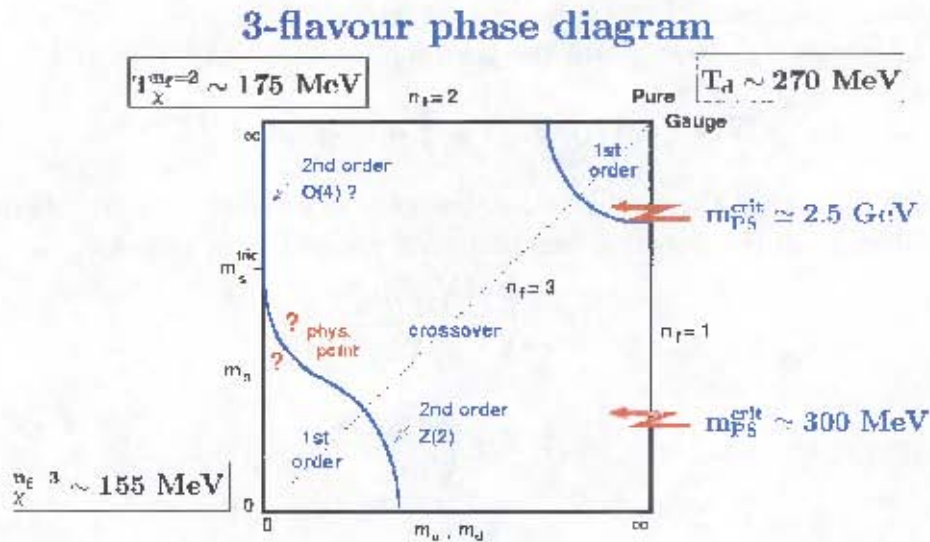


Figure 2.2: A phase diagram of QCD matter, showing the dependence of the phase transition on the mass of the constituent quarks.[29].

2.3 Ultra-relativistic heavy ion collisions

We have seen that a phase transition is predicted both from phenomenological arguments as well as numerical simulations of the fundamental theory of nuclear matter. The conditions for this phase transition from quarks and gluons confined within hadrons, to a state of deconfinement where partons are free to explore a volume larger than their parent hadron, are visible from the QCD $T - \mu_B$ phase plane shown in figure 2.3.

Where there is exceedingly large density (large μ_B), as well as exceedingly high temperature, there is predicted to be a phase boundary which the system will cross, into a state of deconfinement.

These conditions are expected to have been met in the very earliest stages of the evolution of the universe, when both the density and the temperature of the matter were at their most extreme. Of course, the evolution of the universe is unidirectional, but there are other situations which satisfy these conditions. It is thought that in the core of neutron stars, where the density of matter is at its greatest allowable value (precluding the formation of a black hole), a phase transition to deconfinement may take place. However, this scenario is inaccessible to repeatable experiments which aim to investigate the properties of a quark gluon plasma in the laboratory, which leaves heavy ion collisions as one of the only viable alternatives.

Given sufficient energy, these would create a state of large enough volume, under extreme pressure. This is an ideal environment to study the behaviour of nuclear matter, since by varying the centre of mass (cms) energy of the beam, the mass of the target or projectile, and depending on the detecting apparatus, the experimenter is able to vary most of the relevant experimental parameters. Experimental programs

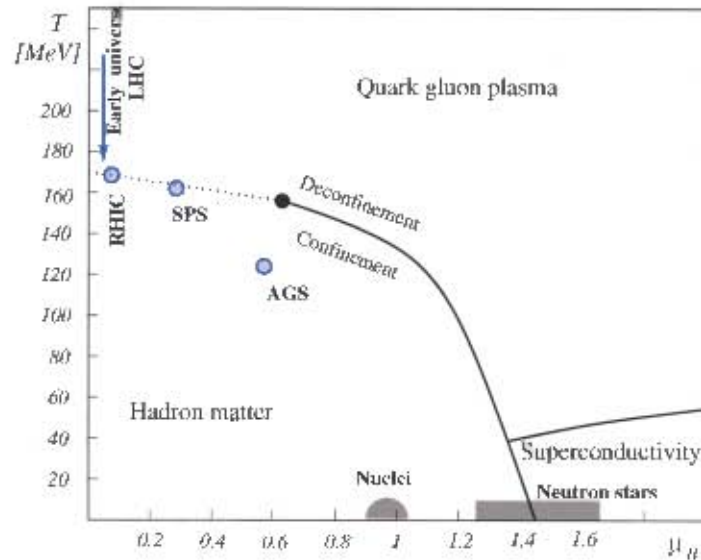


Figure 2.3: Schematic phase diagram showing the boundary between confinement and deconfinement and exotic states at high μ_B . This phase diagram indicates that a tri-critical point may exist around $\mu_B = 700$ MeV and indicates the regions (blue points) explored by current and future heavy ion experiments

at several laboratories have embarked on this kind of very systematic study of the properties of the nucleus, the most recent operational program being the Relativistic Heavy Ion Collider at Brookhaven National Laboratory in New York.

Here, we call values “large” and “high” when compared to the values normally associated with the nucleus, which has a density of about $0.17 \text{ nucleons fm}^{-3}$. The nucleons inside are on average 1.8 fm distant from each other, which is large compared to the size of the nucleon and as such they are considered individual entities. If the pressure of the nucleus is increased 5 or 6 times, or the temperature is increased to around $150\text{-}270 \text{ MeV}$, the number of hadrons per unit volume grows to greater than unity and they can therefore no longer be considered as individual particles, since they overlap.

In order to measure any effects due to this anomalous state, the system created also has to live long enough to thermalise and has to be large enough. As noted above, the only practical way to recreate such a system in a repeatable way in the laboratory under controlled conditions, is the collision of high energy beams of nuclei. Some of the kinetic energy of the beam of nucleons is used to compress the nuclear matter and heat it, creating a system which is favourable to the creation of a QGP.

Heavy ion collision programs have been progressing at a number of laboratories around the world, with an impressive and growing set of data, spanning a wide range of system size, centre of mass energy and phase-space coverage.

2.3.1 Kinematic variables

Before beginning a detailed discussion of the physics of heavy-ion collisions, it is useful to outline the kinematic variables that will primarily be used during this work, and which are standard to the field. Since the vast majority of particles concerned are relativistic, it is not convenient to use velocity as a kinematic variable, but rather a variable called “rapidity”, which will is defined as

$$y = \frac{1}{2} \ln \frac{E + p_z}{E - p_z}. \quad (2.9)$$

This variable has the range $-\infty \geq y \geq \infty$, but differentials of it, dy , are relativistically invariant, whereas those of velocity are not. A related variable is η , the pseudo-rapidity, which is given by

$$\eta = -\ln \tan \frac{\theta}{2}, \quad (2.10)$$

and is a valid approximation to the rapidity when $p \gg m$. For very high-momentum light particles, such as pions, this approximation allows an estimation of the rapidity without having to measure the particle’s energy or momentum explicitly, but only the angle at which it was emitted in the cms frame. The value $y = 0$ is referred to as “mid rapidity”, while $y > 0$ and $y < 0$ are referred to as “forward” and “backward” rapidities respectively.

Also, in a nuclear collision, the energy and momentum of all particles is initially primarily in the longitudinal plane (parallel to the beam axis). This means that the transverse plane contains the energy and momentum *transferred* during the collision; the relevant variables for the study of the produced particles then become not the momentum or mass, but rather the transverse momentum p_T and the transverse mass m_T , defined as [80]

$$m_T^2 = m^2 + p_x^2 + p_y^2 = m^2 + p_T^2. \quad (2.11)$$

These kinematic variables for the longitudinal (y) and transverse (p_T, m_T) planes will be used throughout this work.

2.3.2 Collision dynamics

In a heavy ion collision at ultra-relativistic energies (where the energy of the colliding nuclei is much larger than their rest mass), the energy of the particles in the beam is essentially available for particle production. The amount of the energy available for the particle is determined by two factors: the centrality of the collision and the centre of mass energy of the colliding system. The centrality of a collision is usually quantified by three parameters: the impact parameter (distance of closest approach of the centres of the nuclei), the total number of “participants” in the collision and the number of “binary collisions”, which will be defined below.

Since contemporary large heavy ion colliders, such as RHIC and SPS and the currently under construction LHC, collide beams of nuclei at very near light speed, relativistic mechanics plays an important role in describing the collision. The incoming nuclei experience a contraction in the laboratory frame along the beam direction

proportional to γ , the relativistic factor. At such high velocities, the duration of the collision is extremely short², only lasting a few $\text{fm}\cdot\text{c}^{-1}$. The higher the collision energy of the two nuclei, the more they become “transparent” to each other, resulting in a decrease in the baryonic potential at mid-rapidity. Thus, one is able to make connections between measured data and observables in models which rely on a vanishing baryonic potential.

The main connection between the theoretical predictions and the observed quantities in heavy ion collisions is the transverse energy released. This directly observable quantity is related through a simple model to most theoretical signatures of the QGP. Other global observables are useful for characterising classes of events, but specific signatures can usually only be discussed in reference to a particular model, since the plasma lives for too short a lifetime to be observed directly

2.3.3 The evolution of the QGP

The system created by the collision of nuclei at these high energies undergoes a time evolution, expanding and cooling rapidly as it does so. Initially, it is extremely dense and hot and is dominated by different processes as it evolves. A schematic diagram of this evolution is shown in the Minkowski space-time diagram in figure 2.3.3. This is a representation of the so-called Bjorken picture [31] of the initial state of the system created during the collision (see section “transverse energy” 2.4.1). The interactions of the system are dominated by a cascade of hard parton scattering and it is conjectured that the evolution of this phase can be calculated within the perturbative QCD framework, or by phenomenological processes, like the Lund “string breaking” model [99]. After a short time of about 1 fm, the created system thermalises through re-scattering and, once thermalised, relativistic hydrodynamic transport models are considered to give the relevant description.

As the system expands and cools according to relativistic hydrodynamics, the partons combine into hadrons. Soon after the hadrons have formed, they stop interacting inelastically, after which stage no more particle species are produced. This stage is known as “chemical freeze-out”, where the particle yields and ratios are fixed. The system may, however, continue to thermalise through elastic collisions. The point at which elastic collisions amongst hadrons cease is known as “thermal freezeout” and is the stage where momentum distributions are fixed. The particles no longer interact in any way, except to decay via the weak force and are free-streaming towards the detectors.

2.3.4 Experimental heavy ion collision programs

Over the last two decades a huge amount of experimental data has been collected in the field of relativistic heavy ion collisions, at various colliders throughout the world. Most notably, at BNL’s Alternating Gradient Synchrotron (AGS) and Relativistic Heavy Ion Collider (RHIC) as well as the European Centre for Nuclear

²In ultra-relativistic collisions, where most particles are moving very near the speed of light, times are customarily measured in units of length : $x = v\cdot t = c\cdot t$, where $c = 1$




Figure 2.4: Diagram of the space-time evolution of a heavy ion collision, in the presence (right) and absence (left) of a QGP [21].

Research (CERN)'s Super Proton Synchrotron (SPS), experiments have been running dedicated to the creation and investigation of the QGP. At the SPS, a variety of fixed target experiments, from very light systems (^{16}O and ^{32}S at beam energies of $\sqrt{s} = 200$ GeV) to very heavy (^{208}Pb at a beam energy of 158 GeV) have been used in seven experiments in the West and North Areas of the accelerator complex. At around the same time, the AGS was colliding beams of ^{28}Si at $\sqrt{s} = 14.5$ GeV/A and ^{197}Au at $\sqrt{s} = 11.6$ GeV/A.

From the vast amount of combined information from these experiments, it has been shown that the chemical freeze-out of hadrons occurs at about 170 MeV at the SPS [6] and about 130 MeV at the lower energy AGS [10]. Considering the critical temperature predicted by both phenomenological QCD (MIT Bag Model) and Lattice QCD, there are qualitative thermodynamic reasons to believe that a QGP should be created at these energies. There also are so-called "abnormal" deviations from the normal hadronic descriptions, *i.e.* descriptions with no deconfinement. Most notable among these is the anomalous suppression of the yield of J/ψ meson in central Pb-Pb collisions due to the onset of Debye screening, measured by the NA50 and NA60 collaborations. This yield was termed anomalous because it could not initially be

explained in terms of the conventional hadronic absorption scenario [22], although subsequent analyses have disputed the point, by extending the hadronic absorption scenario [23].

Signals like the anomalous suppression of the quarkonia yields, strangeness enhancement, and others have been and continue to be the topic of intense theoretical debate and experimental scrutiny. Since the QGP lives for far too short a time to be measured directly, its existence and nature has to be inferred indirectly from the properties of the observed hadrons. Some of these indirect signals will be discussed below, with particular emphasis on the suppression of quarkonia yields, which is the observables which the ALICE dimuon spectrometer is most specifically designed to measure.

2.4 Observing the QGP

Since the formation of a QGP depends critically on the initial energy density (and hence temperature), it is important to define certain “global” observations which will give an experimental handle on these quantities. The initial energy density, temperature and other thermodynamic variables are, however, not directly observable, and have to be inferred from other global observables. There do exist, however, specific observables which are directly measurable and are related to the existence and nature of a QGP. The global observables and their relation to the created state will be described below, as well as a few representative direct probes of the QGP.

2.4.1 Global observables

As discussed above, a prerequisite for a phase transition into a QGP is that the energy density of the created state is greater than the critical value predicted by QCD. From simple phenomenological estimates, as well as lattice QCD results, this is known to be around $1 \text{ GeV}\cdot\text{fm}^{-3}$. It has to be established, when analysing collision data, that an energy density of this magnitude or greater has been created.

Centrality

The term “centrality” refers to the distance of closest approach between the centres of the two nuclei being collided. Just as the QGP itself is not directly measurable through any but indirect means, so also the centrality of collisions is not a directly observable quantity. It is usually described in terms of three related quantities :

Impact parameter, b : The vector between the centres of the two nuclei as they collide.

Number of participants: The total number of nucleons which suffered at least one nucleon-nucleon collision during the interaction of the two nuclei

Number of binary collisions: The sum of all collisions between pairs of nucleons.

Particle yields in fundamental nucleon-nucleon collisions can be related to p-A or A-A collisions by scaling by the number of participants or the number of binary collisions, depending on which part of the spectrum is under investigation (hard or soft). These quantities are directly related to the impact parameter b via a Glauber model [30], but neither they nor the impact parameter are directly measurable.

Using the Woods-Saxon distribution of the nuclear density profile as a function of radial distance r

$$\rho_A(r) = \rho_0 \frac{1 - \omega \left(\frac{r}{R_A} \right)^2}{1 + e^{\frac{r-R_A}{z}}}, \quad (2.12)$$

where ρ_0 is the maximum density, R_A is the nuclear radius, z is the surface thickness and ω allows for irregularities, the nuclear thickness function can be defined as

$$T_A = \int dz \rho_A(z), \quad (2.13)$$

which represents the total number of nucleons seen by the incoming projectile. Denoting the target A and the projectile B , the nuclear overlap function of the colliding system (which is related to the number of participants in the collision) can be defined as

$$T_{AB}(\vec{b}) = \int d^2s T_A(\vec{s}) T_B(\vec{b} - \vec{s}), \quad (2.14)$$

where the vector \vec{s} in equation (2.14) is the position in the transverse plane at which the density is being calculated. This is normalised to minimum bias

$$\int d^2b T_{AB}(\vec{b}) = AB. \quad (2.15)$$

The A-A hard cross-section seen at a particular impact parameter can now be written in terms of the p-p hard cross-section, by scaling by the nuclear overlap fraction :

$$\sigma_{AB}^{hard} = \int \sigma_{pp}^{hard} T_{AB} d^2b \sim AB \sigma_{pp}^{hard}. \quad (2.16)$$

Since quarkonium production proceeds via hard scattering, the p-p cross-section can be scaled up to the A-A case in this way :

$$\sigma(AB \rightarrow q\bar{q}) = AB \sigma(pp \rightarrow q\bar{q}), \quad (2.17)$$

This binary scaling formula allows one to compare yields from p-p collisions to A-A, or predict yields from A-A collisions.

The number of participants and binary collisions at a given centrality can be inferred from the particle multiplicity, usually within the framework of a Glauber model [30]. A monotonic (but not necessarily linear) relationship between the number of participants and multiplicity is assumed, since this allows the experimenter to make the assumption that the more central a collision is, the greater the multiplicity. A given fraction of the multiplicity therefore infers a given fraction of the centrality.




Figure 2.5: An example of the centrality determination in the PHOBOS experiment at RHIC [32]. The correlation between the zero-degree calorimeters (at forward rapidities) and paddle counters (nearer mid-rapidity) is shown, with the corresponding calculation of the number of participants from a Glauber model

Experimentally, the correlation between energy deposited in “zero-degree counters”³ and detectors at other rapidities is usually used to estimate the centrality.

The energy deposited in forward detectors (see figure 2.5) is from particles which have not interacted in the collision and is therefore a measure of the number of spectators. Since the total number of nucleons is fixed by the colliding system, the number of participants is calculable once the number of spectators is known. A Glauber model can then be used to determine the number of collisions given the number of participants.

Transverse energy

Since one of the criteria for QGP formation is an energy density above the critical value of $\epsilon \sim 1\text{GeV}/\text{fm}^3$, it is important to reliably estimate the energy density of a collision. This value is something of a matter of definition, since one may artificially construct the initial energy density of Au-Au collisions, even with no interactions: if

³“Zero-degree counters” are detectors which are situated very near the beam line and detect primarily spectator nucleons. The term “zero-degree” here refers to the small angle between the beam line and the vector from the detector to the interaction point upstream.

each nucleus has an energy density ρ_0 in its rest-frame, then the energy density in the region where the nuclei overlap, is simply given by $\langle\epsilon\rangle = 2\rho_0\gamma^2$, where γ is the Lorentz boost factor. With the typical value of $\rho_0 = 0.14 \text{ GeV}\cdot\text{fm}^{-3}$ for the energy density of the nucleus, this estimate for the initial energy density is

$$\begin{aligned}\langle\epsilon\rangle &= 106 \text{ GeV}\cdot\text{fm}^{-3} && \text{(Full RHIC energy : 200 GeV per nucleon)} \\ \langle\epsilon\rangle &= 2931 \text{ GeV}\cdot\text{fm}^{-3} && \text{(Full LHC energy : 5.5 TeV per nucleon)}\end{aligned}$$

This value is almost absurd compared to the $1 \text{ GeV}\cdot\text{fm}^{-3}$ that is thought necessary for the formation of a QGP, but it is also not entirely relevant. The $1 \text{ GeV}/\text{fm}^3$ is the energy of the *created system* which refers to the secondary particles. These, by definition, are not present in the very initial stage of the collision and have to be produced via interactions, which have associated with them a characteristic time, depending on their q^2 . In the laboratory frame, the two colliding nuclei would appear as thin slabs of matter, contracted by their large Lorentz factor. From the moment of impact of these thin slabs, until some formation time τ_F has passed, the system cannot be said to have “formed”. Once this time has elapsed, we are left with essentially a cylindrical tube in the longitudinal (y) direction containing the newly created system. This is essentially the idea first put forward by Bjorken [31] in 1983.

The transverse energy measured at 3 RHIC energies by PHENIX [61] is shown in figure 2.4.1 as a function of the number of participating nucleon pairs (as described above). Also shown are the corresponding energy densities calculated in the Bjorken picture.

Assuming a formation time and longitudinal boost invariance, Bjorken showed that energy density of a given collision could be calculated from the transverse energy at mid-rapidity (defined as the energy of measured particles in the transverse direction). The relation

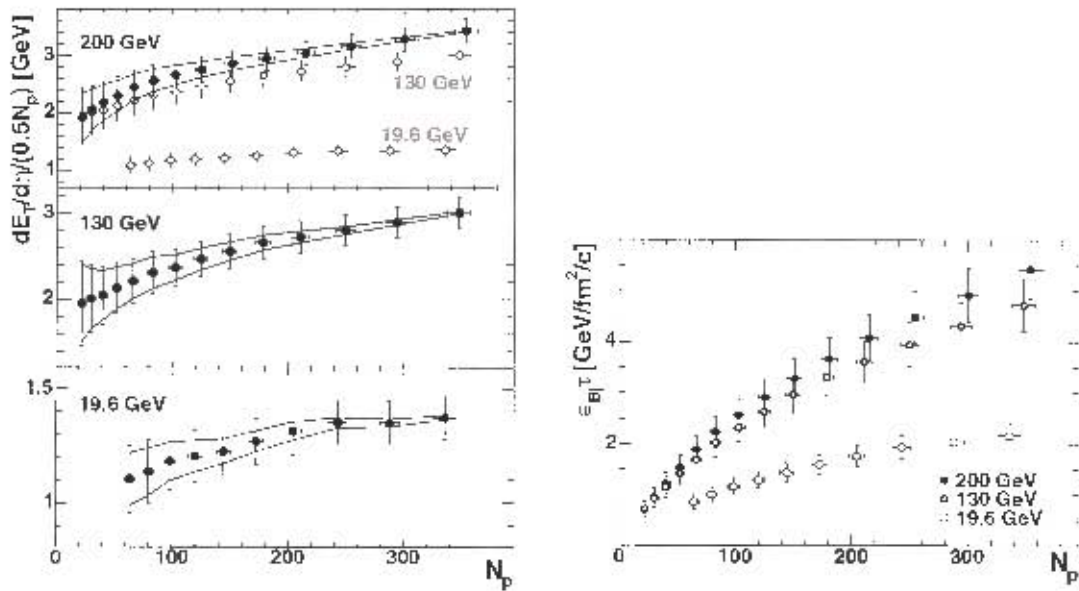
$$\epsilon = \frac{1}{\pi R^2 \tau_F} \left[\frac{dE_T}{dy} \right]_{y=0}, \quad (2.18)$$

where R is the radius of the colliding nuclei ($R = 1.2A^{1/3}$, for a nucleus of mass A) and $\tau_F \sim 1 \text{ fm}/c$ is the formation time of the system. The transverse energy measured *e.g.* at PHENIX [62] implies an initial post-thermalisation energy density of about $5 \text{ GeV}/\text{fm}^3$. This is far above the estimate of the initial energy density required for the creation of a QGP, from Lattice QCD and the Bag Model.

Particle multiplicities

Tightly related to the created energy density is the total particle multiplicity of heavy-ion collisions. This is one of the most fundamental observables, since it constrains the energy density which is itself used as an input to several other theoretical models. Despite its prime importance, there is no first-principles method of calculating the expected multiplicity at the LHC from the QCD Lagrangian, since this depends on non-perturbative soft processes. The integrated hadron multiplicity in p-p collisions is given by

$$\rho_h(y) = \frac{1}{\sigma_{\text{inel}}^{pp}(s)} \int \frac{d\sigma^{pp \rightarrow h+X}}{dy d^2 p_T} d^2 p_T, \quad (2.19)$$



(a) Transverse energy density as a function of participating nucleon pairs (b) Energy density calculated in the Bjorken picture, corresponding to the transverse energy in 2.6(a)

Figure 2.6: Transverse energy density as a function of the number of participating pairs in Au-Au collisions (left) and corresponding energy density calculated in the Bjorken picture, as measured by PHENIX [61].

where $\sigma_{\text{inel}}^{pp}$ is the total p-p inelastic cross-section. This starts to increase after about 20 GeV, but has a poorly-known dependence on \sqrt{s} [73]. The measured multiplicity in p-p collisions can not discriminate between a \ln or \ln^2 dependence on \sqrt{s} from the AGS to RHIC. Although the relevant parameter to scale p-p multiplicities by in order to attain the corresponding A-A values at a given \sqrt{s} is thought to be the number of participants in the nuclear collisions, the evolution with \sqrt{s} is not certain at all, and remains highly model-dependant. Figure 2.7, taken from [141], shows the charged particle density per participant pair in p-p and A-A data from SPS to RHIC and several model predictions, which are extrapolated to $\sqrt{s} = 5.5$ TeV. The fact that the expected multiplicity at ALICE had (and still has) such large uncertainties led to the fact that the experiment was designed with a lot of room for error in this margin. ALICE has been designed for an expected charged particle multiplicity at mid-rapidity of $\frac{dN}{dy}|_{y=0} = 8000$, which represents a safety factor of about 3 with respect to the current estimates, which range from 1000-2000⁴.

⁴In the case of the dimuon spectrometer, however, the main source of charged particles and hence background is secondary interactions in the beam shield, since this detector is situated near the beam-line. For this detector, primary particles from the Pb-Pb collision are emitted in the rapidity range of $4 < y < 7$, however the multiplicity estimate in this rapidity range has a larger error than that at mid-rapidity. Therefore, the ALICE dimuon spectrometer has been designed with a mid-rapidity multiplicity of up to 12 000.




Figure 2.7: Model predictions for charged particle multiplicities per participant in A-A and p-p collisions as a function of \sqrt{s} . The dashed line is the fitted function $0.68 \ln \frac{\sqrt{s}}{0.68}$, on the nuclear data, while the dotted line is a fit to $0.7 + 0.028 \ln^2 \sqrt{s}$. The long-dashed line is a prediction based on a saturation model [74]. The estimate for ALICE in these models is given by the intersection of the curves with the grey arrow on the right.

Baryon density

Recall that most simulations of QCD on the lattice are done in the regime of vanishing baryonic chemical potential ($\mu_B = 0$). In order to make a rigorous connection between the measured observables and the predictions of the theory, it is necessary to have a good handle on the baryon density of the created system.

One way of estimating μ_B is the baryon to anti-baryon ratio, which has been measured at the AGS [75], CERN SPS [76] and RHIC [77]. The rapidity density of net protons, which is (difference between the number of protons and anti-protons), is shown in figure 2.8.

It is seen that already at RHIC energies, for the most central collisions and in a rapidity region around $|y| < 2$, the baryon density is very low. This, coupled with recent advances in lattice QCD which can give predictions at non-vanishing baryon density [78], allows to compare confidently the observed data to the theoretical signatures.

2.4.2 Strangeness production

Anomalies, especially enhancement, in strangeness production has long been pointed to as a possible signature of formation of a QGP. Enhancement in the yield of strange particles, as well an increase in the ratio of strange baryons to baryons. Since there is initially only a negligible amount of strangeness present in the colliding nuclei,




Figure 2.8: Rapidity density of net protons measured at the AGS, SPS and RHIC energies, for central collisions. The beam rapidity of each collider is shown. Note the rapid fall-off of net proton density with beam energy.

all strange quarks have to be created by hard processes during the evolution of the system. As mentioned before, the relative abundances of particle species is well described by statistical hadronisation models, mostly inspired by thermal and chemical equilibrium. This is shown, *e.g.* for the system produced in Au-Au collisions at RHIC in figure 2.9.




Figure 2.9: Comparison between the particle ratios from Au-Au collisions measured by the four RHIC experiments, and statistical model calculations [102], with $T = 174$ MeV and $\mu_B = 46$ MeV.

At high energies where a QGP may be created, the $s\bar{s}$ creation can proceed mainly via gluon-gluon fusion, while in lower energy hadronic collisions they have to be

produced via pair creation or rescattering reactions of the type

$$\pi + N \rightarrow K + \Lambda \quad (2.20)$$

$$\pi + \Lambda \rightarrow K + \Xi \quad (2.21)$$

$$\pi + \Xi \rightarrow K + \Omega \quad (2.22)$$

These processes have a higher energy threshold and while the timescale for gluon fusion is of the order of 1-3 fm, that of pair creation and hadronic reactions is much longer. It has been argued that the strangeness production rate in an equilibrated hadron gas is similar to that of a QGP, but the equilibration time of a hadron gas (about 10 fm) is much longer than that of a QGP (about 3-5 fm). Apart from the increased rate of strange quark production in a QGP, it might also be energetically more favourable to create a bare $s\bar{s}$ pair than a light $q\bar{q}$. This is due to the fact that in the asymptotically-free perturbative (pQCD) plasma, the quarks take on their current masses (which can also be seen as a consequence of the mass-dependant component of the strong force potential falling away) – a phenomenon known as “chiral symmetry restoration”.

One of the most striking results of the SPS heavy ion program was the measurement of the yield of strange particles (particles containing one or more strange quark) as a function of “wounded nucleons”. Here, the wounded nucleon concept is similar to that of “participant” and designates a nucleon that has been “struck”, or suffered a collision. In light ion and p-A collisions at SPS, measured with the WA98 apparatus, the strange-particle yield shows no appreciable deviation from that expected of a hadronic system. In Pb-Pb collisions however, there was a significant enhancement in the strange particle yield – up to about a factor of 10 in the case of the Ξ^- . The yields were also enhanced differently, depending on the particle species. Figure 2.10 shows the enhancement of strange particles per event produced at mid-rapidity by Pb-Pb collisions at the NA57 experiment, compared with p-Be and p-B collisions. If the yields were due to a simple superposition of nucleon-nucleon interactions, one would expect the lines to be flat at unity.

The K^+/π^+ “horn”

A curious feature in the ratio of K^+/π^+ yields from heavy-ion collisions was discovered by the NA49 experiment. An unusually sharp peak, referred to as “the horn” [37] (see fig 2.11), is seen in this ratio as a function of $\sqrt{s_{NN}}$, which is not present in pp collisions. Statistical models of particle production predicted since many years a fast rise and subsequent plateau in this quantity [7, 100, 35], but a maximum was not predicted. It has also pointed out, however that the “horn” appears to co-incide with a transition from baryon-dominated to meson-dominated hadronic matter, which occurs at a temperature of about $T = 140$ MeV and a baryonic chemical potential of $\mu_B = 410$ MeV and $\sqrt{s} = 8.2$ GeV [36].

While thermal models calculations [42] reproduce quite well the peak in the $\Lambda/\langle\pi\rangle$ ratio as a function of energy, they do not reproduce the sharp peak in the K^+/π^+ ratio [36]. This sharp peak is not reproduced in transport codes either [43] (see figure

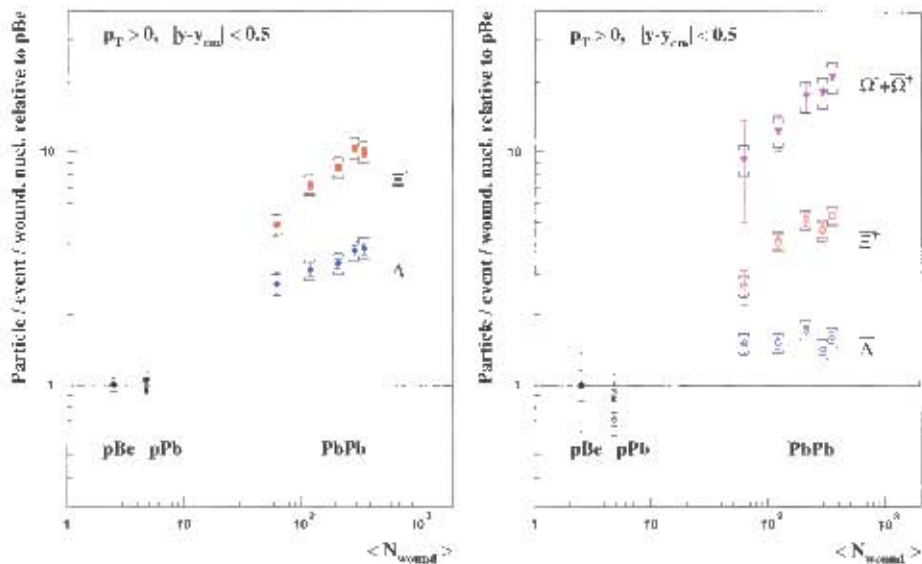


Figure 2.10: The yield of Λ , Ξ and Ω (containing one, two and three strange quarks respectively) in Pb-Pb collisions relative to the yield in p-Be collisions, as measured by the NA57 experiment [33], as a function of the number of wounded nucleons $\langle N_{\text{wound}} \rangle$. The yields are relative to the number of participating nucleons. Shown in the left panel are the yields for particles which may have had at least one quark in common with the projectile or target; the right panel shows the yield for particles whose constituent quarks would have had to be created during the collision.

2.12(a)). However, a sharp peak in the excitation function of the total strangeness to pion ratio E_s was predicted in the Statistical Model of the Early Stage [38], assuming that a phase transition to deconfined matter occurred and this indeed describes fairly well the “horn”, as shown in figure 2.12(b). The sharp rise in the K/π ratio is very suggestive evidence that a transition from a confined state of hadronic matter to a deconfined state of quarks and gluons is seen. There is, however, much debate around this point and several other explanations in terms of purely hadronic matter are under study.

2.4.3 Elliptic flow observables

In heavy-ion collisions, the initial spatial anisotropy is converted into momentum-space anisotropy during the initial phase of the evolution of the state created. The elliptic flow, which is defined as the second Fourier harmonic of the azimuthal anisotropy, is usually defined as

$$v_2 = \frac{p_x^2 - p_y^2}{p_x^2 + p_y^2}. \quad (2.23)$$

Important insights into the evolution of the system can be obtained by investigating the v_2 generated by various particles and as a function of centrality, because it is a measure of how strongly the system can convert a pressure gradient (produced by the anisotropy in spatial dimensions) into into a momentum anisotropy. Therefore, hy-

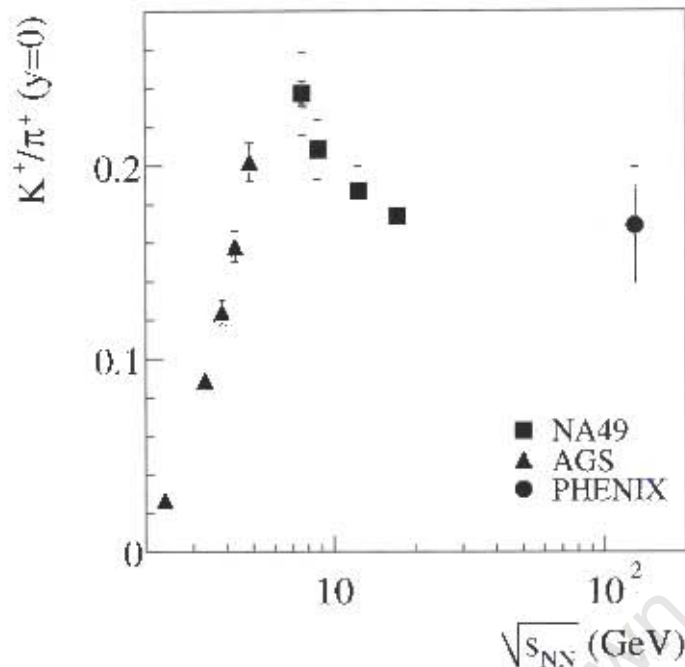



Figure 2.11: Ratio of $\frac{K^+}{\pi^+}$ yield at mid-rapidity as measured by the AGS RHIC and SPS (see references 23 and 24 in [39]).

hydrodynamic models are used to describe the behaviour of the flow of particles. Figure 2.13 shows that v_2 values of pions measured by STAR in Au-Au collisions at $\sqrt{s} = 200$ GeV are much higher than that of pions measured at the lower-energy SPS experiment NA49. Hydrodynamic calculations which describe ideal relativistic fluid flow describe the mass and p_T dependence of elliptic flow very well. This is yet more evidence of the formation of a deconfined state [47], as calculations using a hadron-gas equation of state severely underestimate the v_2 values obtained by STAR for π^\pm and protons at $\sqrt{s} = 130$ GeV, as shown in figure 2.14.

Elliptic flow studies also managed to give some insight into hadron formation. Models of hadron formation by coalescence or recombination of constituent quarks predict that at intermediate p_T , the values of v_2 of various particles will approximately scale with the number of constituent quarks (n) of the given particle. In this case, it was predicted that v_2/n vs p_T/n for all hadrons would lie on a single curve, and this curve would then represent the momentum-space anisotropy of the constituent quarks prior to hadron formation [67]. This is indeed seen to be the case - as is shown in figure 2.16, the scaled v_2 curves for protons, kaons and lambdas all lie very close to a single curve, over a very wide p_T range [68].

2.4.4 Specific signatures

In this section, a short representative list of specific observables which are predicted to be signatures of the creation of a QGP will be given. Some of the important recent



(a) Measured ratio of $\frac{K^-}{\pi^-}$ yields at mid-rapidity as a function of \sqrt{s} and predictions of UrQMD [43] and hadron gas models [35]. (b) Strangeness-to-pion ratio E_s as a function of F compared to a prediction from the Statistical Model of the Early Stage [38].

Figure 2.12: Measured ratio of $\frac{K^+}{\pi^+}$ yields as a function of $\sqrt{s_{NN}}$ (left) and strangeness-to-pion ratio $E_s = \frac{\langle \Lambda \rangle + \langle K + \bar{K} \rangle}{\langle \pi \rangle}$ as a function of $F = \frac{(\sqrt{s_{NN}} - 2m_N)^{3/4}}{\sqrt{s_{NN}^{1/4}}}$ (right).

measurements from RHIC will be shown, most notably the suppression of high p_T hadrons in central Au-Au collisions, however the focus will be on signatures most relevant to the ALICE dimion arm, which deal with heavy quarkonia ($c\bar{c}$ and $b\bar{b}$ families) suppression.

Direct photons

Photons are exceptional probes of the created matter, because once created they only interact via the electromagnetic force (and very weakly via the weak force). This means that if one measures the photon spectrum, one is in fact measuring components from the entire evolution of the created state, without further rescattering. Photons may be created during the initial phases of the collision, via parton scattering, or during thermalisation - the so-called “prompt” and “direct” photons. They may also be produced later in the evolution of the system, most abundantly from π^0 decay.

The direct photons would come from two stages of the evolution of the collision :

1. Very early interactions, in the pre-equilibrium phase of the created system
2. Photons created during the thermalisation of the system.

During thermalisation, electromagnetic process which could contribute to the photon spectrum include quark-antiquark annihilation ($q + \bar{q} \rightarrow \gamma + g$) and Compton scattering ($qg \rightarrow \gamma q$). Photons could also be produced during hadronisation. The measured photon energy spectrum contains contributions from all of these sources, which have to be disentangled in order to make statements about the temperature of the QGP. This is a nontrivial task, given the fact that the QGP photon yield is much smaller

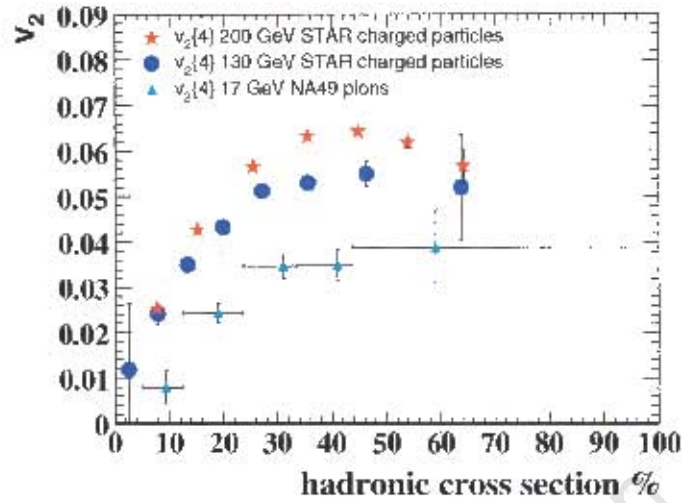


Figure 2.13: Centrality dependence of v_2 , integrated over p_T . The triangles are the NA49 measurements for pions at $\sqrt{s_{NN}} = 17$ GeV [44], while the circles and crosses are measured by STAR at $\sqrt{s} = 130$ GeV [45] and 200 GeV [46], respectively (figure from [47]).

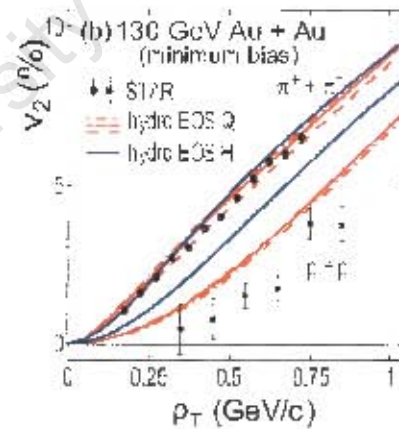


Figure 2.14: Hydrodynamics calculations for a hadron gas (EOS H) and QGP (EOS Q) equation of state [48], compared to STAR v_2 measurements for pions and protons in minimum bias 130 GeV Au-Au collisions [49] (figure from [47]).

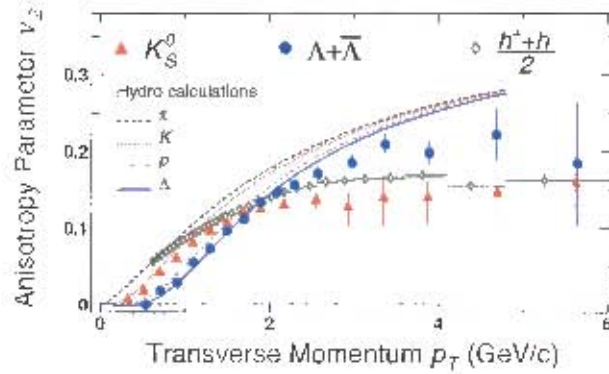


Figure 2.15: Minimum-bias v_2 for identified particles as a function of p_T [68] and hydrodynamic calculations [69].

at high p_T than the prompt and decay photons, as well as the fact that the spectra of the direct and hadronisation photons have a very similar shape.

Recent results on the measurement of photons from the PHENIX experiment [52] have proven interesting, however. The double ratio of

$$\frac{\left[\frac{\gamma}{\pi^0}\right]_{\text{signal}}}{\left[\frac{\gamma_{\text{decay}}}{\pi^0}\right]_{\text{noise}}} \quad (2.24)$$

is constructed in order to cancel out experimental uncertainties in the γ and π^0 measurements. The numerator in this ratio is the measurement of the yields, while the denominator is the simulated photon yield from hadronic decays, based on a fit to the measured π^0 spectrum, then divided by that fit [50]. Here, $\left[\frac{\gamma}{\pi^0}\right]_{\text{noise}}$ denotes the fraction of photons from the decay of pions to total number of pions. The systematic error on the measure decreases somewhat with p_T due to the high p_T suppression of pions observed by all four RHIC experiments in central Au-Au collisions.

Any deviation from unity of the above double ratio is due to the yield of direct photons above background. The recent results from PHENIX p-p collisions on this observable agree quite well with a QCD prediction as shown in figure 2.17. This double ratio for direct photons produced in Au-Au collisions shows a marked increase as a function of p_T for central events and there is a smooth transition as a function of centrality, as measured by PHENIX [50] and shown in figure 2.18.

A large yield of direct photons over contributions from meson decays is observed in more central Au-Au collisions, and the magnitude is consistent with a binary-scaled pQCD calculation.

Jet quenching

The phenomenon of hadron jets, first seen in p-p collision experiments, is due to the hard scattering of two partons, which acquire large transverse momentum. After fragmentation into hadrons, the relics of these initial hard partonic scatterings are




Figure 2.16: Identified particle v_2 , scaled by the number of constituent quarks, fitted to a polynomial. The ratio of data/fit is shown in the lower panel [68].

seen as jets of high- p_T particles. At leading order, jets are seen back to back, as the initial hard scattering conserves four momentum.

The situation is somewhat different in proton-nucleus or nucleus-nucleus collisions, compared to proton-proton collisions. In order to quantify the differences in the two systems rigorously, the so-called “nuclear modification factor” is used. This is defined as

$$R_{AA}(p_T) = \frac{\frac{1}{\langle N_{bin} \rangle} \cdot \frac{dN}{dp_T} \Big|_{AA}}{\frac{1}{\langle N_{bin} \rangle} \cdot \frac{dN}{dp_T} \Big|_{pp}}, \quad (2.25)$$

where $\langle N_{bin} \rangle$ is the average number of binary nucleon-nucleon collisions (defined as 1




Figure 2.17: Direct γ excess double ratio (see text) and cross-section in 200 GeV p-p collisions at RHIC [50]. The photon cross-section is compared with pQCD calculations for various scale factors from [51]




Figure 2.18: Direct γ excess double ratio (see text) in 200 GeV Au-Au collisions at RHIC [50] as a function of p_T for various centralities. The solid line represents a Next-to-leading order (NLO) Perturbative QCD (pQCD) calculation [51]

in the case of p-p or p – p collisions) and $\frac{dN}{dp_T}$ is the multiplicity spectrum measured in A-A or p-p collisions. This ratio cancels out the trivial deviations due to the larger system and provides a consistent way to compare nucleus-nucleus data to nucleon-nucleon data. Any deviation from unity of this measure is a sign that nuclear effects are coming into play during the collision, or put another way, that the nucleus-nucleus collision can no longer be seen as the simple superposition of nucleon-nucleon collisions, scaled by $\langle N_{bin} \rangle$.

This ratio has been measured in most heavy ion experiments which measure particle spectra. At the lower energy AGS and SPS experiments, a huge increase in the ratio was observed at intermediate p_T . The ratio rises sharply between 1 GeV/c and 2 GeV/c and remains high out to fairly high p_T . This has been attributed to initial state p_T broadening due to multiple scattering and has been dubbed the ‘‘Cronin effect’’ [53]. In stark contrast to these low energy experiments, when the RHIC experiments measured the ratio, they found a large suppression as a function of p_T all the way out to very high values (10 GeV in the case of PHENIX). The comparison between the p_T distribution of the nuclear modification factor at the CERN SPS and RHIC for nucleus-nucleus collisions is shown in figure 2.19.

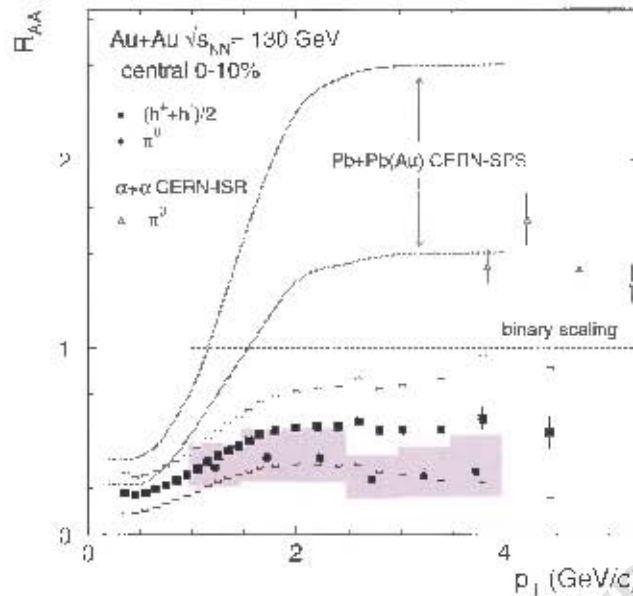


Figure 2.19: Nuclear modification factor R_{AA} for charged hadrons and neutral pions in the most central 10% Au-Au collisions at RHIC (solid points). Error bars indicate statistical errors while surrounding bands indicate quadrature sum of systematic errors. Also shown are the modification factors for $\alpha + \alpha$ at $\sqrt{s} = 31$ GeV [54] and central Pb-Pb collisions at the SPS [55], indicated by the dotted uncertainty band.

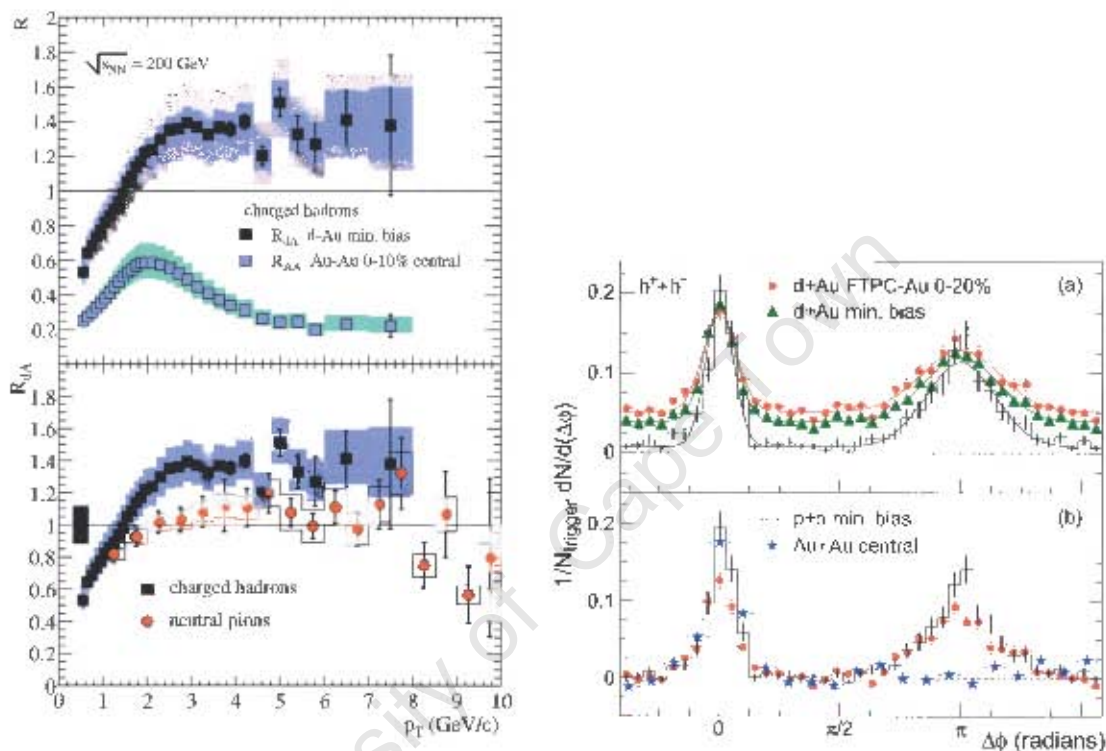
High p_T suppression results

The suppression factor of about 5 in the case of pions was attributed to parton energy loss in a deconfined medium created in the exit channel. It was argued that initial hard-scattered partons were travelling in a dense coloured field and were thus induced to radiate away their energy before fragmentation, which would clearly deplete the high p_T hadron yield. This scenario was widely accepted and proposed as direct evidence of the creation of a deconfined medium, identified as the QGP during the first two years of RHIC's running. There were, however, competing views [56] on the origin and cause of the high p_T suppression. Contrary to the interpretation of the creation of a deconfined medium absorbing the jets, it was proposed that the initial gluon density reached saturation, resulting in a lack of scattering centres for initial hard scattering, which could also deplete the high p_T yield.

For this reason, a control experiment was planned, which would collide deuterons on gold nuclei. In contrast to the hot and dense system created when gold nuclei collide, the lighter d-Au system could not create final state hot enough to melt the hadrons into a QGP. Since the initial state wavefunctions of the gold were the same in both Au-Au and d-Au collisions, any difference in the results between the two experiments would be due to the difference in the final state.

It was with great excitement that the results of the d-Au run from all four RHIC

experiments were announced in June 2003, which showed no suppression of high p_T particles, as in the Au-Au case [57, 58, 59, 60]. Also, there was a reappearance of back-to-back jets in the d-Au data. Jets of hadrons are formed in Au-Au collisions due to hard parton scattering, but while the jet near the surface of the plasma escapes, the far-side jet has to travel through a deconfined medium. The induced gluon bremsstrahlung dilutes the jet energy and no jet is observed on the far side. The effects of jet suppression, observed as the complementary phenomena of high- p_T suppression of the nuclear modification factor R_{AA} and the disappearance of back-to-back jets, are shown in figure 2.4.4.



(a) The comparison of the nuclear modification factor R_{AA} for central Au-Au and minimum-bias d-Au data taken with the PHENIX detector [58]. The suppression of the ratio at high p_T is indicative of the phenomenon of jet quenching and is clearly visible in the Au-Au data. This large suppression is not present in the d-Au data, suggesting that a QGP is not formed in that system.

(b) The observed angular correlation of high- p_T hadrons produced in d-Au collisions (top) [59], and central Au-Au and p-p (bottom). Note that the “away-side” peak at 180° is present in the p-p data, slightly broadened in the d-Au data, but is not seen in the Au-Au data, suggesting a suppression of the away-side jet.

Figure 2.20: Evidence of a deconfined medium created in Au-Au collisions at $\sqrt{s} = 200$ GeV, from the RHIC experiments PHENIX [58] and Solenoidal Tracker At RHIC (STAR) [59].

The azimuthal correlation function is defined as

$$D(\Delta\phi) = \frac{1}{N_{trigger}} \frac{1}{\varepsilon} \frac{dN}{d(\Delta\phi)}, \quad (2.26)$$

where $N_{trigger}$ is the number of hadrons with transverse momentum $4\text{GeV} < p_T^{trig} < 6\text{GeV}$ triggered on and ε is the tracking efficiency of the detector. The distribution is built from each high p_T trigger particle with the associated particles in the event with $2\text{GeV} < p_T < p_T^{trig}$. The near-side peak around $\phi = 0$ which is typical of jet structure. In the p-p and d-Au data, an away-side jet appears, which is also typical of the back-to-back production of high- p_T jets. In the central Au-Au data, however, the away-side peak is not seen, which reflects the absorption of the away-side jet.

Further investigation of the away-side jet

It was seen that the away-side jet disappears almost entirely in central Au-Au collisions. If the away-side jet really did disappear, then either this was due to a selection bias, by setting the near-side trigger threshold too high, or since energy is conserved in the di-jet creation process, the energy of the away-side jet is transferred to the medium. Also, initially, the reaction plane was not taken into account. Since the jets are found in azimuthal correlations, if there is a non-zero azimuthal correlation from the transverse flow component v_2 , this could influence the observation of the away-side jet. Further analysis of the RHIC data aimed at resolving the questions of how the away-side jet loses its energy, how that energy is redistributed, how the medium responds to the jet, and how these effects change with centrality and trigger-particle p_T .

A modification of the correlation function described in 2.26 takes into account these effects of flow in the background:

$$C(\Delta\phi) = \text{Jet}(\Delta\phi) + \text{Bknd}(\Delta\phi) \quad (2.27)$$

which is specifically written as

$$C(\Delta\phi) = D(\Delta\phi) + \xi(1 + 2v_2^t v_2^a \cos 2\Delta\phi) \quad (2.28)$$

where $D(\Delta\phi)$ is the correlation function defined in 2.26, ξ normalises to the level of flow in the background distribution [63], and t and a refer to the triggered and associated particles respectively. Figure 2.21 shows the azimuthal correlation function for pairs of charged hadrons with a trigger particle of $2.5\text{GeV} < p_T^t < 4.0\text{GeV}$ and associated particle of $1.0\text{GeV} < p_T^a < 2.5\text{GeV}$, as measured by the PHENIX collaboration [63] in Au-Au collisions at $\sqrt{s} = 200\text{ GeV}$. Another aspect which should be taken account of is the effect of the reaction plane, which has the effect of modifying the v_2 of the triggered particle [64], notably that v_2 of the triggered particle modified when it is calculated with respect to the reaction plane. While there is not yet a clear, consistent picture of how the away-side jet is modified by the medium (and conversely, how the medium reacts to the influence of the away-side jet), some effects are identifiable and are currently under study. Among these is the splitting of the away-side peak in the $\Delta\phi$ as a function of centrality (see figure 2.22). Models which implement jet broadening cannot describe the dip or the flat jet shape [65], while models which invoke Cerenkov gluons or medium dragging effects from flow do not describe the data much better [66]. On the other hand, Casalderry *et al.* proposed a ‘‘Mach-cone’’ or ‘‘shock-wave’’ mechanism to explain the away-side jet properties [70] which seems to hold much promise.




Figure 2.21: Correlation function of pairs of charged hadrons as measured by the PHENIX Collaboration in Au-Au collisions at $\sqrt{s} = 200$ GeV for different centralities. The solid bands indicate the estimated background contribution.

2.4.5 Heavy quarkonia suppression

Matsui and Satz are traditionally credited with predicting the suppression of heavy quarkonia [79] - the J/ψ ($c\bar{c}$) and Υ ($b\bar{b}$) families - as a signal of the formation of the QGP. Owing to their large mass, these mesons are created primarily at the earliest stages of the evolution of the system, so they would be suitable to probe the earliest state. These mesons also have a significant decay branching ratio in the dilepton, specifically dimuon channel (see table 2.1). Decay muons are only very slightly affected by subsequent interactions in the hadronic medium which the quarkonia decay in and carry the information about the state of the system (temperature, energy density) to the detectors. It was also thought that “There appears to be no mechanism for J/ψ suppression in nuclear collisions, except the formation of a deconfining plasma, and if such a plasma is produced, there seems to be no way to avoid J/ψ suppression” [79]. If this statement held strictly true, then any deviation of the yield of J/ψ from the expected value could be attributed to a deconfining plasma. As we shall see, the suppression of heavy quarkonia can also be attributed to purely hadronic processes. Before detailing these, a brief description of the formation process is in order.

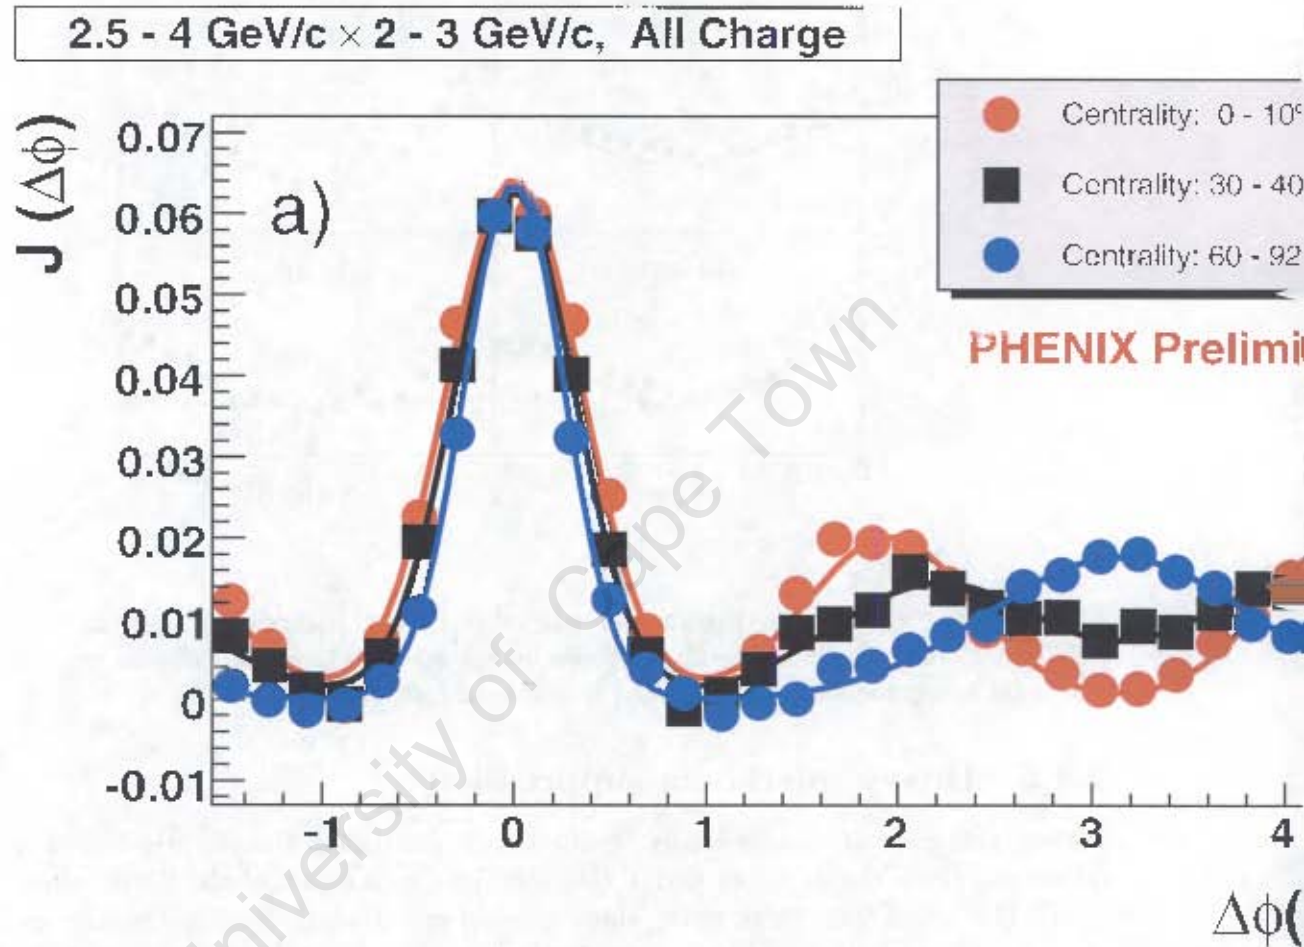


Figure 2.22: Correlation function with the underlying event subtracted for central (red), mid-central (black) and peripheral events (blue), showing the splitting of the away-side peak in central Au-Au collisions at $\sqrt{s} = 200$ GeV.

Resonance	$J^G (J^{PC})$	Mass (MeV/c^2)	Full Width (keV)	BR ($\mu^+\mu^-$) (%)
J/Ψ	$0^- (1^{--})$	3096.916 ± 0.011	93.4 ± 2.1	5.93 ± 0.06
ψ'	$0^- (1^{--})$	3686.093 ± 0.0034	337 ± 13	7.3 ± 0.8
Υ	$0^- (1^{--})$	9460.30 ± 0.26	(54.02 ± 1.25)	2.48 ± 0.05
Υ'	$0^- (1^{--})$	10023.26 ± 0.31	(31.98 ± 2.63)	1.93 ± 0.17
Υ''	$0^- (1^{--})$	10355.2 ± 0.5	(20.32 ± 1.85)	0.443 ± 0.008

Table 2.1: Zero-temperature properties of the lowest-lying heavy quarkonia spectrum ($b\bar{b}$ and $c\bar{c}$) [80].

Heavy quarkonia formation

Before a discussion of quarkonia suppression mechanisms, it is appropriate to outline the basic properties of heavy quarkonia and their dominant formation and decay mechanisms. The $c\bar{c}$ and $b\bar{b}$ vector meson families are produced not only in their ground states of $J/\psi(1S)$ and $\Upsilon(1S)$ respectively, but also in states of higher excitation, like $\psi'(2S)$, $\chi_c(1P)$ and similarly for the $b\bar{b}$ mesons. The properties of these lowest-lying $c\bar{c}$ and $b\bar{b}$ mesons are given in table 2.1. Note that the production mechanisms of closed charm and beauty mesons are qualitatively similar and so the discussion here will apply implicitly to Υ mesons, even if the present discussion may focus at points on only J/ψ 's explicitly. The formation mechanisms in p-p, p-A and A-A collisions will be discussed.

Formation of $b\bar{b}$ and $c\bar{c}$ pairs.

The formation of heavy quarkonia pairs involves an interplay between perturbative and non-perturbative QCD. Since the charm quark mass m_c sets a scale which is large enough to make the QCD coupling constant $\alpha_s(m_c^2)$ sufficiently small, perturbative QCD may be used to describe the parton-level production of a $c\bar{c}$ pair. Thereafter, the formation of a bound state proceeds non-perturbatively and these effects have to be taken into account. The pQCD-based Colour Singlet Model (CSM) [104] underestimated the high p_T yield of charmonium by over an order of magnitude at the Tevatron [105]. These data can be described by the Colour Octet Model (COM) [106] and the Colour Evaporation Model (CEM) [107] however.

The production of quarkonia is best described with a mix of perturbative and non-perturbative approaches, separating out the hard and soft parts of the interaction, based on the QCD factorisation theorem. The leading-order QCD processes for $c\bar{c}$ production are $gg \rightarrow c\bar{c}$ and $q\bar{q} \rightarrow c\bar{c}$, shown in figure 2.23. However, heavy quarkonia production is known to have significant contributions from Next-to-leading order (NLO) diagrams [109]. These contributions, such as virtual corrections to the leading order processes and soft collinear gluon emissions, increase the total cross-section, but can be approximately described as an overall multiplicative factor increase to the LO cross-section. The formation of $c\bar{c}$ pairs may be described well by NLO pQCD, but the formation of bound states of charmonium occurs through processes with small q^2 ,




Figure 2.23: Feynmann diagrams of leading-order (LO) processes for $c\bar{c}$ ($b\bar{b}$) production.

which makes the perturbative approach less applicable.

The Colour Evaporation Model (CEM)

The process of gluon fusion primarily forms a $c\bar{c}$ pair in a colour-octet state, denoted as $(c\bar{c})_8$. In order to produce a real J/ψ singlet state $(c\bar{c})_1$, the system has to neutralise its colour. This process is well described in a simple model called the Colour Evaporation Model (CEM), where the colour is evaporated by soft gluon exchange with the surrounding colour field. The transition from $(c\bar{c})_8$ to $(c\bar{c})_1$ involves an intermediate pre-resonance state of $(c\bar{c}g)_1$.

Quarkonia suppression in the plasma - the screening mechanism

According to Satz and Karsch [81], from the properties of finite temperature quarkonium, the inter-quark potential for non relativistically bound quarkonium at zero temperature is usually given in the Cornell form :

$$V(r, 0) = \sigma r - \frac{\alpha}{r}, \quad (2.29)$$

where r is the inter-quark separation and the parameters are $\sigma = 0.192 \text{ GeV}^2$, $\alpha = 0.471$, $m_c = 1.32 \text{ GeV}$ and $m_b = 4.746 \text{ GeV}$. This potential is modified at higher temperatures due to colour screening, and becomes

$$V(r, T) = \frac{\sigma}{\mu(T)} (1 - e^{-\mu(T)r}) - \frac{\alpha}{r} e^{-\mu(T)r}, \quad (2.30)$$

where the screening mass $\mu(T)$ is an increasing function of temperature. At finite temperature, when $r \rightarrow 0$ the $\frac{1}{r}$ behaviour is dominant, while as $r \rightarrow \infty$, the range of the potential decreases with $\mu(T)$, which makes the binding less effective with temperature. Semi-classically, the energy is given by

$$E(r, T) = 2m_Q + \frac{c}{m_Q r^2} + V(r, T), \quad (2.31)$$

where $\langle p^2 \rangle \langle r^2 \rangle = c = 1$. Minimising $E(r, T)$ with respect to r gives the radius of bound state at each T . It is noted that for a screening mass $\mu(T)$ above a certain critical value $\mu_D(T_D)$, there is no longer a minimum and bound states cannot form. Note also that μ_D is independent of T_D , but T_D depends on the functional form of μ_D . The screening mass is smaller for higher mass resonances, which means that

higher resonances are “melted” before the more tightly bound lower mass ones. One therefore expects the sequential suppression pattern of $b\bar{b}$ and $c\bar{c}$ families, with *e.g.* the Υ'' melting before the Υ' before the Υ . Taking into account feed-down from the decay of higher-lying resonances, the suppression patterns calculated by Digal, Petreczky and Satz [111] are shown in figure 2.24.

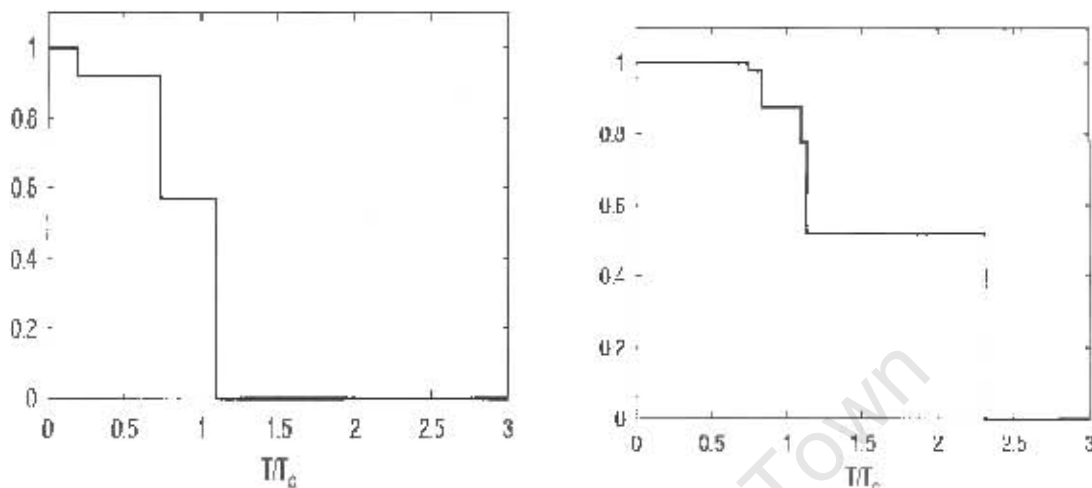


Figure 2.24: The J/ψ (left) and Υ (right) suppression patterns as a function of temperature relative to the critical temperature of the plasma from [111].

2.4.6 Non-plasma suppression mechanisms

As noted above, it has been shown that the yield of quarkonia may also be suppressed by purely hadronic means. It is essential to be able to have a reliable handle on these mechanisms, in order to describe properly the extent of plasma suppression (for an overview of J/ψ suppression mechanisms, see for example reference [94]). Two conspicuous mechanisms exist for the suppression of quarkonia which will be described below. The first, nuclear absorption, is a final-state effect, while the second, “shadowing” is an initial-state effect.

Nuclear absorption

According to the “binary-scaling” assumption of hard processes described above, the cross-section of J/ψ production in nucleus-nucleus collisions should scale as

$$\sigma_{J/\psi}^{AB} = AB\sigma_{J/\psi}^{pp}, \quad (2.32)$$

where A and B represent the mass number of the projectile and target respectively and σ_{pp} is the cross-section of quarkonia formation in p-p collisions. This simple parametrisation was modified somewhat in the light of J/ψ production data from

p-A collisions at the SPS [83, 85] and Fermilab National Accelerator Laboratory (FNAL) [93, 88] have shown that this scaling relation is modified by a factor α :

$$\sigma_{pA} = \sigma_{pp} A^\alpha, \quad (2.33)$$

where α takes all of the nuclear effects into account. This nuclear dependence is illustrated for the WA38 data in figure 2.4.6 and typical values of α at the SPS are around 0.92 [85].

This suppressed yield of quarkonia production relative to expected binary-scaling predictions is found to be similar for the J/ψ as for the ψ' . This discovery lead to the conclusion that it arises from pre-resonance absorption of the $c\bar{c}$ state in the surrounding nuclear matter created during the collision [83, 84, 86], since the resonance and ground-state of the $c\bar{c}$ pair have significantly different radii and would thus be expected to have different hadronic absorption cross-sections.

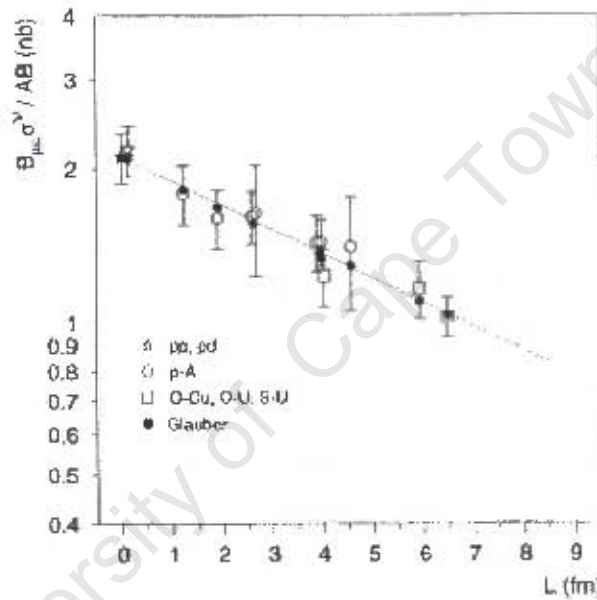


Figure 2.25: J/ψ cross-sections per nucleon, times dimuon branching ratio, as a function of nuclear size taken from [85]. The line corresponds to the best fit to the function $B_{\mu\mu} = \sigma_0 (A \times B)^{\alpha-1}$.

Co-mover absorption

Due to the very large number of produced particles in heavy-ion collisions, it is not unreasonable to expect that the J/ψ mesons formed during the evolution of the produced state suffer final-state interactions. These final-state interactions are often referred to as “co-mover” interactions, since the J/ψ is a tightly bound state and would have to interact with particles very near to it in phase space. This process was initially proposed by Brodsky and Mueller [89] as an effect at the partonic level on the production of J/ψ particles. Kharzeev *et al* described a model [90] which expresses

the co-mover suppression of J/ψ as a function of the number of wounded nucleons, while Armesto, Capella and Ferreiro use an expression for the co-mover suppression based on the Dual-Parton Model (DPM) model [91]. This description of the co-mover suppression seems to describe most the NA50 data very well, at 158 GeV [92]. Figure




Figure 2.26: Ratio of J/ψ to Drell-Yan production as a function of E_T measured by NA50. The dashed line is the normal nuclear absorption model [112]. The solid line is a co-mover absorption calculation [92]

2.26 shows the NA50 data for the ratio of measured to expected⁵ is shown along with a Glauber nuclear absorption curve, as well as the co-mover absorption described in [91], using a co-mover absorption cross-section of 1 mb. The co-mover absorption describes the peripheral bins well, but under-predicts the suppression at high- E_T .

Shadowing

Apart from the final-state nuclear effect of hadronic absorption, there exists another important modification of the yield of quarkonia from non-plasma mechanisms. This is an initial-state effect, due to the modification of the gluon structure function at high density. Shown in figure 2.27, is the nuclear parton distribution function (nPDF) calculated by Eskola *et al* [95, 96, 97] for different values of four-momentum exchange (Q^2). This function shows the ratio R_G^A as a function of x and Q^2 . Here, x is the fraction of the total nucleon momentum carried by the gluon and Q^2 is the transferred four-momentum. The ratio R_G^A is defined as

$$R_G^A(x, Q^2) = \frac{f_g^A(x, Q^2)}{A \cdot f_g(x, Q^2)}, \quad (2.34)$$

⁵Proportional to binary-scaling Drell-Yan production, as in figure 2.28

where $f_g(x, x^2)$ is the gluon distribution function within the nucleon and $f_g^A(x, Q^2)$ is the gluon distribution function within the nucleus. From the distributions shown in figure 2.27, it is observed that there are regions, at high- x , where the ratio is greater than unity and regions, at low- x , where it is less. These regions are referred to respectively as “anti-shadowing” and “shadowing”. Since one of the most important production mechanisms of quarkonia is gluon fusion, if the number of gluons is decreased in a given range, the yield of quarkonia will correspondingly be suppressed. This mechanism is an initial-state effect, as opposed to the final-state effects previously described, but contributes nonetheless to the overall quarkonia suppression pattern.




Figure 2.27: Calculation of the nuclear modification of the gluon density due to the effect of “shadowing”, taken from [95]

2.4.7 J/ψ suppression from SPS and RHIC

The results from the SPS experiment NA50 on J/ψ production as a function of energy density (which is related to centrality as discussed above) showed in 2000 that there is a threshold effect. The ratio of the J/ψ yield (measured in the dimuon decay channel [112]) compared to the yield expected from a fit to “normal” nuclear suppression described above is shown figure 2.28. It is clear that in the most central Pb-Pb collisions (corresponding to the highest energy-density), there is an “anomalous” suppression beyond what is expected by normal nuclear suppression.

The PHENIX collaboration has since made several measurements of the J/ψ distributions as a function of y , p_T and centrality, [113, 115, 114, 116]. It can be seen from figure 2.29 that there is significant suppression in the yield of J/ψ mesons in Au-Au collisions at $\sqrt{s} = 200$ GeV as a function of centrality. The yield is also more suppressed at mid rapidities than at forward rapidities [116] - a trend which is




Figure 2.28: Comparison of measured suppression of J/ψ to that expected from normal nuclear suppression, showing a threshold of “anomalous” suppression near $2.2 \text{ GeV}/\text{fm}^3$.

contrary to predictions using regeneration models [117] and co-mover models [118]. It is still not yet clear to what extent regeneration and normal nuclear suppression are present in the RHIC data. Recent In-In data from NA60 may help to clarify the situation, but it is thought that regeneration would play a more significant role in Pb-Pb collisions at the LHC due to the higher abundance of charm and beauty quarks.

2.5 The challenge for ALICE

It has been shown that the existence of a quark-gluon plasma can never be detected directly and experiments have to rely on indirect information and theoretical signatures in order to make statements on the presence of this state. Of the experimental signatures discussed above, the most relevant for this study is the suppression of quarkonia, since the ALICE dimuon spectrometer is specifically designed to measure the spectra of muons with high accuracy and acceptance.

The ALICE experimental setup will be discussed in the next chapter, with emphasis on the dimuon spectrometer. The challenge for ALICE will be to separate out hadronic from plasma effects in order to make reliable statements on the existence and nature of the quark-gluon plasma if it is indeed formed in collisions at the LHC. In order to achieve this goal, the experiment has been designed to measure several observables simultaneously, with an array of dedicated systems. These will be described in detail in the next chapter.

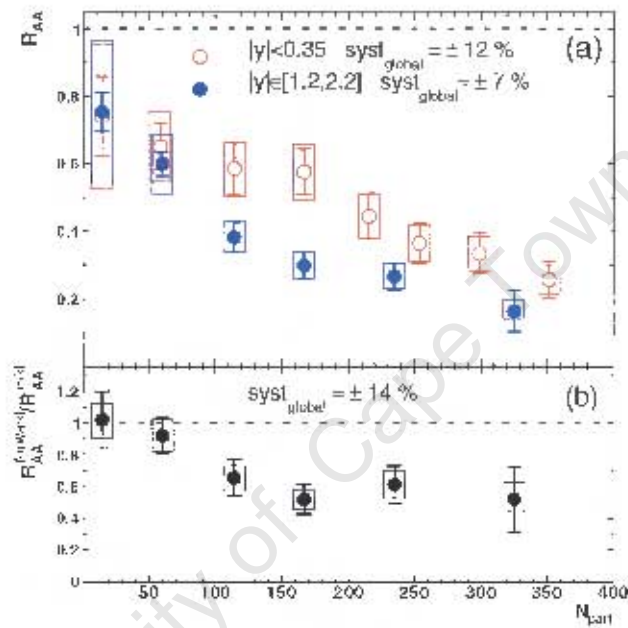


Figure 2.29: R_{AA} of J/ψ mesons measured in Au-Au collisions at $\sqrt{s} = 200$ GeV as a function of centrality (N_{part}) [116]. The top pane shows R_{AA} at mid- and forward rapidity, while the bottom frame shows the ratio between R_{AA} measured at forward rapidity, to that measured at mid-rapidity.

Chapter 3

The ALICE experiment at CERN's Large Hadronic Collider (LHC)

ALICE is one of four major experimental installations - and the only experiment dedicated to heavy-ion collisions at the LHC. The Large Hadronic Collider (LHC) will be able to collide a variety of symmetric and asymmetric systems, ranging from light systems (p-p, Ar-Ar and p-Pb), to very heavy systems like Pb-Pb. In this chapter, a brief description of the collider and the components of the ALICE experiment will be given.

3.1 The Large Hadron Collider

The LHC will be the largest accelerator in the world to collide heavy ions and is being built in the same tunnel as the 26.659 km-circumference Large Electron-Positron (LEP) ring. As such, the physics regime opened up by this machine will represent experimentally uncharted territory to the theory of QCD. The nominal operating parameters of the accelerator are given in table 3.1.

Apart from running Pb-Pb to create the largest, densest, hottest system possible, control experiments of lighter ions are scheduled to investigate the system size dependence of the results. These include :

- Intermediate mass ions, such as Ar-Ar. Depending on the results of the Argon runs, the LHC can accelerate symmetric beams of other intermediate mass ions,

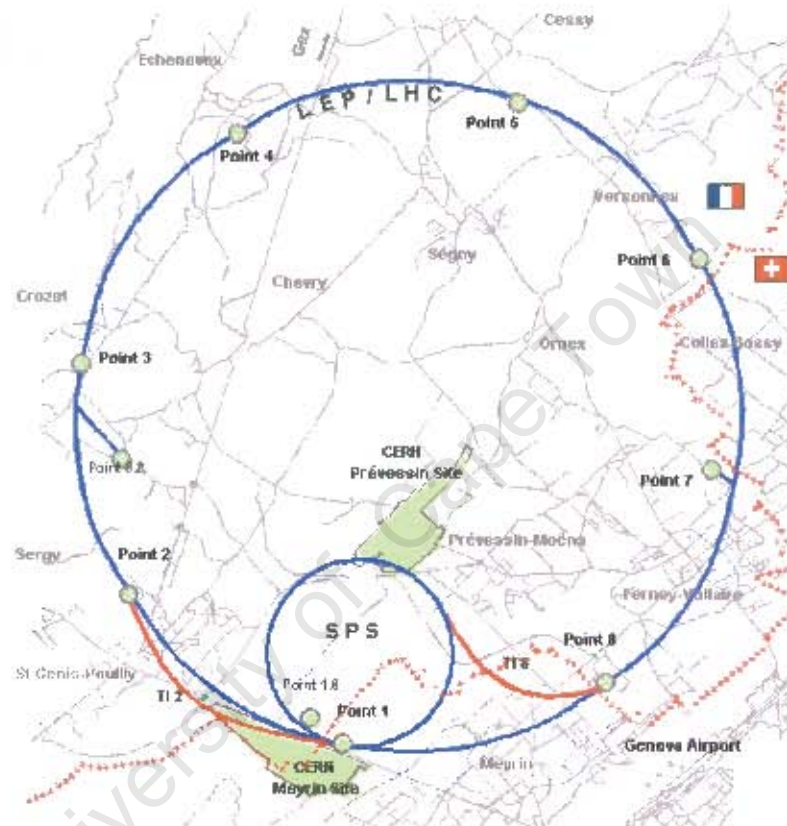
	pp	Pb-Pb
Energy per nucleon (TeV)	7	2.76
Bunch length (cm)	7.55	7.94
Beam radius (μm)	71	15.9
Ions per bunch	$1.1 \cdot 10^{11}$	$7.00 \cdot 10^7$
Luminosity ($\text{cm}^{-2}\text{s}^{-1}$)	$< 3 \cdot 10^{30}$	10^{27}
Time between bunches (ns)	24.95	99.8

Table 3.1: Principle parameters of the Large Hadronic Collider at CERN [120], [141].

such as Sn, Kr, N and O.

- Elementary nucleon-nucleon collisions. These include pp, p-A and other “p-like” systems, *e.g.* d-d or α - α . The data from these experiments is necessary to provide the relevant reference information to the heavier systems.

The LHC will have several interaction points - regions where the counter-moving beams of nuclei overlap - where the large experimental halls will house the four experiments : CMS, ATLAS, LHCb and ALICE. A schematic diagram of the layout of the accelerator complex can be seen in figure 3.1, with ALICE at point 2.



Map of CERN sites and LHC access points

Figure 3.1: The LHC accelerator complex at CERN, showing the location of the four experimental halls, including the ALICE experiment (Point 2).

The first three experiments in this list will be focused on probing signals predicted by the Standard Model, while the last is designed with a very different goal in mind, the search for and exploration of the QGP.

3.2 The ALICE experiment overview

The ALICE experiment will be placed in the cavern previously occupied by the L3 experiment at LEP [121], at interaction point 2. It consists of a wide variety of detector subsystems, which will allow it to study many different observables simultaneously. It will have excellent tracking and particle identification in the mid rapidity region, and good coverage in y and p_T for specific observables, like dimuons and photons. The ALICE detector consists of a central part containing most of the subsystems, a dimuon arm at backward rapidity in the LHC co-ordinate frame, and several subsystems for event characterisation and triggering at forward rapidities. A schematic diagram of the overall ALICE experimental apparatus is shown in figure 3.2, where the detector components are indicated. A description of the detector subsystems is given below, and will be grouped into various sections - the central systems, the forward detectors and the dimuon spectrometer which the nature of this thesis requires a special focus on. In order to provide the experimental context to these systems, a short description of the other important sub-detectors and systems will be given too, including the trigger and DAQ systems and the computing framework.




Figure 3.2: Schematic diagram of the ALICE detector, to scale, showing various subsystems.

3.3 The central (barrel) ALICE systems

The central “barrel” of ALICE consists of a dense variety of integrated detector types around the interaction vertex, dominated by a TPC. Complementing the TPC is an ITS based on 3 different silicon detector technologies. The TPC and ITS are

Layer (Detector Type)	Single Track		Double Track	
	$\sigma_{r-\phi}(\mu\text{m})$	$\sigma_z(\mu\text{m})$	$\sigma_{r-\phi}(\mu\text{m})$	$\sigma_z(\mu\text{m})$
1,2 SPD	12	100	100	850
3,4 SDD	38	28	200	600
5,6 SSD	20	830	300	2400

Table 3.2: Resolution parameters for the Inner Tracking System for single and double tracks

situated inside the large L3 magnet, which provides the solenoidal magnetic field required to provide momentum information on the tracked particles. These other detectors help in measuring specific signals, or with overall particle identification (PID), event characterisation or triggering. They include a Transition Radiation Detector (TRD), Time of Flight (TOF), High Momentum Particle Identification (HMPID) and PHOTon Spectrometer (PHOS) detectors and will be described briefly below.

3.3.1 The L3 magnet

All of the central detectors described below are situated within the volume of the large central magnet, last used for the L3 experiment at LEP. This magnet has a 5m radius and is 12 metres long, providing a maximal solenoidal field for use in ALICE of 0.5 T. In this magnetic field charged particles are deviated by the Lorentz force, enabling the central tracking systems to perform tracking down to very low transverse momentum, about $p_T > 100$ MeV/c.

3.3.2 The Inner Tracking Chamber (ITS)

The ITS [122] has six detection layers covering the rapidity region of $|\eta| < 0.9$, with high spatial resolution used mainly for tracking and decay vertex determination. The radius, technology and segmentation of the detector are optimised to satisfy certain constraints, including physics observables and efficiency. The components of the ITS, seen in figure 3.3 are pixel, drift and strip detectors. The Silicon Pixel Detector (SPD) provides accurate vertex information, while the Silicon Drift Detector (SDD) provides analog dE/dx for PID purposes. The Silicon Strip Detector (SSD) provides a track-matching capability which the TPC uses to improve its tracking.

With this setup, the ITS is capable of determining a very precise event vertex (see table 3.2), as well as secondary vertices from heavy flavour decay. These include heavy flavour (charm and beauty) mesons, as well as strange baryons (Λ , Ξ , Ω) and mesons (K_S^0 , ϕ , etc). It has increased acceptance compared to the TPC for low momentum ($p_T < 100$ MeV) tracks and increases the resolution on high p_T tracks with the addition of the TPC.




Figure 3.3: The Inner Tracking System, showing the various component layers of silicon

3.3.3 The Time Projection Chamber (TPC)

The TPC [123] is the heart of the inner detectors, providing track identification, momentum and energy loss information for most of the charged particles created in the nuclear collision. It has a cylindrical shape (see figure 3.4), with an inner radius of about 85 cm, an outer radius of about 250 cm and a length of about 500 cm. The inner radius was chosen to give an average multiplicity of 0.1-0.2 particles. cm^{-2} (in the transverse plane) and the outer radius was chosen to obtain a dE/dx resolution for electrons of less than 10 %. This gives the TPC an acceptance of about $|\eta| < 0.9$ in pseudorapidity and $100 \text{ MeV}/c < p_T < 100 \text{ GeV}/c$ in transverse momentum for charged tracks. To keep the occupancy manageable, the detector is segmented into about 560 000 readout pads of three different sizes, depending on the radial coordinate on the endcaps.

3.3.4 The Transition Radiation Detector (TRD)

The Transition Radiation Detector (TRD) [124], shown in figure 3.5, is situated radially between the TPC and the TOF (described below), in order to detect the transition radiation of charged particles as they traverse from one medium to the next. It allows very good $e - \pi$ separation for tracks with $p_T < 1 \text{ GeV}/c$ (where the TPC $e - \pi$ rejection efficiency drops) and will accept almost all electrons in the range $|\eta| < 0.9$. The TRD is also a very fast detector, and as such can be used as a trigger for high p_T leptons. Such a trigger will considerably increase the heavy flavour (Υ and J/ψ families) yield in the high-mass region of the dilepton spectrum.




Figure 3.4: Schematic scale diagram of the ALICE Time Projection Chamber, showing support spokes, and inner and outer sections.

The TRD is a segmented detector consisting of 540 sub-detectors and has a total sensitive area of 750 mm^2 .

3.3.5 The Time of Flight (TOF) detector

The time of flight detector [125] is used for particle identification and covers the full azimuth in the pseudorapidity region of $|\eta| < 0.9$. In conjunction with the ITS and TPC, it can be used on an event-by-event basis to separate out pions, kaons and protons with $0.2 \text{ GeV}/c < p_T < 2.5 \text{ GeV}/c$. The cylindrical device is situated radially just beyond the TRD. The dimensions of the modules match those of the TRD and TPC, and are arranged in such a way that the joining areas of the modules are aligned with the dead areas of other detectors projected from the interaction point, which creates minimal dead zones (see figure 3.6).

3.3.6 The High Momentum Particle IDentification (HMPID) detector

As its name implies, the HMPID [126] is designed to identify very high momentum particles, albeit in a limited acceptance – only about 10 % of the central barrel acceptance. It is optimised to separate out π/K and K/p up to transverse momenta of about $3 \text{ GeV}/c$ and $5 \text{ GeV}/c$ respectively. The HMPID is a Ring-Imaging Čerenkov detector (RICH) design, mounted in an independent support cradle at a distance of




Figure 3.5: Schematic diagram of the ALICE Transition Radiation Detector (to scale)

about 4.8 m from the nominal vertex. The cradle is fixed in the space frame at about the two-o'clock position (see figure 3.6).

3.3.7 The PHOTon Spectrometer (PHOS)

The ALICE photon spectrometer [127] is a high-resolution electromagnetic spectrometer, which is designed to detect photons in a limited acceptance at mid-rapidity. It can provide neutral meson (*e.g.* π^0) and photon identification. As seen in figure 3.6, it is positioned at the bottom of the ALICE central barrel, 460 cm radially distant from the interaction vertex, covering the pseudorapidity region of $-0.12 \leq \eta \leq 0.12$ and 100° in azimuth, giving a total sensitive area of about 8 m^2 .

With the PHOS, ALICE will be able to study the QGP by measuring direct photons (as described in section 2.4.4), which make up about 10 - 15 % of the total photonic cross-section. In addition, it will be able to study jet quenching, by measuring the high momentum light meson yield through the dominant decay channel of $\pi^0 \rightarrow 2\gamma$ and $\eta \rightarrow 2\gamma$. It will also measure event-by-event fluctuations of the photon yield compared to the charged particle yield, which is predicted as a signal of the onset of chiral symmetry restoration.

3.4 The forward-rapidity ALICE detectors

There exist in ALICE a set of forward detectors for the determination of the overall underlying event centrality and multiplicity for global triggering and other collision

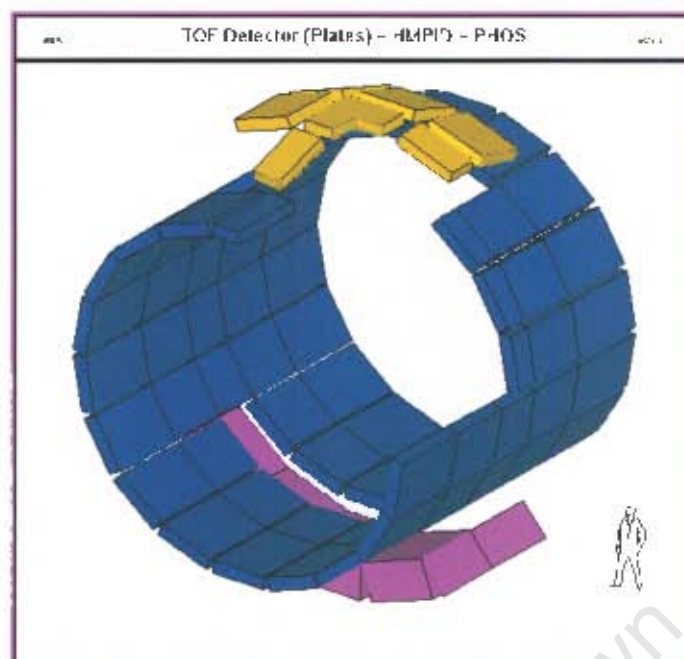


Figure 3.6: Schematic diagram of the TOF (central barrel), PHOS (lower) and HMPID (top) detectors in the ALICE experiment, to scale.

information¹

The Zero degree calorimeter (ZDC)

The spectator nucleons in the collision are detected by the ZDCs placed at 116 m from the interaction point, on both sides of the vertex. Since the ZDC's are so far along the tunnel, they cannot provide a L0 signal in the time required, but they use three levels of discrimination to provide an L1 trigger input for the different centrality classes.

3.4.1 The Photon Multiplicity Detector (PMD)

The PMD is situated 360 cm from the interaction point, in the forward rapidity region (opposite the dimuon spectrometer) and covers the full azimuth in the pseudorapidity region of $2.3 \leq \eta \leq 3.5$. It has an electronics system similar to that of the dimuon arm, and detects photons in the same phase-space as the dimuon spectrometer, making it useful for correlation studies in conjunction with the dimuon spectrometer. The PMD can measure event-by-event photon spatial distributions (in the $\eta - \phi$) plane. This is also possible with the dimuon spectrometer and allows a study of event-by-event fluctuations and flow. Apart from QGP signals, event characteristics such as the transverse electromagnetic energy and reaction plane are measurable with the

¹See section 2.4.1 for an explanation of how these detectors are used to estimate centrality.




Figure 3.7: The ALICE Zero degree calorimeter (ZDC), next to the beam pipe.

PMD, which make it invaluable for making flow and Hanbury-Brown Twiss (HBT) interferometry measurements.

3.4.2 The Forward Multiplicity Detector

This detector consists of 5 concentric silicon strip ring counters placed on both ends of the interaction region, covering the pseudorapidity region of $-3.4 < \eta < -1.7$ and $1.7 < \eta < 5.4$. It is used for measuring the multiplicity of charged particles, but its slow readout time does not allow it to participate in the L0 trigger, although it does have an input into the L2 signal. It is segmented into 5 annuli of silicon strip detectors (2 in the forward and 3 in the backward region of the vertex), with two different types (internal and external), which form 20 and 40 sectors in ϕ respectively. The FMD can be used to study signals which appear in multiplicity variations of charged particles, such as event-by-event fluctuations and flow, but is also useful in determining the centrality of the event. The ITS in conjunction with the FMD gives almost 4π coverage of charged particle production.

3.4.3 The T0 and V0 detectors

Timing is crucial for many detectors, which need a “start” signal. The T0 plays this important role, giving a precise measurement of the collision start time for the TOF and trigger electronics. It is made up of two arrays of 12 Čerenkov counters each, read out by fine mesh photomultiplier tubes. The arrays are placed on either end

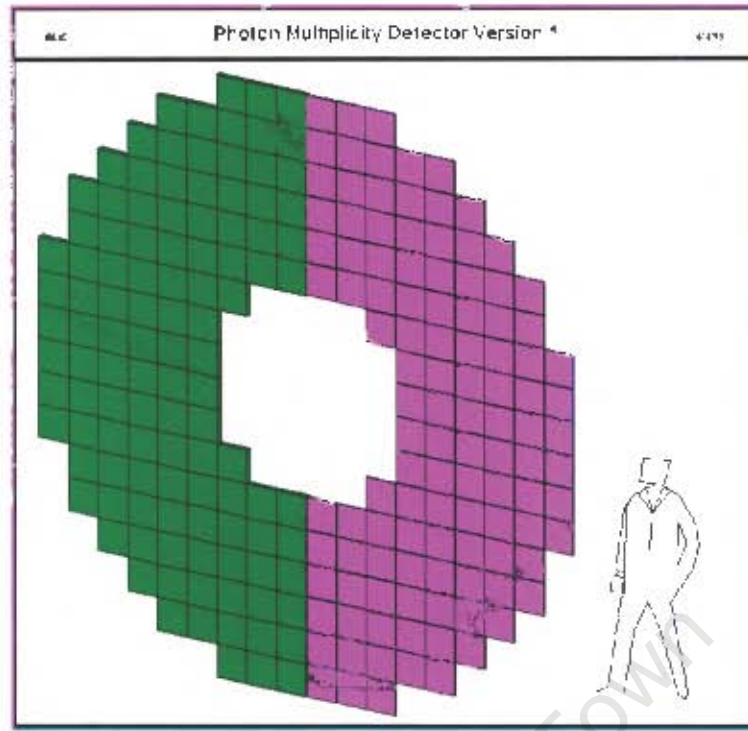


Figure 3.8: A schematic diagram of the ALICE Photon Multiplicity Detector (PMD) (to scale)

of the vertex region, at $z = -350$ cm ($T0_L$) and $z = 70$ cm ($T0_R$). The fast output signal is used as an input to the L0 electronics logic, and is used to trigger on events and reject beam-gas background. There is also an input into the pre-trigger from the co-incidence of the T0 and the TRD. The time resolution of the T0 is about 50 ps and its vertex resolution is of the order of cm.

The V0 detector is composed of two sets of radially and azimuthally segmented scintillators placed on either side of the interaction vertex, with the purpose of providing accurate vertex information. This vertex position information is used to filter out particles which do not come from the primary vertex, in order to improve the centrality estimation by providing a more reliable multiplicity measurement.

3.5 The dimuon spectrometer

The dimuon spectrometer [128] will form the focus of the description of the ALICE detector components, since the physics signals we are interested in will be measured with this detector. It provides a specific capability to study the muons from the decay of particles produced in heavy ion collisions.

The spectrometer is situated at backward rapidities in the polar range of $171^\circ - 178^\circ$ ($-4 \leq \eta \leq -2.5$) and covers the full azimuth. Using the ZDC, PMD and ITS in conjunction with the dimuon spectrometer, ALICE has the capability to directly compare the yield of quarkonia in different regions of phase space, and as a function

of centrality. The possibility also exists to study $e^{+/-} - \mu^{-/+}$ co-incidences with the TRD, in order to enhance an open heavy flavour measurement.

The basic design of the detector must fulfil the criteria imposed on it arising from the experimental environment. Essentially, this means taking account of the particle multiplicity, interaction rate and invariant mass resolution required to separate out various 'onium species². The most stringent requirements are set in the case of central Pb-Pb collisions, where the multiplicity is the highest. The maximum expected hit density in this environment decreases as a function of z from about $\frac{dN}{dA} \Big|_{max} = 4 \cdot 10^{-2} \text{ mm}^{-2}$ on the first station, to about $\frac{dN}{dA} \Big|_{max} = 1 \cdot 10^{-2} \text{ mm}^{-2}$ on station 3³.

The spatial resolution of hit reconstruction on the chambers has a direct influence on the momentum determination and hence the mass resolution. A spatial resolution of about $\sigma_{xy} < 100 \mu\text{m}$ is needed to achieve a mass resolution $100 \text{ MeV}/c^2$, which is in turn required in order to separate out the sub-states of heavy-quarkonia resonances. Other factors which would degrade the momentum resolution are multiple scattering and inefficiencies in the detector. For this reason, the thickness of the active surfaces of the tracking chambers are less than 3% of the radiation length χ_0 and the efficiency of the detectors, taking into account dead areas, is greater than 90%. A further criterion for the dimuon spectrometer is that it should have good acceptance at low p_T for muons.

The spectrometer consists of 5 sets of tracking stations, and 2 trigger stations. The detectors themselves are of a segmented cathode strip chamber design, which will be described in detail below, however it is worthwhile to note here that the the segmentation of the detector has been designed to take into account the expected multiplicity of particle hits on the detector planes. In the case of the trigger planes, Heavy-Ion Jet Interaction Generator (HIJING) simulations have shown that the maximum hit rate for central Pb-Pb is about $3 \text{ Hz}/\text{cm}^2$.

The detector system is roughly made up of 6 components:

- beam shield surrounding beam pipe
- front absorber
- tracking stations
- dipole magnet
- muon filter
- trigger stations

These components will be discussed in depth below.

²It should be noted as well that this invariant mass resolution requirement is a function of the invariant mass.

³These estimates include a safety factor of 2, to take account of uncertainties in the particle multiplicity and possible unknown sources of background.

3.5.1 The absorbers

Leptons, specifically muons, do not interact via the strong force and so are able to pass relatively unaffected through large pieces of dense material. In order to suppress the large number of unwanted particles (hadrons, photons, electrons *etc*) produced in the initial collision and detect only muons, there exist two absorbers in the dimuon spectrometer - the front absorber and the muon filter.

The front absorber is designed to reduce as much as possible the flux of hadrons before the spectrometer, while minimising the effects of energy loss and multiple scattering on muons, as well as reducing the number of hadrons and secondary muons from hadronic decay. The absorber, which has undergone several design changes [129], is a mixture of carbon, and lead, tungsten and polyethylene [130], combined in a way so as to optimise the absorber for the conditions above. It is 4.2 metres long, corresponding to 10 radiation lengths. Since the front absorber is partly within the TPC barrel, it is protected by a lead endcap in order to reduce the contribution of secondary particles scattering into the TPC sensitive volume.




Figure 3.9: Schematic diagram of the ALICE dimuon arm front absorber, to scale, showing the distribution of materials used.

The trigger chambers of the spectrometer have to be shielded from all particles but muons, so that the trigger rates are not contaminated by background due to neutrons and gammas. For this reason, a 1.2 m thick iron wall separates the trigger chambers from the rest of the spectrometer stations. This absorber plays the role of a muon filter, since muons of less than 4 GeV/c are stopped in it, and it only allows high momentum particles through to the trigger chambers.

3.5.2 The spectrometer tracking system

The tracking system is made up of 5 tracking stations, each consisting of 2 detection planes each consisting of two detection planes. Each detection plane consists of a set of multiwire proportional chambers (“quadrants” in the case of stations 1 and 2 and “slats” in the case of stations 3, 4 and 5. Both cathodes of the chamber are segmented

and read-out - the first cathode is finely segmented in the vertical (y) direction and is referred to as the “bending plane”, while the other is segmented with a slightly coarser grain in the horizontal (x) direction and is referred to as the “non-bending plane”. A dipole magnet is used for momentum determination, and the finer segmentation in the y -direction is designed to take advantage thereof. The dipole magnet is centred at a distance of 9.9 m from the interaction vertex, with an integrated field strength of 3 T.m, and a maximum field value of 0.7 T. It has a length of 5 m in the $-z$ direction and covers the polar angular acceptance of the rest of the spectrometer.




Figure 3.10: Schematic view of the ALICE dimuon spectrometer dipole magnet. The front (left) and side (right) views, as generated from aliroot, both show the third tracking station of the spectrometer within the volume of the dipole magnet.

Stations 1 and 2 are situated before the dipole magnet (in the z -axis), while stations 4 and 5 are situated beyond it; station 3 provides a measurement within the dipole magnet (see figure 3.10). Due to the high particle multiplicity, and requirements of good resolution and large sensitive areas by the spectrometer, segmented cathode pad chamber technology was decided upon for the tracking stations. The segmentation of the tracking chambers was designed so that the average occupancy would remain roughly constant as a function of radial distance from the beam pipe (see figure 3.11). Since the hit density decreases with radial distance from the beam pipe, this means that pad sizes are smaller nearer the beam line, and increase further out. The chambers of stations 1 and 2 require special segmentation, due to the very high multiplicity near the interaction vertex.

3.5.3 The dimuon trigger subsystem

The dimuon trigger system has to select dimuons of interest from a significant background. The trigger chambers provide the initial bias to select events with high

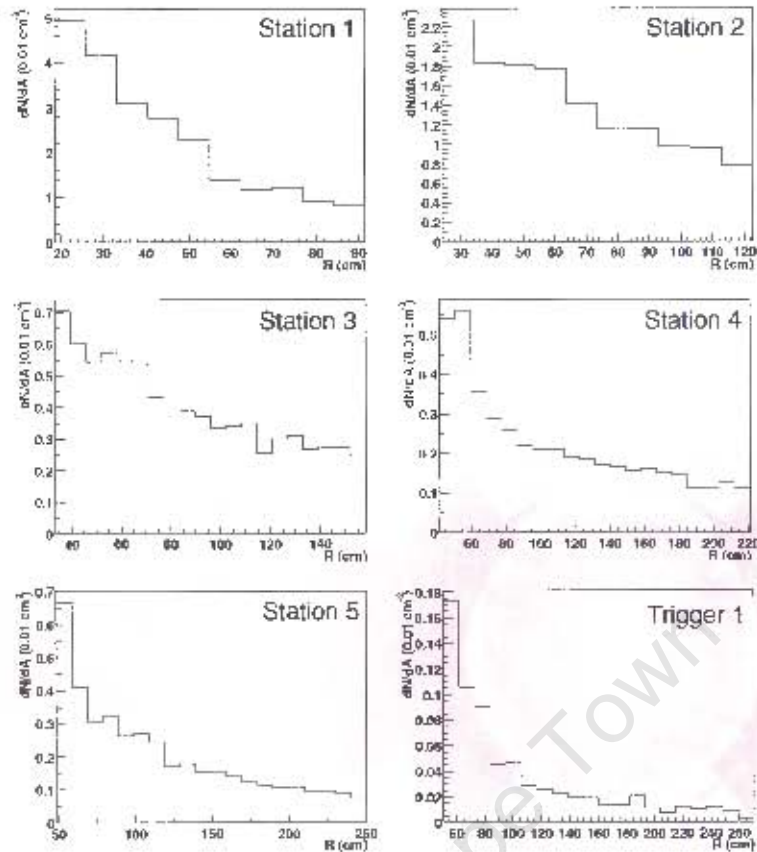


Figure 3.11: Occupancy of the tracking chambers as a function of radial distance in the case of central Pb-Pb collisions [128]

momentum muons and are specifically designed to minimise the impact of a soft background coming from low-energy muons, neutrons and gammas from secondary reactions and decays. In order to provide a rough p_T estimate, the trigger chambers have a granularity of about 1 cm. A very fast response time is required in order to participate in the global ALICE L1 trigger. Single-gap Resistive Plate Chamber (RPC) technology was thus selected to be used for the trigger stations.

3.5.4 Dimuon spectroeter read-out

The detector read-out and raw data format has implications for the dHLT and therefore relevant details of the dimuon tracking and trigger chambers read out will be briefly stated here.

The readout for the dimuon tracking stations involves a series of digital processing of the data which is done to a large extent by integrated electronics situated directly on the detector. Signals generated by the passage of charged particles through the sensitive areas of the detectors are read out by the MANAS Numérique (MANU) Front-End Electronics (FEE), which perform pre-amplification, shaping and filtering of the signal in the 1.1 million channels. An application-specific integrated circuit (ASIC), the MUON Arm Readout Card (MARC) handles the data sequencing as well




Figure 3.12: MUON spectrometer tracking and trigger chambers. Dimensions are in cm.

Station No.	1	2	3	4	5
Z position (mm)	5100	6860	9824	12900	14241
Bending plane	4x6	4x7.5	5x25	5x 25	5x25
Pad size (mm x mm)	4x 12 4x24	4x15 4x30	5x50 5x100	5x50 5x100	5x50 5x100
Max. Occupancy (0.01 cm^{-2}) (central Pb-Pb x2) (cm^{-2})	4.0	1.6	0.6	1.0	1.2

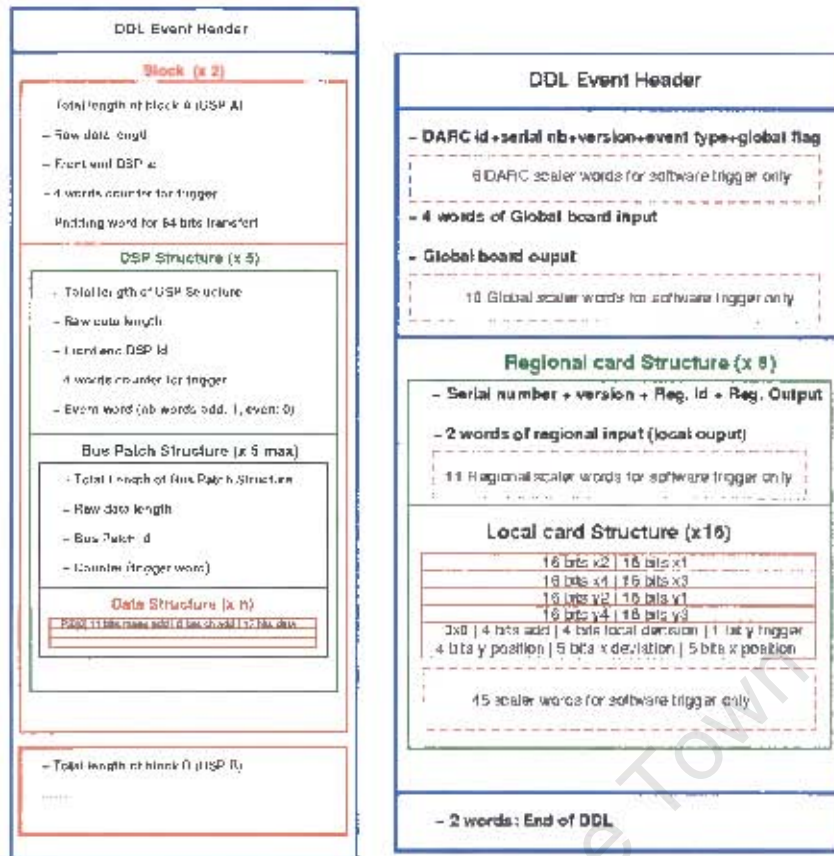
Table 3.3: Tracking chamber physical characteristics. The pad size of the tracking chambers is given for the two segmentation regions (inner and outer).

as the zero-suppression and communication with the Digital Signal Processor (DSP)s. These DSPs, which are grouped in clusters in specific Cluster Read Out Concentrator Unit (CROCUS), transmit the data via the Protocol for ALICE Tracking Chamber (PATCH) electronic buses. Once the data has been transferred from the FEEs to the CROCUS, it is shipped to the DAQ and dHLT compute farm via Detector Data Link (DDL) cables, which are standard to all ALICE subdetectors, along which raw data is transferred.

Four Multiplexed ANalog Signal processor (MANAS) chips, which provide 12 bits of Analog to Digital Converter (ADC) information after pre-amplification and analog-to-digital conversion, are grouped on a MANU board. The data is transferred to front-end CROCUS controllers via PATCH-busses on the detector; each of the 5 CROCUS controllers handles 10 PATCH-busses. A concentrator CROCUS reads out the 5 front-end CROCUS, at which point the data is formatted and sent to a Source-Interface Unit (SIU) of the DDL cable. There are two data blocks of 25 PATCH-busses per DDL event and one half-chamber is connected to one DDL, containing the merged data of 5 CROCUS, corresponding to at most 50 PATCH-busses [137].

The 21000 channels of the four trigger chambers are read out by 3 880 front end cards, which are 8 channel A Dual-Level Threshold (ADULT) chips [138]. The trigger read-out is hierarchical and is controlled by 234 local and 16 regional readout cards. More information on the detailed MUON arm trigger read out can be found in [139], but for the sake of clarity, the readout procedure as it pertains to the DDL-wise raw data will be outlined here. Once read out by these cards, the data is sent to the Dimuon Arm Readout Card (DARC), with one DARC serving 8 regional cards and hence 2 DDLs for the trigger chambers. A global card in the trigger crate gathers information from the regional cards and generates the trigger output [140], based on a lookup table.

The data from the trigger and selected tracking chambers is then used as the input to the dHLT tracking and trigger algorithms, which will be described in detail in chapter 5. Knowledge of the raw data format at this point is essential in order to unpack the raw data and process it.



(a) Dimuon spectrometer tracking chamber raw data format. (b) Dimuon spectrometer trigger chamber raw data format.

Figure 3.13: Raw data format used in offline simulations.

3.5.5 The dimuon spectrometer raw data format

An ALICE raw event consists of an event header (16 words), an equipment header (7 words) and the DDL sub-event. This DDL sub-event begins with an 8-word header (32 bits each) [131] and also contains the actual data payload. The raw data format, given in [153], is defined by the ALICE geometry, mapping and segmentation conventions, and is different for the tracking and trigger chambers. Note that although the geometry and segmentation also varies from chamber to chamber, this requires different mapping definitions, and not an altogether different data format.

Tracking chamber DDL raw data format

The tracking chamber raw data format contains two blocks of data, each of which contain the DSP information. The actual payload consists of data packed into a little-endian 32-bit word with the bits assigned as follows :

31 parity bit

29-30 Must Be Zero (MBZ)

18-28 MANU identification number (11 bit)

12-17 channel identification number in MANU card (16 bit)

0-11 digitised ADC reading of pad charge

Trigger chamber DDL raw data format

The trigger chambers are read out by the same ALICE-standard DDLs as the tracking chambers, but their data payload is formatted differently. There are two sets of data structures, depending on whether only the physics event is read out, or whether additional information is to be read out. The trigger raw data payload is a 32 bit word containing

- 9 bits zero
- 4 bits local decision:
 - 00 No trigger
 - 01 Negative particle trigger
 - 10 Positive particle trigger
 - 11 No deviation trigger (undefined)
- 1 bit y-trigger
- 4 bits y-position
- 5 bits for x-deviation (1 bit for trigger, 4 bits for deviation)
- 5 bits for x-position

The global ALICE High Level Trigger (HLT) is briefly described in chapter 4, to set the context for the dHLT.

3.5.6 Hit reconstruction (offline cluster finder)

The digitised data read out from the detector has to be reconstructed in order to provide positions of particle tracks in spectrometer. This is done by estimating the position a track hit on a spectrometer chamber by investigating the deposited charge distribution, or charge “cluster”. If there are no overlapping charge clusters, as is the case when the occupancy is low and the hits are far from each other, then this task is relatively simple. The charge distribution is fitted to a Mathieson-like function [132] to find the co-ordinates of the hit. In the case of overlapping hits, where one has two tracks contributing to a single charge cluster, a fit is performed with the same function, but with more parameters and the number of hits and their positions are determined from the local maxima. This resolution of this method, however, degrades badly in the presence of background, so a more robust method has been developed by A. Zinchenko [133].

This method is based on the Maximum Likelihood - Expectation Minimization method [135], also known as the Lucy-Richardson method [134]. This new method of reconstructing hits from the charge distributions provides much better position resolution in the presence of large background and therefore better mass resolution. Although the technique described here is very efficient at identifying clusters and reconstructing hits from them, it is however much too slow to be implemented in the time frame necessary for the dHLT, of the order of ms. The fast or “online” hit reconstruction procedure will be detailed more in chapter 5.

3.6 The ALICE trigger and DAQ system

3.6.1 General trigger design considerations

The ALICE trigger has evolved over many years, building on experience from NA57 and other experiments. Experimental running conditions, such as the multiplicity and collision rate, place severe constraints on the trigger design. ALICE has been designed primarily to cope with the most central Pb-Pb collisions, which are the most demanding in terms of produced particles. These are expected at a relatively low rate, about 10 kHz, but very high multiplicities of between 4000 and 8000 particles per unit of pseudorapidity. The considerations of the experimental conditions at ALICE have led to three general features of the ALICE trigger system

global trigger selections Boolean combinations of trigger signals are used to describe whole events, while more detailed treatment of the signals is left to the HLT, which has the capacity for more sophisticated data processing.

checking of time separation of events Also known as *past-future protection*, which rejects out-of-time events

Saturating trigger rates While the TPC is limited in its trigger rate by the drift time of the chambers, other detectors can be read out much faster. Advantage can be taken of their reduced readout time by organising them into “clusters”, which can be read out semi-independently, which would be useful to acquire a better data set. The dimuon spectrometer is one of these detectors.

The sub-detectors of ALICE have vastly different latencies (see table 3.6.1) – for example, the ITS SDD has a sensitive period of 260 μ s while the response of the T0 takes under 50 ns. In order to make the most efficient use of the detector, there is the possibility to read out only certain subsystems for any given event. The sub-detectors are organised into “clusters” and these clusters are read out for events where certain conditions have occurred. These conditions are known as “trigger classes” and are generally constrained by the latency of the detectors.

The trigger system is divided into the Central Trigger Processor (CTP) and several Local Trigger Unit (LTU)s. The detailed functioning of these is not relevant to this work, however certain aspects relate to the integration of the HLT; more detailed information can be found in the ALICE TDR [143].

Detectors relevant to MUON event	past (μs)	future (μs)
ITS Pixels	-0.1	0.1
PMD	-5.0	5.0
Dimuon tracker	-5.0	1.0
Dimuon trigger	-0.1	0.1
FMD	-2.0	2.0
T0, V0, ZDC	-0.025	0.025
Central Detectors		
TPC	-88.0	88.0
ITS Strips	-10.0	3.0
TRD	-10.0	10.0
TOF	-0.5	0.5
HMPID	-5.0	5.0
PHOS	-2.0	2.0

Table 3.4: Table of latencies for the forward subsystems (including the ITS Pixels and PMD, which are included in the MUON event) and central detectors [150].

3.6.2 The Central Trigger Processor

The Central Trigger Processor (CTP) plays the central role in the ALICE trigger. It has the responsibility of combining and synchronising various trigger signals from the sub-detectors and then sending the correct sequence of trigger signals to these sub-detectors, depending on the kind of event. The number of sub-detectors of ALICE restricts the generality of the trigger definitions, since the number of logical combinations scale like 2^N , where N in this case is around 50 - thus a very large number of possible combinations! The strategy for ALICE therefore was to make very rough global selections on events and then leave more sophisticated selections to the HLT, which will be described in the next chapter. A schematic diagram of the basic features of the CTP communication with other systems is shown in figure 3.14 [143].

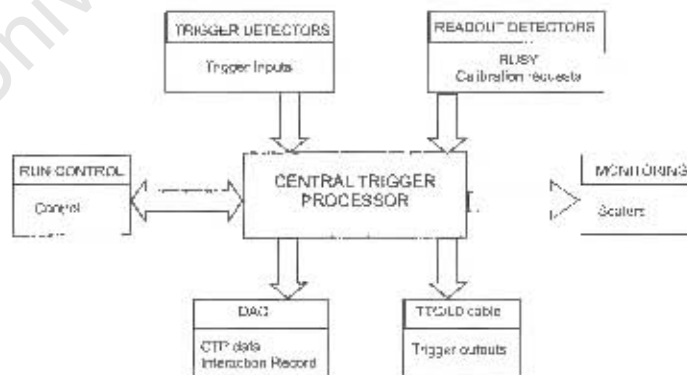


Figure 3.14: Schematic diagram of the ALICE CTP, showing the flow of various data types and the subsystems with which the CTP interacts [143].

The trigger inputs are divided into 3 levels depending on their latency : level 0 (L0), level 1 (L1) and level 2 (L2), in order of increasing latency time. The size of the detector, amongst other factors, such as the detector technology, dictates the latency of the trigger signals, since some cases a triggering detector may simply be too far away from the Central Trigger Processor (CTP) for the signal to arrive in time. Thus, every signal which can arrive within $1.2 \mu\text{s}$ of the event is included in L0, while after $6.5 \mu\text{s}$ the signals which have arrived from other triggering detectors are sent as the L1 signal. The third step, L2 is available after $88 \mu\text{s}$ and its main purpose is to issue the validation of the past-future protection.

3.6.3 ALICE data acquisition

While the trigger has the responsibility to announce the presence of an interesting event, it is up to the ALICE DAQ system to read out the event and write it to the data storage. Since there is a very complex set of triggers classes, as well as running conditions in the experiment, the DAQ has to be able to address them all. The ALICE DAQ has to be able to handle the most severe data rates imposed on it, which is during Pb-Pb collision running during one month of the year. The event sizes are directly related to the multiplicity of events and the multiplicity is in turn related to the centrality (see chapter 2). The average central event size expected, based on simulations assuming at most 8000 tracks per unit of rapidity, is 86.5 MB before compression.

The data read out from the detectors is transferred to the permanent storage, via a series of processing farms and an event-building network, which perform basic reconstruction, calibration, etc. A schematic diagram showing the incremental event-building process and the relevant DAQ hardware and software at each stage is shown in figure 3.15.

The strategy in ALICE has been to use low-cost commodity components for the DAQ processing architecture, instead of specifically-designed hardware. High-quality commodity PC's components, network interconnects and protocols will be used to buffer the data and transport it to and from the HLT processing farm. FPGA's will be used to ship the data directly from the DDL into the host RAM, which is used as a large buffer space directly after the detector read-out. The data from the detector is sent to LDC's via the ALICE-standard optical fibre DDL. The DDL's are terminated by Interface Units – Source-Interface Unit (SIU) on the originating side and Destination-Interface Unit (DIU) on the destination side. The DDL is a point-to-point link, designed to read out ALICE sub-systems in a standard way, consisting of a pair of optical fibres. They are used to transmit data both to and from the detector, with a bi-directional bandwidth of up to 200 MB/s [144]. The DDL is terminated in Destination Read-out Receiver Card (DRORC), which interfaces the DDL to the PCI bus of a computer which serves as the Local Data Concentrator.

The DRORC contains an FPGA and is capable of handling Direct Memory Access (DMA) transactions, giving the LDC full control over the DDL and allows event fragments to be transferred into the host RAM at a very high rate (up to 200 MB/s). The DRORC contains two DDL interfaces which are used to copy raw data to the

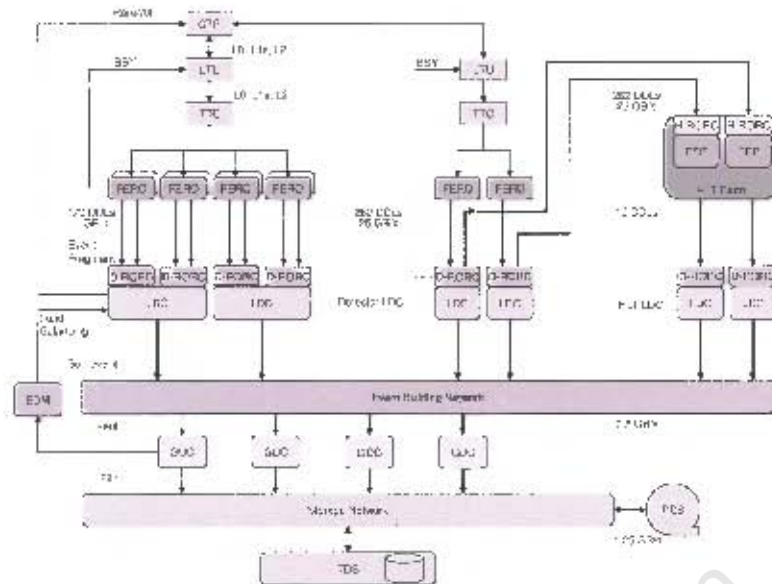


Figure 3.15: Data Acquisition (DAQ) data flow overview

DAQ and if necessary to the HLT farm for higher-level processing. Once the raw data is shipped into the memory of the LDC, it is sent over the event-building network to the appropriate Global Data Concentrator (GDC), where whole events are reconstructed. This event-building network connects the LDCs to the GDCs via TCP/IP over Gigabit Ethernet. The GDCs are also commodity computers running Linux and are responsible for full event reconstruction from the event fragments on Local Data Concentrator (LDC)s. Reconstructed events are shipped to transient storage in the experimental pit so that the DAQ is to some extent protected against problems with the connection to the permanent storage, which is located elsewhere at CERN.

The entire DAQ architecture is controlled by a set of software modules called Data Acquisition and Test Environment (DATE) [145]. A detailed description of the DATE software is not relevant to this work and the reader is referred to [145] for more information. The DAQ architecture described above is very simple, as all detectors (including the HLT) have the same interface. Load balancing is elegantly achieved by simply changing the number of LDCs or GDCs, since the DAQ architecture is completely independent of the layout of these.

3.7 The ALICE offline computing framework

The ALICE computing framework was developed in order to perform both the simulation of physics observables, with the resulting detector response, and the reconstruction and analysis of the data from real and simulated interactions. Since the framework is primarily based on the de-facto standard for data processing in HEP, ROOT [146], the ALICE computing framework was called AliRoot. AliRoot inherits

the basic functionality of ROOT : event generators such as Pythia [147], detector simulation (based on GEANT3 [148], shared memory classes, parallel processing, I/O classes, etc. For a schematic overview of the functionality and dependance of the AliRoot classes, see figure 3.16 [178].

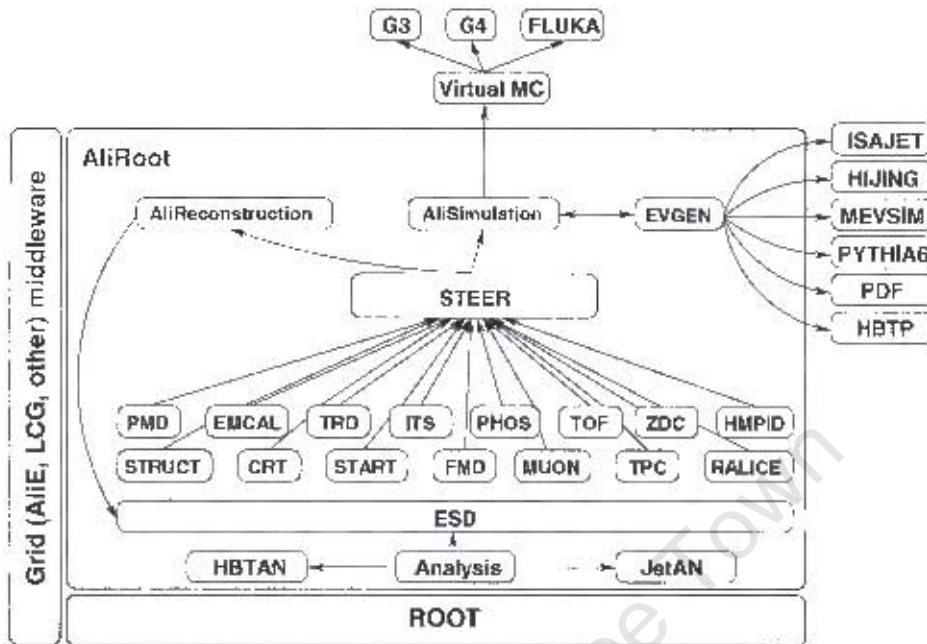


Figure 3.16: Schematic overview of AliRoot functionality and interdependencies.

3.7.1 AliRoot functionality - simulated and real data

In the ALICE framework, simulated and real data are treated similarly, although starting from different inputs. In the simulation case, the input is the physics generation, where basic parameters are used to generate distributions of *particles*. These are then transported through the detector, and the points in the sensitive volumes where *hits* are registered are calculated. The detector response is then applied to these hits, at which point *digits* (digitised detector response) is available. Once these digits are obtained, the data is packed into the raw data format and at this point, we have the intersection with the real data case.

In both the real and simulated data case, the starting point for the reconstruction algorithms is the raw data. These algorithms are clearly detector-specific and the description of the reconstruction data varies from detector to detector. The dimuon spectrometer, for example, has two levels of reconstruction, where the clusters are first obtained and then sets of clusters are fitted to a track model, in order to obtain track candidates. Once the parameters of these track candidates have been extracted, the information is stored in an ALICE-wide format known as the Event Summary Data (ESD). These various stages of data used in the framework, shown as a relationship between the amount of information available and amount of processing required, can be seen in figure 3.17. It is not within the scope of this thesis to describe the details of




Figure 3.17: The ALICE data processing framework, showing the connection between simulated data and real data, with the relationship between the amount of information available and the amount of processing required.

the rest of the AliRoot framework, and the reader is referred to ALICE Computing TDR [178] for further details. In short, the ALICE computing framework provides an integrated platform with which to process all data in the ALICE experiment.

3.8 Summary

This very brief description of the main features of the ALICE detectors subsystems serves to illustrate the scope of the capabilities of the detector. Here the focus has however been on the dimuon spectrometer and the trigger and data acquisition, as these will form the background to the present work on the dHLT. Before describing the dHLT in detail, however, it is necessary to elaborate on the design and capabilities of the more inclusive HLT, which will be done in the next chapter.

Chapter 4

The ALICE High Level Trigger

The ALICE High Level Trigger was conceptualised as a way to increase the significance of low cross-section observables, while respecting the financial constraints of mass-storage. These observables generally have characteristic signatures, making it feasible to write fast algorithms for identifying them. These can be coded into a set of triggers which allow data to be written for only those events or parts of events which may contain them. In this way, the contribution of background can be suppressed relative to the rate of physics signals. The scope of the HLT is essentially determined by the data environment at ALICE. The frequency of events recorded by the detector depends on the LHC beam characteristics (species, luminosity, interaction rate), while the size in terms of bytes of each event is determined by the number of channels in the detector(s) as well as the content of the raw data. The rate of past-future protected events is expected to be about 200 Hz in Pb-Pb collisions, for which the TPC data alone will be read out at approximately 15 GB/s. The mandate of the HLT, as given in the HLT Conceptual Design Report [156] is to **Trigger** based on an online analysis of the event, those events which are candidates to contain physics signals.

Select parts of events which contain physics signals, known as a region of interest (RoI)

Compress without loss and as far as possible, the raw data being read out after having been selected, in order to minimise the mass-storage requirements.

The HLT fits naturally into the trigger system of ALICE as an extra trigger layer between the L2 and event-building stages. It is designed to perform much more sophisticated calculations than the other trigger layers and this processing is implemented on a commodity cluster. The time and data rate constraints imposed on the HLT mean that a large amount of processing has to be completed in a small amount of time, while preserving data integrity. Advantage is therefore taken of parallelism in the data, which is in general hierarchical, with larger regions being built up from smaller ones. The processing algorithms therefore process the smallest possible size of data in a massively parallel fashion and the task is then to reconstitute the event faithfully, before the final trigger decision and event building stages.

In this chapter, an outline of the physics motivations (signals), focusing on those relevant to the dimuon arm, will be given, along with some technical details of the HLT in order to set the context for the next chapter, the dimuon spectrometer-specific high-level trigger, or dimuon High Level Trigger (dHLT). The processing communication framework and architecture will be described as they pertain to the dimuon high-level trigger and some benchmarks of the system will be shown.

4.1 Introduction

As described in the second chapter of this thesis, physics observables in heavy-ion collisions can be grouped into two types, depending on which momentum range they are probing. These are referred to as the “soft” (for low momentum transfer) and “hard” (for high momentum transfer) régimes of QCD. The bulk of hadron production, correlations, fluctuations, spectroscopy, *etc* can be described in terms of the soft régime. On the other end of the interaction scale are the hard processes, which are described by pQCD and describe the various aspects of jet and heavy-flavour production. There is a connection between the momentum scale and the cross-section of processes, with low momentum transfer processes having very large cross-sections and the high momentum transfer processes having cross-sections orders of magnitude smaller. In order to perform an analysis of the soft physics in Pb-Pb collisions a set of around 10^6 events (10^8 p-p collisions) is needed in order for the analysis not to be statistics-limited. Heavy flavour production ($b\bar{b}$ *e.g.*), or jet production with an energy of up to 200 GeV, is expected to occur roughly once every 10^4 events [110]. Clearly, to mitigate the effect of the large discrepancy in the cross-sections for these two kinds of events, some online selectivity is called for to assure a more balanced and significant data set, without overburdening both the DAQ and mass-storage bandwidth. Three ways in which this can be achieved have been identified, described above as the mandate of the HLT.

Due to its relatively long drift time, the TPC, which is the slowest of the ALICE detectors, is the focus of the HLT, since it also produces most of the data. Due to its slow drift time, the TPC records several superimposed events during TPC readout, which can be dealt with by the lower-lying ALICE triggers through the selection of “clean” events in Pb-Pb collisions. This is not the case in the higher luminosity p-p collisions, where only an online filter can separate superimposed events. The option therefore exists with an online HLT, to record data at 200 Hz in Pb-Pb collisions and up to 1 kHz in p-p collisions, both of which would allow ALICE and specifically the TPC to utilise the LHC luminosity to its greatest potential. Both of these situations would result in a sufficient accepted event rate for bottomonium spectroscopy and jet studies. Without the HLT, however, the resulting raw data would be - at 15 GB/s from the TPC alone - far too high for either DAQ or mass storage to handle.

4.1.1 Physics observables and specific triggers

A characteristic of many of the signatures of the QGP described in chapter two is that the most interesting are generally those with the smallest cross-section. The HLT is designed as a means to suppress the uninteresting background and increase the relative rate at which these signals are acquired. A brief description of their relevance and characteristics will be given below, as they pertain to the HLT algorithms and functionality.

Jet physics

The study of high- p_T jet production is one of the main goals of ALICE, as it gives information on the interaction of coloured partons with a deconfined medium. By studying the fragmentation function of jets¹ and the evolution of the jet cross-section as a function of centrality (or E_T), detailed information may be gained on the characteristics of the state.

In this case, the emphasis is on high- p_T jets in order for the fairly simple and fast online jet-finding algorithms to work efficiently. The basic assumptions of the algorithm are:

- Jets have a unique topology, with one “leading” particle carrying most of the energy, in the range of $20 \text{ GeV} < E_T < 30 \text{ GeV}$.
- They have a sufficiently large track multiplicity of $\langle n_{ch} \rangle \sim 10$ in order to stand out about the fluctuating minijet background.

These jets would be detected primarily in the central detectors, with the TPC and ITS providing tracking information.

Quarkonium spectroscopy

Both $c\bar{c}$ and $b\bar{b}$ quarkonia have significant branching ratios for decay into leptons, for example: the $\Upsilon(9460) \rightarrow e^+e^-$ decay produces two electron tracks of about 5 GeV. These are identified well by the TRD, and by using seeds provided by the TRD, the TPC can provide more accurate PID and momentum information on the candidate pair. A trigger rate of about 100-250 Hz for di-electrons from J/Ψ and Υ decay is expected in the TRD acceptance, but this rate is dominated by false positives. The more accurate HLT processing will reduce the level of background due to these “fake triggers” by more than a factor 10, resulting in an event rate of a few Hz.

The detection of the decay of the Υ state into dimuons is the main goal of the dimuon spectrometer. In an environment dominated by muons from soft background particles, such as π and K decays, subtraction of the background is usually done effectively with a simple momentum cut. A more sophisticated invariant mass cut requires more processing power and is the basis of the dimuon spectrometer High-Level Trigger. A specific tracking algorithm applied to the dimuon spectrometer

¹The fragmentation function of jets is defined as the distribution of momentum amongst the particles constituting the jet, relative to the leading particle.

will be described in the following chapter, where it will be shown that background rejection factors of up to 100 are achieved and the data rate can be reduced to 100 MB/s.

Momentum filter and trigger for open charm

The idea of a momentum and track filter is to track all particles online with the TPC information, in a massively parallel fashion, obtaining estimates (p_T and impact parameter) for the track parameters of each track. The predominant hadronic decay channels of the D family are

$$D^\pm \rightarrow K \begin{matrix} - \\ + \end{matrix} \pi^\pm \pi^\pm$$

$$D^0 \rightarrow K^+ K^-.$$

Considering that the average p_T of the decay products of the D mesons is $\langle p_T \rangle \sim 1$ GeV and that the average p_T of the soft pions which make up the bulk of the underlying event is ~ 0.4 GeV, then the fraction of tracks with $p_T > 0.8$ GeV is about 10% of the total. Thus, filtering out low-momentum tracks would be a good filter for selecting D -meson decays. What is more, the tracks from D -meson decays have a characteristic decay length, which allows a further filter on tracks which have impact parameters far from the main nominal interaction vertex. This procedure can reduce the raw data rate by a factor of 5 in the TPC.

Pileup removal in p-p collisions

The cross-section for rare processes, such as jets, Ω and Υ production, requires that ALICE record about 10^{10} events in order to perform a study which is not statistics-limited. If the LHC delivers the nominal p-p luminosity of $10^{30} \text{ cm}^{-2} \text{ s}^{-1}$, the average number of events in the TPC read out time would be 25. Clearly if these events are read out together, there is significant pileup. The average data volume per read-out would therefore be about 2.25 MB and thus DAQ saturation occurs at 300 Hz. In this case, only 10^9 events would be recorded per data-taking year, which is an order of magnitude lower than is necessary. The HLT would have the ability to separate piled-up events at a rate of 1.0 kHz in p-p collisions, a reduction in the raw data rate of about a factor of 6 can be achieved which would allow the resulting data rate of 400 MB/s to be handled by DAQ.

4.1.2 HLT design requirements

From the goals of the HLT described above, it is clear that it can best complete its assigned task by performing a complete online event reconstruction using information from as many detectors as possible. It is limited in its ability to do this by the processing requirements that such a task implies. For the TPC alone, which has the most severe processing requirements, the number of CPUs necessary can be estimated as follows.

The multiplicity in the central region from Pb-Pb collisions ranges from optimistic estimates using recent results from RHIC (about $\frac{dN_{ch}}{d\eta} \sim 3000$) to pessimistic estimates of up to $\frac{dN_{ch}}{d\eta} \sim 6000$. For the optimistic estimates, and taking into account Moore's Law [149], by 2006, about 400-500 CPUs will be required in order to process the raw data from the TPC, TRD, Inner Tracking Chamber (ITS), while the muon spectrometer would require a further 5-20. For the more pessimistic estimates of charged particle density, the number of CPUs needed for online event reconstruction is about 1000. This number includes the nodes required to satisfy the necessary overheads of networking, load-balancing, fault-tolerance, interfaces to other systems such as DAQ, *etc.* This number sets the overall functional requirement of the ALICE HLT. In order for the system to perform efficiently, several other requirements have to be imposed upon this.

Online processing capability The HLT has to be able to reproduce the offline results for momentum resolution and efficiency as closely as possible (generally within 10 %) although in a short enough time to respect the time constraints.

Computer network architecture The design of the computer architecture is dictated by two fundamental features of the experiment and HLT - on the one hand, the raw data has implicitly local information and a high degree of granularity, while on the other hand, for an HLT decision to be made, information from the entire reconstructed event has to be available. This suggests a hierarchical network topology with a high degree of connectivity between hierarchies, but not within them.

FPGA co-processors Some operations performed on the raw data (charge clustering and Hough transforms, *e.g.* are too time-consuming when implemented on PC-CPU's. If the algorithms are implemented on an FPGA situated on the front-end receiver cards, these act as co-processors for specific tasks. They reduce significantly both the size of the data coming from the front-end processors and the time required to do the processing (compared to performing the same task on a commodity CPU).

Communication framework Since the HLT is implemented on a distributed, parallel system, a software framework is required for processing components to communicate with each other in order for data and other information to be available where it is needed. To serve this purpose a communication framework has been developed [151], which is based on the so-called "Publisher-Subscriber", or "Producer-Consumer" paradigm. More on this framework can be found below and its adaptation to the dimuon High Level Trigger (dHLT) in the next chapter.

Interface to DAQ Although the HLT has to interface to many systems, such as the ALICE trigger system, Detector Control System (DCS), sub-detectors, DAQ, *etc.* it only has two fundamentally required interfaces in the data stream. These are the input data interface - the Front-End Processor (FEP)'s where the HLT

Scenario A	DAQ mode	No HLT processing. Data is read out directly to DAQ.
Scenario B	HLT processing, no trigger	HLT processing is present, but does not affect the data stream. HLT decisions are recorded for monitoring purposes
Scenario C	Full HLT mode	Full HLT processing of HLT data, selection of region of interest (RoI) is done event-by-event, as is readout of full events. Includes zero-suppression and data compression

Table 4.1: Foreseen ALICE HLT / DAQ modes of operation

receives raw data from the detectors - and the output data interface - where the HLT sends processed data and trigger decisions to the DAQ. It is thus designed such that the DAQ interfaces with it in the same way as any other sub-detector.

In order to be able to develop and debug the HLT system without affecting the ALICE data chain, there has to be the ability to bypass the HLT stage completely, or branch the data stream into an independent read out stream. Of course, once the system has been optimised, the possibility should exist for all data to be sent only through the HLT. These data taking modes are summarised according to the HLT Technical Design Report (TDR) [150] in table 4.1.2.

4.2 The online HLT

The functional components of the HLT are hierarchical, with the first level of the processing chain highly parallelised, operating on local data, with the results of that processing being merged further down the chain until finally the full event is reconstructed and a trigger decision can be passed. The schematic diagram of the hierarchy is shown in figure 4.1. The first level, marked “local pattern recognition” is detector-specific and could take the form of charge clustering – hit- or track-reconstruction – as in the case of the dimuon raw data. In some cases, a layer before the event reconstruction will perform global pattern recognition, such as in the case of the Hough transform on TPC tracks. The processed, reconstructed event is then passed on to the final level; and a trigger decision is issued.

4.2.1 HLT data flow

The logical data flow of the HLT (summarised in figure 4.2) starts with the arrival at all the sub-detectors of an “L2A” message, which signals all detectors to open their readout channels. The raw data is replicated by the Read-out Receiver Card (RORC), one copy of which is sent to DAQ and one copy of which remains on the Front-End Processor (FEP) nodes.




Figure 4.1: Schematic diagram of the functionality of the HLT

Although the actual processing performed at each step is detector-specific (each detector focusing on different physics), the general procedure is as follows. Local hit reconstruction is performed (possibly on FPGA co-processors) on the digitised raw data. From then on, the hit points may be handed off to a tracking algorithm to deduce track parameters. Once the tracks have been identified, their track parameters may be converted to “physics quantities” (such as momentum, PID, *etc*), which are then merged from all sub-detectors in order to form the global HLT decision. Some detectors, in particular the dimuon spectrometer, do not share much of the event with the other (central) detectors and their processing chain happens in much more independent fashion. The results of that processing are merged back into the data stream just before the global HLT decision. The HLT decision could validate or deny an entire event or only parts of events (RoIs), in which case only those parts of the data selected would be read out.

4.2.2 Software integration into AliRoot

The HLT online code for analysis is written in C++ and has to be adapted to the raw data. In order to compare offline simulations with the performance of the HLT online analysis, there has to be an interface to the ALICE offline code, AliRoot. Of course, this also is necessary for the Monte-Carlo data generated with offline event generators to be sent through the online system. The specific implementation of the dHLT interface to AliRoot will be discussed in detail in the next chapter.

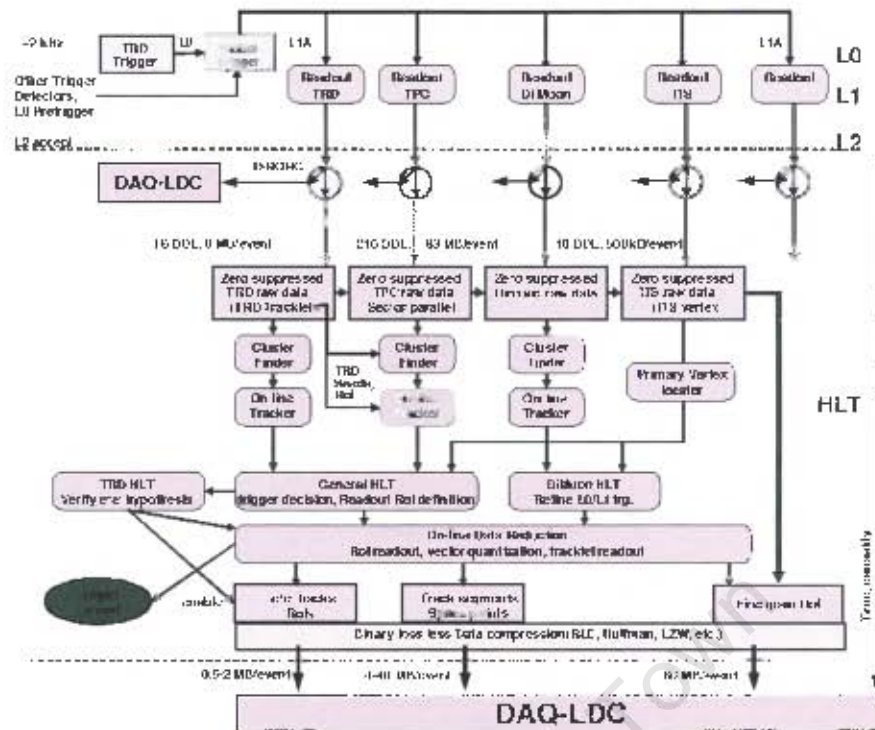


Figure 4.2: A general overview of the HLT data flow, showing the logical functions performed at each stage [156]

4.3 HLT computer network architecture

The HLT computer architecture is driven by the ALICE detector hierarchy. Basically, all ALICE detectors operate in a quasi-independent way, but are synchronised by the ALICE trigger signals - L0 for timing and L2A to set the start of readout. The data read out by the DDL's is transferred to the HLT FEP's while a copy is sent to DAQ). The raw data has an inherent granularity to it which is exploited by the architecture of the HLT; *e.g.* the hit reconstruction processing which is performed on the FEP's only requires local information and there is no need for communication between FEP's. It is possible to terminate more than one DDL in a FEP, but whether or not this is actually done in the experiment depends on the data rate of the particular detector, and the amount of processing required. Thus, it is unlikely that this will occur in the case of the TPC FEP's, where the data sizes and processing times are the larger, but it may be feasible in the dimuon spectrometer, where the event sizes are much lower and the hit reconstruction algorithm much faster, due to the simpler data.

The data delivered by the ALICE DDL's is interfaced to the FEP's via the PCI bus of the host machine, to which the so-called HLT-RORC card is connected. The use of the PCI bus to transfer information between the RORC's and the rest of the host machine ensures that the FEP can be as generic as possible, since almost all modern Personal Computer (PC)'s use PCI.

4.3.1 Connectivity and data flow

The overall connectivity requirements of the HLT are given by the aggregate input and output rates. Although the data flow into and through the HLT is defined by the physics requirements defined above, the maximum input rate is set by the number of DDL's connected to the system and their readout data rate. There are about 250 DDL's connected to the HLT (see table 4.2) from the sub-detectors involved. The DDL number and payload of each sub-detector is shown in table 4.2

Detector	Number of DDL's	Sub-event event size		
		Pb-Pb central (kB)	Pb-Pb peripheral (kB)	p-p (kB)
TPC	216	352	90	10
TRD	18	39	10	
Dimuon	10	15		
ITS	56	35		
Trigger	1			

Table 4.2: DDL connections and data event payloads for the detectors in the HLT system. The rates are limited in TPC co-incidence mode to 200 Hz in Pb-Pb running and 1 kHz in p-p running (see text). The dimuon spectrometer can be triggered at 1kHz if no co-incidence with the TPC is required.

Although it is a parallel-processing computer, the HLT does not require a generic symmetric network, where every node has equal latency and bandwidth to every other node. This is due to the fundamentally hierarchical nature of the processing framework and reduces significantly the number of inter-node connections and the switch sizes needed. The HLT will however also be used as an offline data processing farm when ALICE is not taking data, and so satisfactory bandwidth between nodes has to be guaranteed. For this reason, the inter-node connectivity is designed to be 1 Gbit per node. The resulting topology can thus be described as a generic backbone (of 20 Gbit/s) with some appropriate add-ons when required for online-specific HLT operations.

Several network technologies are available for the HLT to run over, Myrinet Gigabit Ethernet (GbE) and Fibre Channel being some of the more common ones. Not only are several networks available for study, but also several new interconnects such as Universal Serial Bus (USB) and IEEE 1394. The main criteria for choosing a network and interconnect are the bandwidth provided and the load on the CPU, since these are vital aspects of the HLT. Due to the hierarchical nature of the system, the main emphasis is on throughput (hence, bandwidth) and not latency as is traditional in massively-parallel systems. As network technologies and their associated price are rapidly evolving, it was decided to keep GbE as a baseline.

4.3.2 The HLT compute nodes

The HLT relies heavily on PCI interconnects in the compute nodes, and since almost all commercially available computers are based on PCI (or PCI-express), this means that there is a very wide range of choice of platform for the compute nodes. The main constraint is the operating system, which is chosen to be Linux. The current versions of Linux, in both the 2.4 and 2.6 kernel series support a large number processors in symmetric multiprocessor (SMP) mode. The HLT is flexible enough to support a heterogeneous set of compute nodes and scalable enough to allow the impact lower performance (but cheaper) nodes to be mitigated by simply adding more of them if the need arises. The decision on which CPU architecture and number of processors per node to use will be made as late as possible in order to take advantage of new developments in the industry and speed increases due to Moore's Law.

The setup for the dHLT computer architecture follows the general principles of that of the full HLT, with some simplifications. These will be described in the following chapter, along with the specific setup used for the dHLT prototype.

4.4 HLT communication framework

As mentioned above, the HLT is not what could be termed a "traditional" massively parallel compute system. While these are generally latency restricted, the HLT requires rather a high throughput. Processes run almost independently and proper buffering can result in very close to 100 % CPU utilisation. The optimisation then is for zero unnecessary network copying of data and not low latency.

4.4.1 Publisher-Subscriber paradigm

A distributed data processing framework was developed [151] as an interface between the different processing steps. The framework was based on the principle of a flow in a "data pipe", with processing along the way. The raw data at the start of the chain would be read (consumed) by the FEP's and once a step of processing has been finished, would be released (produced) for the consumers of the next step in the chain. This general paradigm of distributed processing, known as the "Publisher-Subscriber" paradigm has a few general principles which apply to the HLT :

- Data producers should be able to feed many subscribers. This is also known as a "fan-out" procedure.
- Processing on a single node is made much more efficient by passing only data descriptors, specified by the framework, between publishers and subscribers. The data itself remains in a reserved shared memory until needed. This reduces significantly the amount of bandwidth needed, both on the bus and on the network.
- In order to be flexible enough to support the different ALICE running conditions and triggers, the framework has to be able to switch between configurations at

runtime. This is also essential to ensure fault-tolerance in the case of a crash of processing nodes

- In order for debugging and monitoring to be built into the system elegantly, there are two kinds of data consumers defined :

Blocking consumers also called “persistent subscribers”, which reserve the input data until they are finished processing it.

Monitoring consumers or “transient subscribers”

- As much of the functionality as possible should be provided in a way that hides the operating system and network protocol behind abstract interfaces. Neither publishing nor subscribing components communicate directly, but rather via “proxy” components which deal with the communication protocols.

The HLT publisher-subscriber framework is an object-oriented design taking into account these principles, which allows for distributed fault-tolerant inter-process communication.

4.4.2 The HLT analysis object

In between the steps of subscribing to data and publishing the results is the actual processing of that data, which is done by the HLT analysis object. A schematic diagram of the functionality of the HLT analysis object and its relation to the *Publisher* and *Subscriber* components are shown in figure 4.3.

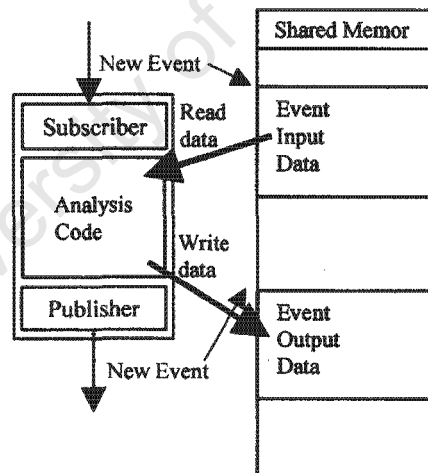


Figure 4.3: Functionality of the HLT analysis object. Data is subscribed to, processed by the analysis object and then published back into the event stream

Note that the data (published or subscribed) does not have to be contiguous, but may be fanned out or collated by so-called “scatter” and “gather” operations. The functionality for these operations is built into the HLT in for two reasons. The first is that some processing components require data from several sources (such as the

TPC track finder), which requires the ability to merge data from different streams. The second reason is that there should be the ability to distribute processing across nodes in order to avoid bottlenecks and provide a means for load-balancing.

4.4.3 Load balancing

Since the required compute time as well as the number of processing elements for any given step can vary significantly, in order to efficiently process the data, the ability for the data to be “scattered” into multiple streams has been maintained. The reverse is also possible, where multiple events are gathered back into one. These load-balancing components are the so-called *EventScatterer* and *EventGatherer*² respectively.

If an event is broken up into blocks and scattered to be processed in parallel by several processing components, a mechanism for rebuilding the event from its parts is required. This is the so-called *EventMerger*, which consists of multiple subscribers each subscribing to a block of a given event. The *EventMerger* merges all of the data from a given event and creates a new event descriptor for that event, which is then published for subscribers upstream.

4.4.4 Publisher-Subscriber bridges

The above data flow components are sufficient if the data resides entirely in the shared memory of a single node, where communication is only over the bus connecting several processors. If processing is to be done inter-node, there has to be a network-code interface in order to satisfy the general design principles laid out above. This is implemented by a so-called “Bridge” between the HLT framework code and the network code. The networking bridge is implemented in a transparent way, so that the processing components are not aware whether they are connected directly or over a bridge. This allows any processing to be moved to any node by inserting the relevant Publisher-Subscriber *Bridges* into the chain.

4.4.5 Publisher-Subscriber performance and fault tolerance

The HLT communication framework has been tested on 3 benchmarked machines [150], in order to test the timing, scalability and fault tolerance of the system. This was done in order to identify possible bottlenecks in the system as well as to predict future behaviour. Although the performance of the HLT code and existing network and computing technologies are sufficient to deal with expected event rates, a serious concern for the HLT is how it will handle failures in the system. It is estimated that with the number of nodes foreseen to be used in the HLT farm, severe failures will be experienced at a rate of about 2 per week [152]. In a system such as the ALICE HLT that is processing data continuously, these failures should be minimised as far as possible. The HLT nodes will contain a sensor which will monitor the status of the

²A convention used in this work has been chosen to put all HLT and dHLT software framework components in *italics*.

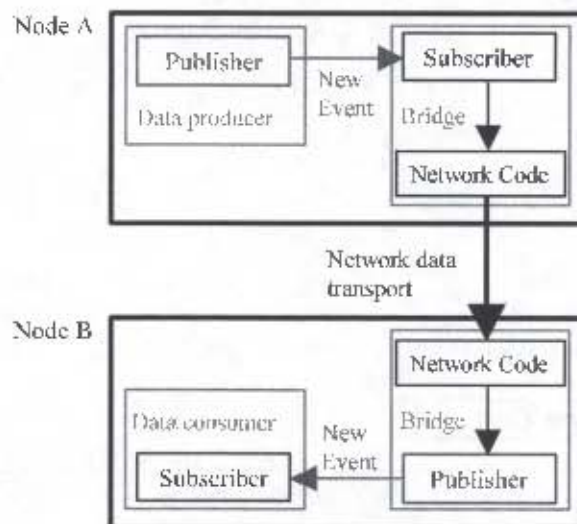


Figure 4.4: Schematic diagram of the network bridge of the HLT communication framework, showing the abstraction layer.

node components and issue alerts if necessary. The software framework is also fault-tolerant and if any particular node goes down, the jobs assigned to that particular node will be rescheduled elsewhere – either on a spare node or, if none are available at that particular, on currently available nodes when resources permit.

This fault-tolerance capability of the HLT was tested by disconnecting a processing node from the HLT prototyping cluster and observing the effect on the data output and input rates. Figure 4.5 shows that even a complete failure of a node has little effect on the total data processing rate, as a spare node was quickly assigned to process the overhead.

This general outline of the purpose and functionality of the HLT serves as an introduction to the specific application of the system to the dimuon spectrometer. The next chapters describe the design and implementation of the dHLT, along with some performance tests.

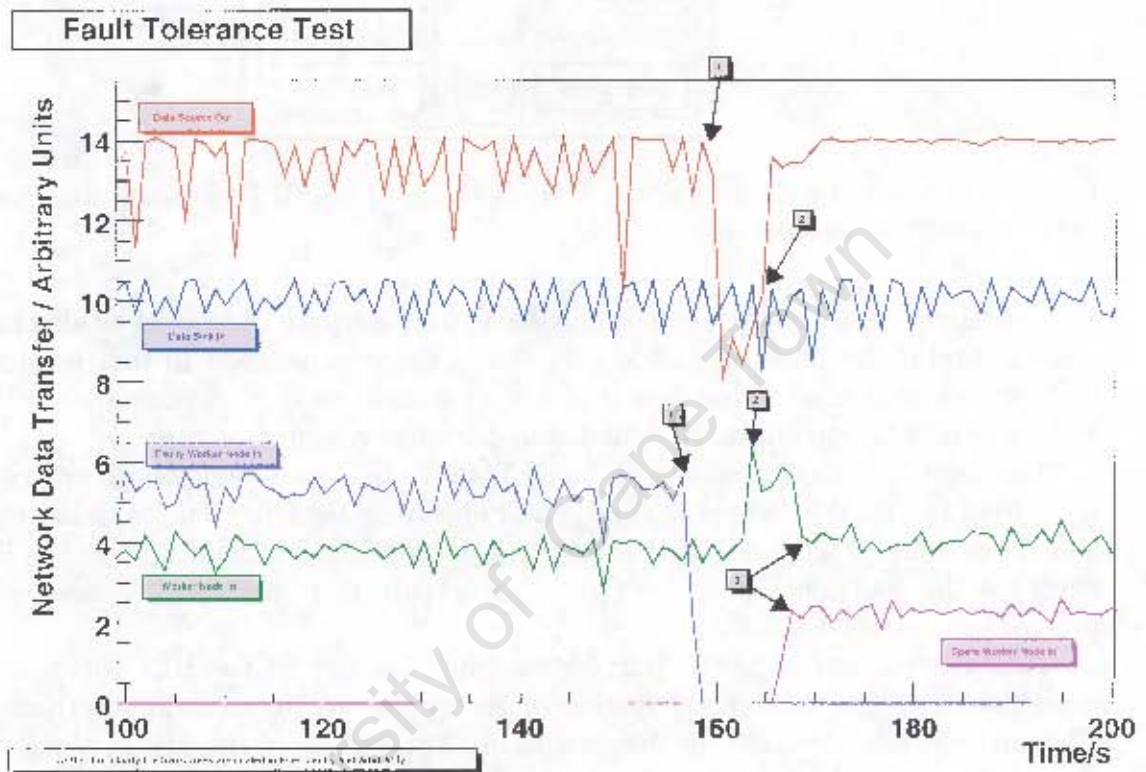


Figure 4.5: Test of the HLT failure tolerance with one node failing and one spare [150]. The events marked by "1" and "2" are the network interruption on the original node and the re-distribution of the task to the spare node respectively. Note that for clarity, the curves have been scaled vertically in order see them better.

Chapter 5

The design of the dimuon High Level Trigger (dHLT)

5.1 Introduction

The dimuon spectrometer is limited in the amount of data it can acquire, due to the restriction on the DDL bandwidth assigned to it by ALICE, to a maximum readout rate of 1 kHz (in stand-alone readout mode, for Pb-Pb collisions). As described in Chapter 2, one of the best candidate signals for the QGP is the anomalous suppression of heavy quarkonia - $b\bar{b}(\Upsilon)$ and $c\bar{c}(J/\Psi)$. The decay of these mesons into pairs of high- p_T muons provides, in principle, a fairly clean trigger with which to detect their presence and select events in which they are produced. This signal is, however, buried in a large muonic background mostly from the semi-leptonic decay of π and K mesons. Although this background is present over a wide range of p_T , it is largely at low p_T , below 1-2 GeV. At higher p_T , combinatorial contributions from semi-leptonic decay of open charm and beauty mesons dominates the background (see section 7.3 for a more detailed discussion of background sources). In order to reduce as much as possible the contribution from background sources, the dimuon spectrometer L0 trigger relies on a p_T cut done at hardware level from hit positions on the trigger chambers. The accuracy of the trigger trips is limited due to their size, which makes it somewhat inefficient around the cut value. This has the unwanted effect that a certain amount of background is accepted into the limited bandwidth. The consequence of this is that the relative level of the signal, which already has a small cross-section, is reduced.

By using more accurate hit positions from the tracking stations, it was proposed [155] that the calculation of the p_T of muons could be done much more accurately, which would allow for a more selective trigger. Instead of being forced to set the p_T cut very low in order to prevent loss of signal, one could reject a much larger amount of background, while still accepting a very large fraction of the signal. This is the current basis of the dimuon High-Level Trigger, or dHLT.

The dHLT is designed to run in three modes -

Pass-through where the dHLT flags events which should be accepted, but does not actually intervene in the data stream.

Yes/No Where the dHLT accepts or rejects entire events based on whether they pass the trigger conditions or not.

Filter Where only regions, or tracks of interest defined by the trigger are readout by the DAQ.

The dHLT is designed with 3 important functional features, which aim to take the most advantage possible of the level of parallelism in the data stream. For example, since each detector plane is readout independently and hit reconstruction is done effectively per-chamber, this means that these processes can be done in parallel. This is done in order to process data as fast as possible with the least amount of overhead or processor waiting time. The three aspects of the dHLT which permit this are

- the online hit-reconstruction algorithm¹
- the online track-finding and triggering algorithm and
- the data transport framework.

Various optimisations in these functionalities are possible depending on which implementation is being used and will be discussed further below. In order to access this information from the tracking stations and still make a trigger decision in time, a processing framework is needed which can produce trigger signals with the required latency – under 1ms for a 1 kHz trigger rate. This is achieved by processing the data readout in one DDL per plane in parallel, using an inter-process communication framework. Although this framework, which already existed at the time of the dHLT development as part of the global ALICE HLT, was not designed with the dimuon spectrometer in mind, the communication framework of the HLT is general and flexible enough to allow it to be modified for this specific case.

This chapter will describe the dHLT design. The global constraints on the design of the trigger system and the components of the dHLT are discussed and a description of the relevant components of the communication framework and its components is given. A description of the tracking algorithm for the trigger are detailed and its performance compared using offline simulations. The dHLT integration test and the results thereof are described, including the hardware and software components.

5.2 Data flow of the dHLT

In order to issue trigger decisions on a whole event, the data which is read out and processed in parallel has at some stage to be recombined, in order to have it available locally. This implies a flow of data from several nodes to a single one. In this section, an overall description of the data flow will be given, from the detector to the dHLT decision. The details of the processing components and data reformatting at each stage will be described in detail thereafter.

¹This is somewhat similar to the concept of “cluster finding” in the offline case, however the algorithm used in the online case does not do charge clustering - it only reconstructs hit positions.

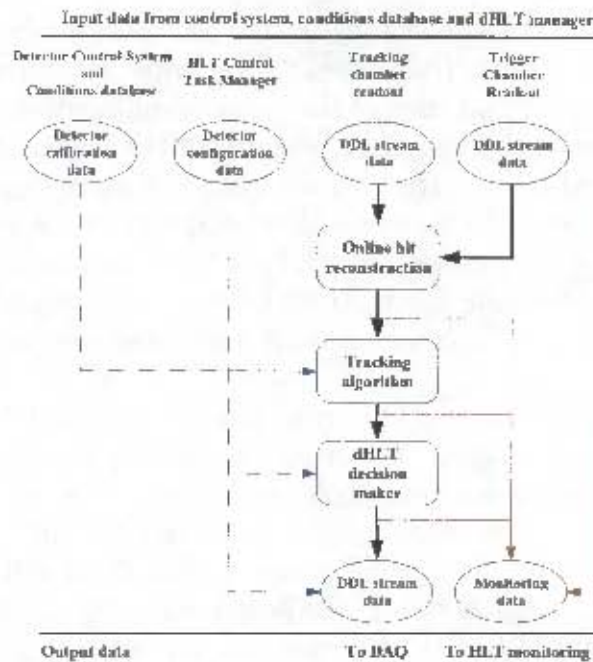


Figure 5.1: Schematic diagram of the data flow of the dHLT. The thick connections in black show the normal flow of data processed by the dHLT. The dashed blue connections show the input of calibration and control information from DCS and the conditions database. The thin connections in red show the extraction of monitoring data from the dHLT processing chain for quality assurance purposes.

As described in the previous chapter, the data is read out from the dimuon spectrometer in a raw data format specified by an ALICE internal note [136] and is packed in a tree-like format, “Russian puppets”-wise (see figure 5.2). When running during the actual experiment, the raw data has to be unpacked and reformatted into a format that the dHLT system can handle. This means that during the development stage, that raw data has to be generated from digitised simulated data, and packed in the correct format in order to test the performance of the dHLT components which relate to the raw data properly. The offline data is generated in a series of operations which are performed in the Monte-Carlo input. The first of these operations is the simulation of the actual Monte-Carlo hits on the detector, which is performed with the GEANT-3 [148] Monte-Carlo transport code. Once these hits are known, the detector response function is applied to the deposited energy, in order to produce the digital signal which would be expected, known as “digits”. The digits are then clustered into reconstructed hits, or “clusters” by a cluster-finding algorithm and these clusters are then fitted to a track model by the track-reconstruction algorithm. This hierarchical data model (hits, digits, clusters, tracks) is specific to the offline dimuon spectrometer software and is used for simulating the dHLT. The expected raw data is obtained from the digitised data, packed into the same format as would be done for the actual data stream by an algorithm in the offline code.

Once the raw data has been unpacked, one is able to perform hit reconstruction. Hit reconstruction is the process of finding local clusters of fired pads, and finding

the centre of gravity of these charge clusters, in order to increase the resolution of the chambers. In this way, resolutions of the order of microns can be achieved, as opposed to the detector pad sizes of the order of millimetres.

Once hit reconstruction has been performed, the higher-resolution reconstructed hit points are passed on to a tracking algorithm. This is a generic component which will perform tracking on the set of reconstructed hits, in a way defined by the specific tracking algorithm. The purpose of the tracking component is to search for track candidates, given very basic assumptions which will be described below, in the spectrometer chambers, and then to calculate the transverse momentum (p_T) of those tracks.

After track identification and filtering, it is the responsibility of the decision component to perform the trigger. The triggering is done at a local and global level; on individual tracks and events as a whole respectively. The result determines whether the event (or part thereof, if the trigger is running in “filter” mode) is accepted to be written out, or rejected. Finally, since ALICE DAQ will interface to the dHLT data streams as any other detector, the post-formatting of the data from dHLT data structures, back into DDL data structures similar to those shown in figure 5.2 has to be performed by the decision component as well.

5.3 Raw data generation and unpacking

The ALICE DDL raw event structure is only pertinent to this work insofar as it is necessary to understand the dimuon-specific data payload in order to unpack the data stream and reformat it into the necessary sub-event structures. Also, in order to realistically test the dHLT in the absence of actual data, it was necessary to have a means to generate a set of realistic raw data on which to perform hit reconstruction. For these reasons only a brief description of the DDL raw data structure will be given here; more information can be found at [131].

5.3.1 Raw data generation from ROOT trees

In order to realistically simulate the efficiency and speed of a software hit reconstruction, it is necessary to work on a realistic data set. Since the offline simulations store data in a per-cluster format in ROOT [146] trees, this data is far removed from what can be expected in realistic running conditions. The data in tree format was useful for testing whether the algorithm performed as expected, *e.g.* giving the expected resolution of hits, but it was not at all reliable for performance measurements. Since the real data will have to be unpacked from the raw format into a more convenient format before the hit reconstruction algorithm can be performed, it is necessary to start from the correct format. The real raw data stream does not have the convenient ROOT structuring and is read out per-chamber, instead of one reconstructed hit at a time. It is thus necessary to reformat the digitised data created with AliRoot into the DDL-wise raw data format specified by the ALICE-MUON collaboration [136]. This was done with an AliRoot macro which read in the simulated, digitised data and

then, with an accurate and realistic mapping file, packed the resulting raw data into a binary data file according to the format specified by the internal note [163, 154].

5.4 The dHLT hit reconstruction

The raw data generated from the ROOT trees has first to be converted into what can be generally described as “geometric” format, since the spatial positions, as opposed to the identifiers of the electronic components (channel number, MANU card, etc), are required for hit reconstruction.

Clustering algorithms, including the simple “centre of gravity” finding algorithm used for the hit reconstruction, are common in fields such as digital image processing, as well as various other applications of pattern recognition.

5.4.1 Datastream model and choice of reconstruction algorithm

The datastream of the tracking chambers is unordered in the sense that readout is not ordered according the spatial correlation. This means that charges from channels contributing to a given charge cluster can be situated far apart from each other in the data stream.

In order to identify clusters efficiently in an unstructured data stream, it is generally necessary to specify a unique property which all clusters in the data have. Usually, this is either the size of the cluster, or the total number of clusters, but these properties are not applicable to the dimuon data stream. The cluster size may vary from hit to hit (as shown in figure 5.3), and the number of clusters also contains a large spread, due to fluctuations in the number of particles produced in the collision. The solution is found in a feature of the electronic read-out of the detector. The high resolution of the MANU chip, which has a 12 bit ADC channel, assures that in each cluster there will be a unique *maximum* channel. This is denoted as the “central charge”, or “central pad” and since it is unique, the number of clusters is constrained to be the same as the number of central pads [157]. The situation is simplified by the segmentation of the detector strips on the slats of chambers 4 and 5 which have fine-grained segmentation in the y direction and lower resolution in the x direction (see table 3.3). This means that the hit reconstruction can essentially be performed independently in the x (column-wise), and y (row-wise) dimensions, since the probability of having two clusters on adjacent columns and rows is very small, of the order of 5%.

The spread in the number of channels fired by a single particle grows inversely with its momentum; lower energy particles deposit larger amounts of energy in the chambers and creating larger charge clusters than those of a higher energy. The hits from low-energy particles are considered as background. One problem with this algorithm is its performance without a lower threshold. Since it only searches to one neighbour distant, it may incorrectly reconstruct two maxima of charge, if very small charge values are not discarded. This behaviour will artificially increase the efficiency,

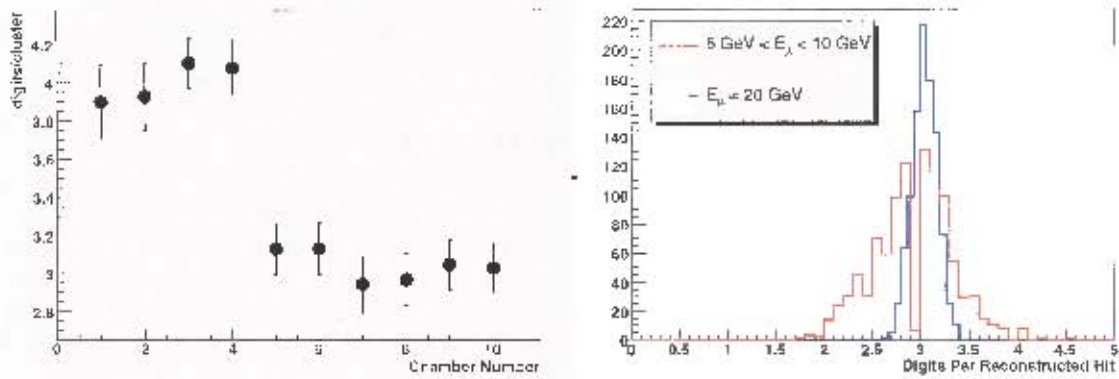


Figure 5.2: The average number of digits per reconstructed hit for the different spectrometer stations (left) for muons of 20 GeV. The distribution of digits per reconstructed hit for muons of energy 5 GeV - 10 GeV and 20 GeV is shown for comparison on the right, for chamber 8.

because the algorithm will find more than one maximum for a given hit - in some cases, efficiencies over 100% are possible, which is clearly unphysical.

This artificially high number of hits, with a significant fraction of overlapping hits, will slow down the track-finding algorithm. To overcome this problem, a simple threshold, or “DC cut” was applied to all ADC channels and only those with values larger than the threshold considered by the algorithm. The optimisation of this cut was done by considering the trade-off between execution time necessary and the efficiency of the algorithm at reconstructing hits. It was known that all central charges are above 100 ADC channels, while most of the non-central charges are below this value.

Once the threshold is applied, centres of gravity are found row- and column-wise for the bending and non-bending plane respectively. These are then merged where possible with the centres of gravity in neighbouring columns (rows). This is only done for central pads which are within one pad-dimension in the $y(x)$ -direction from each other, to find the full 2-dimensional hit [157].

The low occupancy, coupled with the low number of pads fired per particle means that the probability of overlapping hits is very low, about 2%. The hit reconstruction algorithm can therefore sacrifice sophistication for speed. In fact, for the purposes of the dHLT, where the time allocated in order to make a decision is below 1 ms, it is sufficient to do hit reconstruction only in the bending plane. This would give high spatial resolution in the bending plane, but only very coarse-grained resolution in the non-bending plan (of the order of the size of the chamber pads). Since the chambers are read out per-plane (with one DDL per plane), this means that in fact hit reconstruction can be done for each plane independently and plane-merging is not necessary.

During the development of the algorithm, the specific implementation of the algorithm was discussed and two scenarios were identified. These differed mainly in the sequence of processing and the amount of intermediate data storage (memory)

needed for the processing. The choice algorithm to be used to reconstruct the hits from raw data depends on the processing architecture it is to be implemented on. The original conceptual design of the full dHLT system called for hit reconstruction to be done on an FPGA in the front-end processing nodes. The algorithm to be used in that case should conform to the requirements of the FPGA architecture, with the major constraint is the footprint of the code, including the lookup-tables required for mapping and the amount of RAM needed during processing. It is interesting however, to note that this will not be the fastest algorithm for the PC CPU, which is not limited by the severe footprint constraints that an FPGA is, since the host memory can be used efficiently and this may be up to several GB. An algorithm optimised for the PC CPU could be made to run much faster on that CPU than one optimised for the FPGA. Since for the present the only architecture available for testing and benchmarking is the PC CPU, it is instructive to test two different algorithms to see the speed of one optimised for the CPU.

5.4.2 PC-CPU optimised hit reconstruction algorithm.

The algorithm optimised for the PC-CPU relies on a fairly large lookup table containing the geometrical and electronic mapping and segmentation information. This information consists of the MANU and channel numbers for each pad the logical pad position (iX, iY) in the single detection element local reference frame (ie, the slat), as well as the real position of the pad in the global reference frame and the bending/non-bending plane information. Due to the large number of channels in the dimuon spectrometer, this lookup table is very large and the algorithm takes full advantage of the presence of high-speed RAM on the PC to store this information. There is one lookup table per DDL.

The algorithm processes the raw data in a per-DDL way, in the following procedure, shown schematically in figure 5.3 :

- read in raw data in the DDL
- check each fired pad and consider whether or not the charge crosses the threshold.
- sort pads with charge above threshold by detection element (slat)
- central hits are found in the bending and non-bending plane independently by scanning the series of pads
- the hit is reconstructed by a centre-of-gravity calculation from the three pads belonging to a central pad.
- reconstructed hit positions in the bending and non-bending planes are merged to form a 2-dimensional hit.

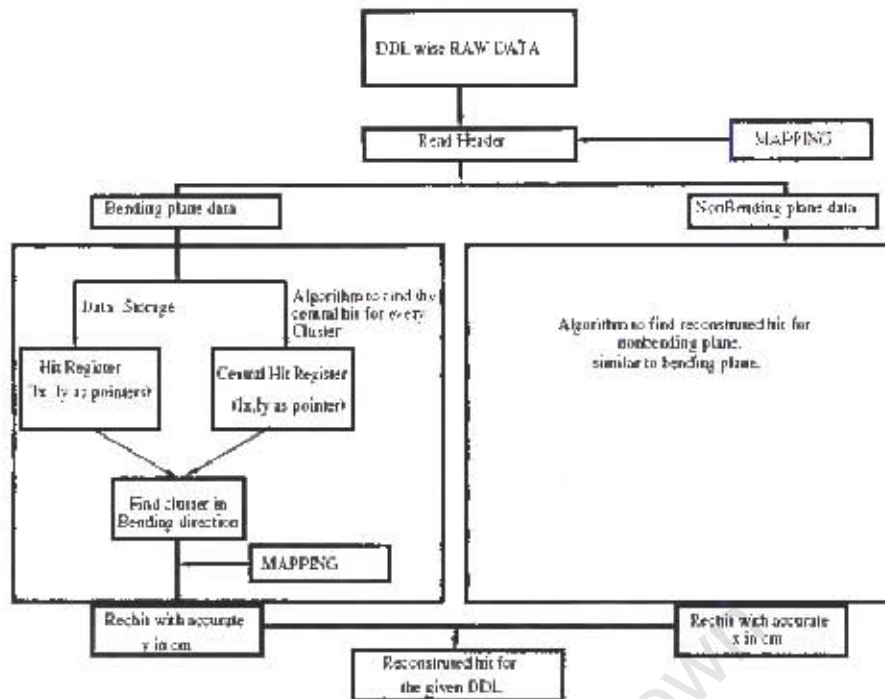


Figure 5.3: Schematic diagram of the PC CPU-optimised online hit reconstruction algorithm.

5.4.3 FPGA optimised hit reconstruction algorithm.

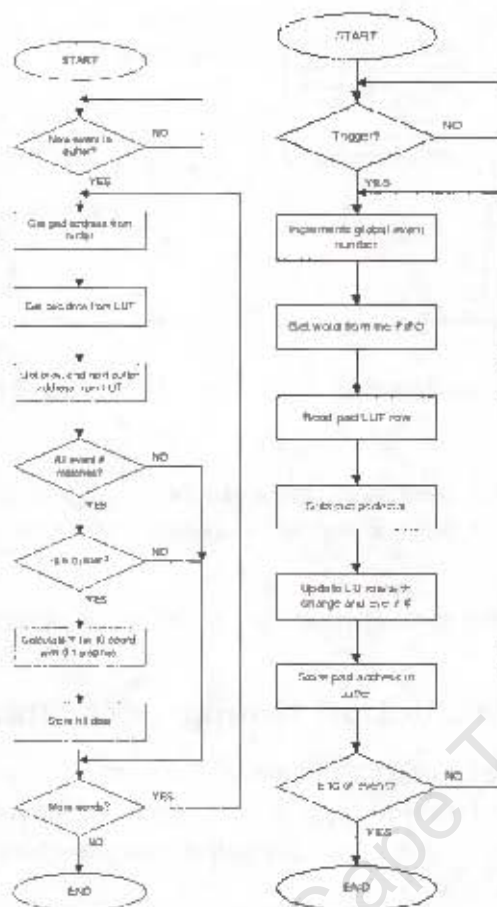
The algorithm optimised for the FPGA is constrained by the fact that it cannot have a very large footprint. This means that there is space for the mapping lookup table and the implementation of the reconstruction algorithm, but very little else. Therefore, the data has to be processed in two steps :

Read and map In this step, the data is read in from the DDL, placed into a buffer and mapped from the raw data format into the geometrical data format using the mapping lookup tables.

Nearest neighbours and centre of gravity search The data is then passed over (in the bending plane strips) and pads with high charges in them are found. When such pads are found, a neighbour with high charge is searched for and if found, the centre of gravity calculation is done.

Clearly, the data has to be passed over twice in this case, as opposed to only once in the case of the PC-CPU optimised case. At the time of writing, the algorithm optimised for the FPGA had been coded and tested only on the PC-CPU, in order to check its efficiency.

An outline of the algorithm is shown in figure 5.4(a), with the corresponding state machine given in 5.4(b). The timing of the algorithm is not realistic, as it has not been tested natively with synthesised code on the FPGA, although this is in progress. Since the first algorithm has been shown to respect the timing constraints set on it, it is




(a) Schematic outline of the FPGA-optimised hit reconstruction algorithm.
 (b) State machine of the FPGA-optimised hit reconstruction algorithm.

in principle sufficient for the final dHLT. The process of testing faster implementations is a work in progress. For the remainder of this work, the hit reconstruction will refer to the PC-CPU optimised algorithm, unless specifically stated otherwise.

5.4.4 Hit reconstruction resolution performance

The resolution of the hit reconstruction algorithm has been estimated by comparing hit positions found with those generated. A large sample of events was generated with singles muons, and then transported through the detector apparatus with the GEANT toolkit, and the detector response applied. The digitised data were packed as usual into the DDL format and then processed with the hit reconstruction software. The positions found by the hit reconstruction software were then compared hit-by-hit to that indicated by the simulation.

A sample distribution of these "residuals" - the differences in position between the generated and found hits - is shown in figure 5.4. Assuming a Gaussian distribution



(c) Residuals in the non-bending ($y - z$) plane. (d) Residuals in the bending ($x - z$) plane.

Figure 5.4: Measured hit reconstruction residuals ($x_{\text{hitrec}} - x_{\text{GEANT}}$) in the bending (left) and non-bending (right) planes, fitted to Gaussian distributions.

of the residuals, we have a resolution in the bending plane of less than $100\mu\text{m}$.

5.4.5 Hit reconstruction timing performances

In order to see whether the time constraint on the hit reconstruction algorithm is satisfied, the hit reconstruction algorithm was coded into ROOT macros and timed. There are three factors which could influence the time needed for the hit reconstruction


1. Channel threshold cut (DC cut)
2. Processor clock speed
3. Hit reconstruction algorithm : which algorithm is used, how efficiently it is implemented, and on which platform

Benchmark results of the hit reconstruction algorithm

The algorithm was and benchmarked by running it on a 2 GHz computer. The time needed was measured as a function of the DC cut imposed with fixed number of input particles (150 per event), is shown in figure 5.5(a). The efficiency is also shown in figure 5.5(b) as a function of DC cut. The results are satisfactory, in that with the nominal DC cut of 50 channels – chosen to ensure a maximal efficiency-to-timing ratio – the time required to reconstruct the hits of an event is well below the 1ms limit. [158]

Implications for the implementation of the hit reconstruction algorithm

The choice of the final architecture on which the dHLT hit reconstruction algorithm will be implemented is yet to be decided. Although, from the results of the tests run on the PC-CPU shown above, it is clear that this will suffice. For the foreseen



(a) Timing of dHLT online hit reconstruction (b) Efficiency of the dHLT hit reconstruction algorithm as a function of DC cut with a algorithm as a function of DC cut. nominal number of particles of 150.

Figure 5.5: Timing and efficiency of the online hit reconstruction algorithm on tracking chambers 4 and 5, as a function of DC cut, for a nominal number of 150 muons in the spectrometer.

dimuon spectrometer data rates, the chinese-checkers algorithm processes the data with sufficient speed, efficiency and accuracy. The number of front-end PC's necessary is much less than that of the ALICE TPC, which requires several hundred. By contrast, the dimuon spectrometer requires one CPU per DDL, which corresponds to only 20 CPUs - or 10 dual-CPU front-end nodes - for the tracking chamber hit reconstruction. The final choice of how to implement the hit reconstruction on the FEP's awaits the outcome of the implementation and testing of a fast algorithm on an FPGA².

5.5 The dHLT tracking algorithm

In the general sense, a trigger algorithm has to be designed in order to make a decision on an event or part thereof. This involves setting some condition for the data under investigation to pass, and checking whether it actually does so. In the case of the ALICE dimuon spectrometer, an example of this condition may be³ that the particle which resulted in a track in the spectrometer, assumed to be a muon once it has traversed the muon filter, has a transverse momentum greater than some threshold, the p_T cut. This is the bare essence of the trigger algorithm - a cut on the p_T of candidate tracks. The manner in which these tracks are selected is up to the specific tracking algorithm used.

²Testing is currently in progress

³This is the only trigger criterion discussed at length in this thesis, however it is not the only conceivable scenario. The conception and testing of other triggering logic is an ongoing work. Naturally, as the spectrometer is designed to accurately measure the decay products of mesons, which decay into high p_T muons (see chapter 1), the first trigger scenario was conceived as a p_T cut.

A first algorithm for selecting tracks based on seeds from L0 trigger records from the trigger data was proposed and studied by Frank Manso [155], but this is by no means the only algorithm capable of identifying tracks and making a cut on their transverse momentum. The tracking algorithm that is used with the dHLT has to fulfil a few very general requirements:

- It has to be as efficient and inclusive as possible : as many tracks as possible should be identified correctly as candidates.
- It has to be as accurate as possible : the p_T of the identified tracks should be calculated to high precision.
- It has to integrate well into the data transport framework : the algorithm should be able to be coded in a way that makes it compatible with the modularity and interfaces of the data flow framework.

The first investigation into the possibility of a dimuon High Level Trigger done by Manso [155] was based on a track-finding algorithm which used L0 information as seeds to search for hits on tracking stations 4 and 5, in order to construct tracks. This was a fairly linear and iterative (with only 2 iterations) procedure to which there exist in principle alternatives. This so-called “Manso” algorithm was simulated in rough terms with offline code and provided the baseline for what a tracking algorithm of this kind could achieve.

An alternative to “a first algorithm for the dHLT”, “a second algorithm for the dHLT”, is currently under study [160]. This is a hierarchical algorithm which works on a local level to build up tracks. It is inherently more parallel than the first algorithm and is a possible alternative. The details of this algorithm are not within the scope of this thesis, and the author refers the reader to the work in progress [160].

5.5.1 The “Manso” tracking algorithm

The basic principle of the tracking algorithm is to use the L0 records as seeds to search for and validate candidate tracks in the spectrometer, in a very efficient and fast way. The raw data from each local trigger circuit recorded by the muon trigger electronics contains estimated information on the position, deviation (in the y direction) and transverse momentum of the particle which caused it to fire. An implicit assumption is made here, that the particle has originated at the Interaction Point (IP) and has passed through the dipole magnet. This assumption, along with the deviation reported by the trigger circuit, allows the tracking algorithm to open a search area, or region of interest by extending a vector towards tracking station 5 of the spectrometer. This vector is calculated taking into account the dipole magnetic field downstream and assuming that only muons pass through the muon filter immediately preceding the trigger chambers.

For each candidate, an intersection point R_p on each of the tracking chambers on station 5 is determined, from the L0 information and the z -position of the chamber. This intersection point then forms the centre of a search area, with radius R_s . Since




Figure 5.7: The last tracking chamber of the dimuon spectrometer (chamber 10) is shown with a sample (central) event. The digitised hits are shown, along with an example of a possible search radius from the L0 information. The search circle has its centre at R_P and a radius of R_S given as a function of R_P (see text).

Once the hit positions of the track candidate are determined, the transverse momentum of that track is then calculated. Assuming a flat magnetic field distribution in the dipole, as well as straight tracks in the $x - z$ plane and neglecting multiple-scattering effects, the transverse momentum of a track is given as a function of the deflection angle (θ_d), the magnitude of the magnetic field B (T) the z -length of the field area (L) by the by [128]

$$p_T = 0.3 \frac{qBL}{\theta_d} \cdot \frac{\sqrt{x_F^2 + y_F^2}}{z_F}, \quad (5.1)$$

where x_F, y_F and z_F are the co-ordinates on the detector, inside the dipole field, as shown in figure 5.8. The positions of the hits on the third station, inside the dipole magnetic field are estimated, using the above assumptions on the track, from the positions on the first and second trigger chambers (x_1, y_1, z_1) and (x_2, y_2, z_2) respectively :

$$x_F = \frac{z_F}{z_1} \cdot x_1, \quad y_F = y_2 - \frac{y_2 - y_1}{z_2 - z_1} \cdot (z_2 - z_F), \quad \theta_d = \frac{1}{z_F} \left(\frac{y_1 z_2 - y_2 z_1}{z_2 z_1} \right).$$

Using standard error propagation, the uncertainty on the transverse momentum of

tracks calculated in this way can be shown to be proportional to the uncertainties in the positions x_1 , y_1 and y_2 ⁴




Figure 5.8: ALICE dimuon tracking algorithm p_T calculation principle, adapted from the L0 trigger principle [128]

Since this tracking algorithm starts with the L0 candidates as seeds for track finding, it cannot recover tracks that L0 has not accepted. In order to overcome this shortcoming, it was proposed to set the L0 cut very low, far below the dHLT cut at the J/Ψ threshold, in order to maximise the efficiency. Since the global L0 efficiency (integrated over p_T) is around 85%, this tracking algorithm can never be more efficient than this value and can only reject background.

The efficiency of the L0 cut is a factor contributing to the total acceptance of dimuons, and the level of the cut with the dHLT in place is a matter for discussion. An issue which at the time of writing was under investigation is how the dHLT tracking algorithm will handle high hit density environments. With no pre-filtering of the L0 information, the dHLT tracking algorithm may be flooded with too many seeds to process in the allotted time. However, if the dHLT tracking algorithm is fast and efficient enough at high multiplicities, there may be no need to make a p_T cut at the L0 level - all the fired circuits of the L0 trigger may simply be passed on to the dHLT. This method may in principle accept more signal than with the pre-filtering of L0 information, but it is doubtful how valuable the information would be, since the filtered tracks would be at low- p_T and hence have a high probability to be rejected by the dHLT anyway. The prospect is under study[160].

5.6 Benchmarks of the Manso tracking algorithm

The fundamental properties of the dHLT algorithm having been defined, it was decided to perform some basic benchmark tests, using simple input distributions, in order to better understand the intrinsic efficiency of the algorithm. The efficiency is defined as an output distribution divided by an input distribution. In the case of the dHLT efficiency, this is the dHLT-calculated p_T spectrum divided by the input p_T spectrum.

⁴The uncertainties in the z position are assumed to be very small since the detectors are aligned in this direction very accurately with the Geometrical Monitoring System (GMS) system.

5.6.1 Cut width and momentum reconstruction accuracy

Of course, the efficiency of the cut depends on the accuracy of the p_T determination of the algorithm. The uncertainty on the p_T of a track calculated from the dHLT can be expressed in terms of the uncertainty of the hit positions as well as the inherent uncertainty of the dHLT algorithm. The former is well-defined by the segmentation and hit-reconstruction code described above, while the latter refers to the simplifications made in the dHLT tracking algorithm expression for the p_T of a track.

The systematic uncertainties of the L0 and dHLT can be estimated by resorting to high-statistics Monte-Carlo simulations. A large number of muons with a flat distribution in p_T were generated and then L0 and dHLT processing performed on them [160]. For the simulation the entire dimuon spectrometer was included, including those parts of the ALICE detector which share the dimuon rapidity region : the forward detectors (V0, T0 and ZDC), the front absorber, the dipole and L3 magnets and the muon filter. Assuming that the uncertainty of the dHLT and L0 p_T estimates are Gaussian distributions, the systematic uncertainty can be found from fitting the differences between the input muon p_T and the p_T calculated by either the dHLT and L0 respectively. These fits, shown in figure 5.9, serve as a useful estimate of the systematic uncertainty. From figure 5.9 we can see that from the simulations, the

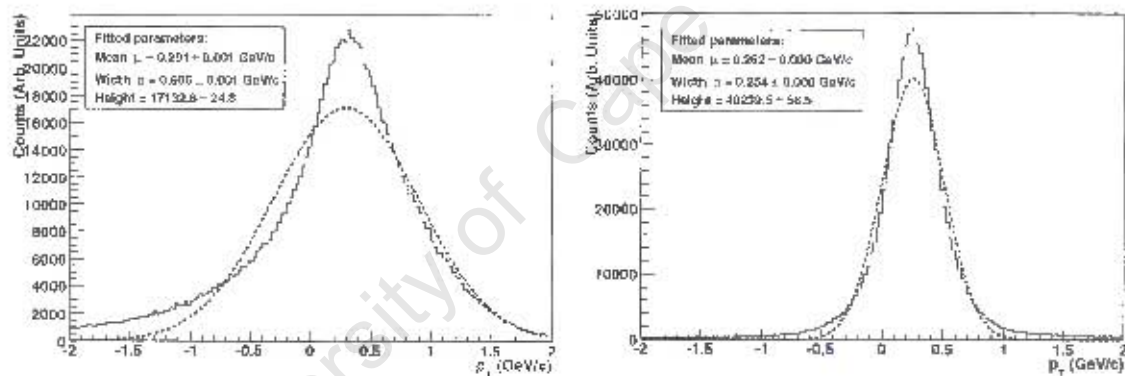


Figure 5.9: Fits of the difference between the p_T calculated by the L0 (left) and dHLT (right) algorithms and the input p_T of muons. Note that the width of the dHLT distribution is significantly smaller than that of the L0 distribution.

accuracy of the dHLT tracking algorithm is about 250 MeV at a p_T of 2 GeV. In order to investigate the sharpness of a p_T cut, assuming that the uncertainty on the dHLT p_T calculation is a Gaussian distribution, the efficiency of the dHLT can be parametrised as an error function :

$$E(p_T) = \frac{A}{2} \left[1 + \text{Erf} \left(\frac{p_T - p_T^{\text{cut}}}{\sqrt{2}\gamma} \right) \right] \quad (5.2)$$

Here, the parameters A (amplitude of the function), γ (the “skin width” of the cut) and p_T^{cut} are to be fitted from simulated distributions. The results of these (normalised) fits are shown in figure 5.6.1. For the given accuracy of our tracking and




Figure 5.10: Fits of the L0 (left) and dHLT (right) efficiency functions. Note that the fitted skin width (γ) of the dHLT function is smaller than that of L0.

hit-reconstruction algorithms, the skin width is fitted to be 722 ± 406 MeV for a 2 GeV cut.

Previously mentioned was the possibility of other tracking algorithms which had different basic assumptions. Since the current “Manso” algorithm uses only the bare minimum of information from the tracking chambers necessary to predict RoIs and calculate the p_T of tracks, it is possible that it does not have the most accurate p_T prediction possible. A systematic study was undertaken [160] to investigate the relationship between the skin width of the efficiency curve, the intrinsic resolution of the hit reconstruction algorithm and the tracking algorithm used. A general expression for the skin width in terms of the spatial resolution $\sigma_{spatial}$ of the hit reconstruction algorithm and the transverse momentum p_T resolution of the tracking algorithm is given by :

$$\gamma(\sigma_{spatial}, \sigma_{algo}) = \sqrt{\frac{2p_T^{cut2} \lambda^2 \sigma_{spatial}}{\pi z_F^2 (\tan(9^\circ) - \tan(2^\circ))^2} + 2\pi \sigma_{algo}^2} \quad (5.3)$$

where λ_k is an integral over θ at a given p_T cut :

$$\lambda_k = \int_0^{2\pi} \left[\left(\frac{\sigma_{rec}^x z_F}{\sigma_{rec}^y z_1} \cos \theta \right)^2 \cdot \left(\frac{z_2 - z_F}{z_2 - z_1} \sin \theta - \frac{z_2}{z_2 - z_1} \cdot \frac{p_T^{cut}}{0.3qBL} \right)^2 + \left(\frac{z_F - z_1}{z_2 - z_1} \sin \theta + \frac{z_1}{z_2 - z_1} \cdot \frac{p_T^{cut}}{0.3qBL} \right)^2 \right]^{\frac{1}{2}} d\theta. \quad (5.4)$$

Here, q , B and L are the charge of the particle, the dipole magnetic field strength and distance the particle travels through that field respectively, as defined above. The positions z_1 , z_2 and z_F are those of the tracking stations 5, 4 and 3 respectively (as defined above). The values σ_{rec}^x and σ_{rec}^y refer to the resolution of the hit-reconstruction algorithm in the x- and y-directions respectively.

The equation for the skin width is parametrised in terms of the resolutions of the tracking and hit reconstruction algorithms. It is thus possible to investigate the behaviour of this parameter as these resolutions are changed, which would reflect possible developments of the hit-reconstruction and tracking algorithms. It was found that significant improvements could be made by having a more accurate hit-reconstruction (and hence track-finding) algorithm ($\sigma_{spatial}$), but that in order to improve the skin width by orders of magnitude, an improved intrinsic accuracy of the algorithm (σ_{algo}) is needed.

5.6.2 Background rejection factor

With symmetric p_T cuts at the L0 and dHLT level, it has been shown to suppress by roughly a factor of 5 the amount of data accepted by the spectrometer, by rejecting background based on its tracking algorithm. This factor has been estimated to be even higher (around 6-7) if asymmetric cuts are placed on track candidates at the L0 and dHLT level, with a low p_T cut at L0 and a high cut at the dHLT level.

5.6.3 Acceptance

The acceptance of dimuons from the decay of J/Ψ and Υ mesons with this tracking algorithm has been investigated and compared to the acceptance of L0. While the acceptance of J/Ψ dimuon decay products is about 10% less (around 69% for dHLT compared to 78% for L0), the acceptance of dimuons from the Υ decay is decreased by only a couple of percent (81% for dHLT compared to 86% for L0).

5.7 Summary

The design of the dHLT builds closely on that of the HLT, which allows it to take advantage of the HLT's software modules as well as hardware components. It has been shown, with very simple input distributions, that the dHLT can reconstruct tracks more accurately than the L0, and that with the current tracking algorithm and hit-reconstruction accuracies, a sufficient enhancement is achieved. The possibility exists to improve these results by implementing improved tracking or hit reconstruction algorithms and a general expression for the skin width of the efficiency curve has been given in terms of these factors. We will now describe the actual implementation of the systems both offline and online, in the following chapter.

Chapter 6

Implementation of the dHLT

6.1 The online and offline cases

The dHLT has been developed as a set of software components with interfaces to each other as well as data sources and sinks. This chapter aims to describe these software modules and how they relate to each other. The dHLT is designed to be used in two modes

online with full integration into the data transport framework, operating in real-time

offline with interfaces to the ALICE offline simulation framework, AliRoot [162].

Both of these environments are essential to the development of the dHLT system and have different data input and output formats. The interfaces are designed such that the processing components are independent of the data sources or sinks, which makes them easily interchangeable and allows for testing the dHLT in different modes. All of the above description of the dHLT tracking algorithm has been tested in the so-called “offline” computing environment, which is to say that the algorithm was tested on simulated data, with no interface or connection to the HLT data transport framework. This analysis is valuable in order to check global properties of the algorithm, such as the maximum efficiency, background rejection, or the time taken for the dHLT tracker to identify tracks and compute their transverse momentum. Furthermore, by isolating the dHLT modules, it is much easier to identify problems related directly to a specific module or another.

In the online case, the emphasis is on realistic data transport and timing, in order to understand whether the system can physically handle the load placed on it. Interfaces to the data transport framework were used to simulate the dHLT in a semi-realistic way, similar to the envisaged situation at the actual experiment¹. The input terminus (raw data) is left open-ended and several options are being developed [160, 165], including input from data files or data published from RORC’s. In the final system, this will be provided by the data streams from the dimuon spectrometer

¹As work on the project progresses (see [160, 165]), these interim processes and data processing components will be replaced with the definitive versions.

DDL's, while during prototyping this has to be provided by files with generated raw data.

6.2 Interface to the HLT data transport framework

As mentioned before, the data transport framework of the dHLT is the same one as the ALICE HLT, with dimuon specific processing components. A more detailed description of the full functionality of the framework is given in Chapter 4, and only the relevant parts will be highlighted, where they pertain to the development of the dHLT system. For a full description of the data transport framework, see refs [161, 151]

The input terminus to the publisher-subscriber framework depends on whether the system is being used in the online or offline mode. In the online case, the data input is from the detector DDL's via the FEE's. In the offline case, the data is published from data files which reside in memory. A file- or DDL-publisher receives the L2A, either from the detector trigger electronics or from a mock-up which simulates the signal at a given rate. Once the data has been published, the processing chain can proceed, starting with the reconstruction of hits from the raw data.

6.3 Offline implementation of the ALICE dHLT software

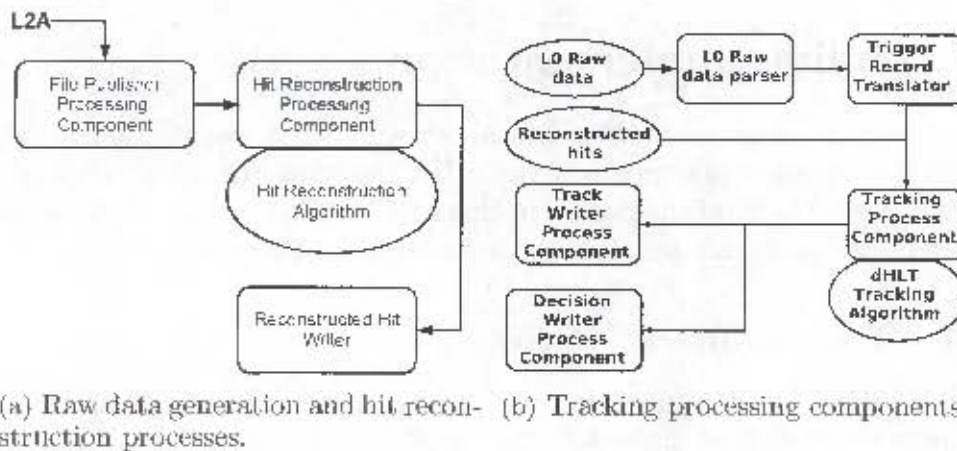
The processing chain of the dHLT consists of several processing and data transformation steps. These processes and their dependencies will be outlined below.

6.3.1 Raw data and reconstructed hits

The first function of the dHLT is to acquire the raw data and transform it into the reconstructed hits format. The components which are required for this task are

- L2A mock signal generator
- File publisher
- Hit reconstructing process component
- Hit reconstruction algorithm
- Reconstructed hit writer process component

The L2A signal generator would provide the "start" signal to the file publisher, in a similar way that the real L2A signal would open the gates on the readout electronics. The actual rate at which this is done can be tuned, in order to simulate different



(a) Raw data generation and hit reconstruction processes. (b) Tracking processing components.

Figure 6.1: Schematic diagram of the reconstructed hit (left) and tracking (right) processing components.

readout rates. The File publishers then publish the data contained in the files and this data is then handed off to the hit reconstruction processing component subscribers. Once the hits have been reconstructed, the data is published to a reconstructed hit writer, which writes the data into a stream. This stream is then published awaiting subscription by a tracking component.

6.3.2 Tracking component and Decision component

The reconstructed hits that are published are subscribed to by tracking components, which perform the track finding and p_T cut. The dHLT components for this are :

- PubSub Framework
- Trigger Record raw data
- L0 Raw Data Parser
- Trigger Record Translator
- Tracker Processing Component
- Tracking Algorithm
- Track Writer
- Decision Record Writer

The L0 trigger data, once parsed and translated into the format necessary for the tracking algorithm, provides the seeds for the tracking. This is then passed by the data transport framework to the tracker processing component, where the tracking algorithm is executed. Once the track has been constructed and the p_T has been calculated, the dHLT decision can be made. Since the dHLT decision is very simple, it is coded as part of the tracking component and the results of this are written with the Decision Record Writer.

6.4 Online implementation

The processing components described above were what was deemed to be necessary to fulfil the design requirements of the dHLT, however the actual implementation is not discussed. The implementation in this case was discussed and two scenarios were envisaged - a “callback” architecture and a simplified monolithic one.

6.4.1 The “callback” option

Since the dHLT has the inherent concept of regions of interest, it is possible to forward only these data of interest to the decision components. An obvious advantage of this scheme would be that there would be a reduction in the data rate over the communication network, since only fractions of events would be sent, however there were disadvantages too. The most severe of these disadvantages was that there is a requirement for a “callback” mechanism, where the decision component would request the RoI's from the nodes where the reconstructed hit information would be stored. This feature is not present in the HLT data transport framework, since in this framework data is envisaged only to be transported in one direction. An elaborate scheme of message passing would have to be constructed over the functionality of the HLT framework in order for the RoI callback mechanism to work. Since this represented a move away from the philosophy of using standard framework components and the actual level of the reduction in the data sent over the network was not clear, it was decided instead to follow simpler and more practical approach.

6.4.2 The “monolithic” option

Although the callback option offered a very elegant and efficient solution in principle, the technical problems stated above meant that the implementation of this method would take some time. A different solution was found, by removing the RoI-passing functionality from the dHLT decision component. The tracker would pass on entire events instead of only parts of them to the decision component. This meant that there was no need for the decision component to communicate its desired region of interest to the tracking processes. This option was dubbed the “monolithic” option since the entire data of an event is passed from the track-finding processing components to the decision processing component, instead of only parts of events.

Note that separate tracking components do not have to run on separate nodes, thanks to the network-transparency of the HLT data transport framework. This also means that the decision component may be on the same node as one or more tracking components. The only constraint was the amount of resources required at the expected data rates by the various processing components. This would dictate how the processing components would be spread amongst HLT compute nodes. For this reason, and in order to integrate all of the software components to determine if there were any design flaws at the time, an integration test was undertaken.

6.5 The dimuon High Level Trigger integration test

It was decided to test the system on the cluster at the UCT-CERN Research Centre. Although at the time of the integration test, the full data chain had not been completed – in particular, raw data generators had not been included in the chain, as well as the RORC publishers – the exercise would be valuable to test the interfaces between the processing components. The interfaces to the data input would also be defined for later integration and the design of the whole system could be tested to see if it would actually be able to handle the load placed on it in ALICE running conditions. What is more, with a test bench setup, it would be possible to stress the system in various ways, in order to determine its behaviour under varying states of network, Central Processing Unit (CPU), disk Input / Output (I/O) load, etc.

6.5.1 Test-bench technical specifications

The test was performed on the UCT-CERN Research Centre's computing facility, Cluster for African Research in Massive Energetic Nuclei (CARMEN) [164]. At the time of the integration test, the cluster shown in figure A.3 consisted of 20 compute nodes running a Linux 2.6 series kernel and a head node with attached Redundant Array of Inexpensive Disks (RAID) storage.

All the computers were Intel Pentium IV based, with the compute nodes running on 2.6 GHz Hyper-Threaded CPU's nodes, with 512 MB of Random-Access Memory (RAM) and 512 kB Level 2 cache. The server was a 3.2 GHz CPU with 1 MB of L2 cache and 1 GB of RAM . The operating system of both compute nodes and head nodes was a customised version of RedHat Linux 9, with the head node running a 2.4.25 based kernel. Several patches, including those for security and the PCI and shared memory interface driver necessary for running the HLT were applied to this kernel. The compute nodes' network interfaces were D-Link 520 Fast Ethernet PCI cards, connected via a SureCom 10/100 Mb/s switch. The switch had a Gigabit backplane and two Gb/s ports, one of which the head node was connected to. This meant that the maximum bandwidth available to each node would be around 10 MB/s, while to the head node there was a maximum of 100 MB/s available. An Network File System (NFS) [166] was exported to the nodes from this port, but they also had 80 GB hard drives with a local ext3 scratch partition with roughly 50 GB of free space for data. More details on the system can be found in Appendix A.

6.5.2 Integration and stress test goals

Apart from the actual test of the integration of the dHLT code modules described above into a coherent, working form, the November 2004 test aimed to stress the dHLT system in two independent ways :

1. Determine whether the dHLT architecture could handle the foreseen rates, in as realistic as possible as scenario.
2. Determine whether the dHLT tracking algorithm gave sensible results

In the first case, since final versions for all components were not available at the time, realistic values for the time required to process data would be inserted as dummy process components, where the tracking would normally have been done and this would test whether it was possible with our network architecture to handle the foreseen event rate. In the second case, the tracking components would be inserted, but reconstructed hit data would be fed in from a file instead of from DDL streams. This would check whether the tracking component would efficiently and faithfully reconstruct the simulated tracks.

6.5.3 Dummy load stress test.

15 Node Dummy Tracker Stress Test Setup – dHLT Integration Test

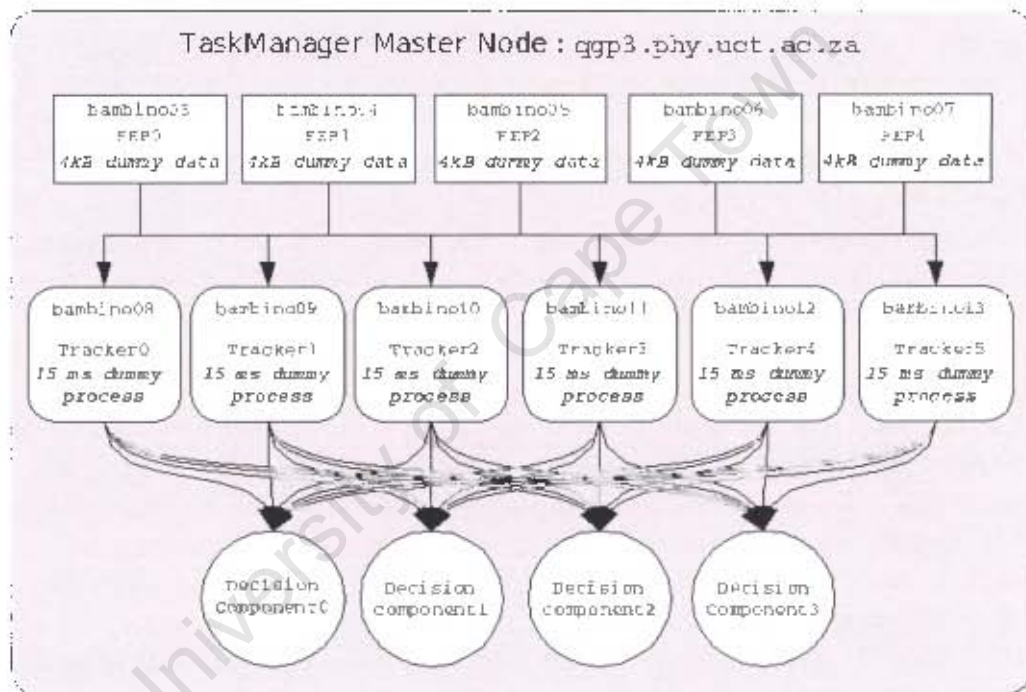


Figure 6.2: Schematic diagram of the physical data flow of the dummy stress test. Connections and framework processes are shown, but not individual dHLT framework components (see figure 6.3)

In this test, the framework of the dHLT system would be tested to see whether it could handle the data rate foreseen during the highest data rate foreseen - Pb-Pb collisions at 1.0 kHz. The raw data and tracker were not included, but the time taken to run the hit reconstruction process from raw data and tracking components was estimated and included in the framework as “dummy processes”. These are processes which do not actually perform tasks, but use CPU time in the same way as the actual processing components. Nominal event sizes and rates were chosen to correspond to

the respective maximum values, including the ALICE-wide safety factor of 2 on the expected multiplicity, namely 500 kB/event at 1.0 kHz for Pb-Pb collisions. Note that the situation under study corresponds to central Pb-Pb collisions, where the particle multiplicity is the highest, while the trigger rate is the dimuon L0 rate limit specified by the ALICE trigger system, not the rate of central events foreseen. This scenario assumes that all events are central, which is clearly not the case and is a very pessimistic assumption. However, since the purpose of the test was to see to what extent the system could be stressed, the very worst case scenario was chosen.

The total event size of 500 kB for the dimuon spectrometer was assumed to be shared equally between 5 FEP's, with 2 *FilePublisher* components each. This corresponds to the situation in the experiment, where 2 RORC's will be hosted per FEP, each reading out a DDL of the spectrometer. Each of these *FilePublishers* was then subscribed to by an *EventScatterer*, which fanned out the data into 6 streams, for load-balancing, since the events are assumed to be published asynchronously. Upstream, a set of 6 dummy dimuon tracking processing components subscribed to these events, making the number events published and subscribed to equal on the average. Since the aim of the test was to check *only* the tolerance of the system to a high event rate, the actual tracking processes were substituted for dummy processes, which eliminated a source of uncertainty in the timing. The dimuon tracking component had been tested with offline code and was found to take on the average 15 ms to find tracks in simulated central Pb-Pb events. Therefore, the dummy tracker consisted of a dummy process which lasted on the average 15 ms, with a slight Gaussian spread of 1 ms.

Once these processes had completed, the data was then published to 4 streams via another *EventScatterer*. Corresponding to these four streams were four Decision Components which represented the end of the processing chain and signalled the end of the event to the *TaskManager*. The full layout of the physical processing chain is given in figure 6.2. The HLT framework components besides the dummy tracker and decision components is shown in figure 6.3.

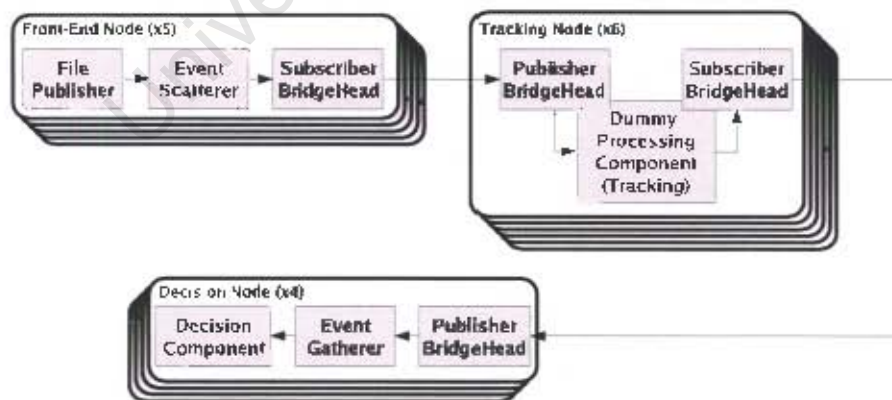



Figure 6.3: The HLT framework components implemented in the dummy stress test, along with the data flow.

The above setup was run at a low event rate (below 1 kHz) and was found to be


stable, with all the components working as expected.



(a) Data output rate for a file-publishing node at an event rate of 1.2 kHz (b) CPU utilisation for the file-publishing node at an event rate of 1.2 kHz

Figure 6.4: File publisher node resource utilisation during the dHLT stress test.

The event publishing rate was then pushed first up to 1.0 kHz and then to 1.2 kHz, in order to fulfil the goals of the stress test - to see whether the same behaviour at low rates would be exhibited by the system at realistic rates. At 1.0 kHz the system




(a) Data input and output rates for a tracking node, at an event rate of 1.2 kHz (b) CPU utilisation of the tracking node at an event rate of 1.2 kHz

Figure 6.5: Tracking node resource utilisation data during the dHLT stress test.

again performed as expected, with the load on the CPU's reaching around the 70% - 90% mark. At 1.2 kHz, the network bottleneck started to become apparent and rates above this could not be sustained for more than a few minutes without the system failing.

The CPU and network utilisation for the nodes in the various levels of the chain are shown in figures 6.5(a) and 6.5(b). In figure 6.4(a), we can see that the network bandwidth is saturated from the file-publishing nodes, in this case `bambino03`. The data output rate on the file publishing nodes is roughly equivalent to the data input rate measured on the tracking nodes, as shown for `bambino08` in figure 6.5.3. Shown alongside is the CPU utilisation on this node while data is being processed. The CPU utilisation is not saturated, but nevertheless quite high. It is also shown that the data output rate is smaller than the input rate. This is not surprising since dummy data



(a) Data input and output rates for a (b) CPU utilisation for the same node Decision node, at an event rate of 1.2 at 1.2 kHz event rate. kHz.

Figure 6.6: Resource utilisation for a *DecisionComponent* node during the dHLT stress test.

was used, but this was to represent the situation in reality, where the data is converted from raw data to reconstructed hits. The output data of the tracking nodes is then published to the decision components running on the decision nodes. An example is shown in figure 6.5(b), in this case *bambino18*. In this test, the decision component performed very little processing, basically only signalling the end of an event. In the real case, the decision component would have to perform slightly more intensive processing, basically performing a cut on the data set. The output data rate might also not reflect reality, in the “callback” case where ROIs are sent back to tracking nodes to select regions to be read out, since this information has to be transmitted to the publishing nodes. The actual size of the increase in data rates depends heavily on the encoding of the ROI’s, and the mechanism used to transmit them, but as stated before, the callback option has been left open as a possible refinement of this work and was not implemented for the current results.

The length of the dummy stress test run was limited by the data storage capacity of the processing nodes, since a very verbose log level was chosen and these messages were spooled to log files which could reach several hundred MB. As a proof of principle, the test was successful, showing that even with a low-bandwidth network, like Fast Ethernet, it is possible to run the dHLT at full rate. The bottleneck was clearly with the network bandwidth as nowhere along the chain was CPU utilisation saturation reached. If the hit reconstruction algorithm is eventually implemented on the RORC FPGA, then the CPU utilisation on the publishing nodes will reduce significantly. Given that the final dHLT will operate on a GbE network, with more powerful machines, the dHLT stress test can safely be assumed to show that the system will manage the load placed upon it in the final experiment. The question of whether it will be as accurate and efficient as offline predictions was attempted to be answered in the tracker integration test described below.

6.5.4 Monolithic dHLT tracker integration

The purpose of the second test was to see whether the tracks generated by simulated events could be recovered and properly triggered on by the actual dHLT algorithm. For this test, it was required to isolate as many aspects of the framework that did not depend on the actual algorithm as possible. As it was already shown that the system could handle the full network load, the tracker integration test was set up in a way so as to test only the software components. The data flow across nodes was kept to a minimum, making the simplest possible connections between software components. Since the network interfaces were transparent to the dHLT processing components, this has no influence on the algorithm.

The monolithic implementation of the dHLT tracking and decision components was run on a single node and were found to perform satisfactorily.

6.6 Global HLT test over public Internet network

Taking advantage of the fact that most of the members of the dHLT collaboration were together in Cape Town in November 2004, it was decided to test the limits of the performance of the HLT data processing framework, in a field which it was never designed to work for, the public Internet. Since there are no hard latency limits in principle on the data processing framework, it was possible that it would perform even over thousands of kilometres. Although many of the processing components were not available at the time, it was decided as in the case of the dummy stress test to use mock-up “dummy” versions of these² in order to test the HLT framework over as far a distance as possible. The test was envisaged as a proof-of-principle demonstration, in order to show the feasibility of transcontinental real-time distributed computing system. The obvious limitations of the study, which are essentially network congestion-related, should be kept in mind. The primary aim was to produce a system in a working state, while a secondary aim was to study the level of the impact of these limitations, if a working system could be achieved.

6.6.1 Global test setup

The sites in the test were the University of Bergen (Norway), the Kirchhoff Institute of Physics in Heidelberg (Germany), the Joint Institute for Nuclear Studies in Dubna (Russia) and the UCT-CERN Research Centre at the University of Cape Town. It was decided that the data flow should mimic as closely as possible the real data flow of the ALICE experiment, with sub-events being built into full events hierarchically. Of the various participating sites in the test, three were in the northern hemisphere connected on a reliable and comparatively fast network; the bottleneck was clearly in the link to the southernmost site, Cape Town. As this was also the site of the dHLT implementation on CARMEN, not to mention the fact that the dHLT collaboration

²The dimion raw data publishers, *e.g.* were not available, and many of the dHLT and HLT processing components were still under development.

was gathered there at the time, it was decided that the final steps in the chain – merging dHLT track information and passing HLT trigger decisions – would be performed in Cape Town. The other parts of the chain were linked via Heidelberg in the following fashion, shown schematically in figure 6.7 : The three sites in the

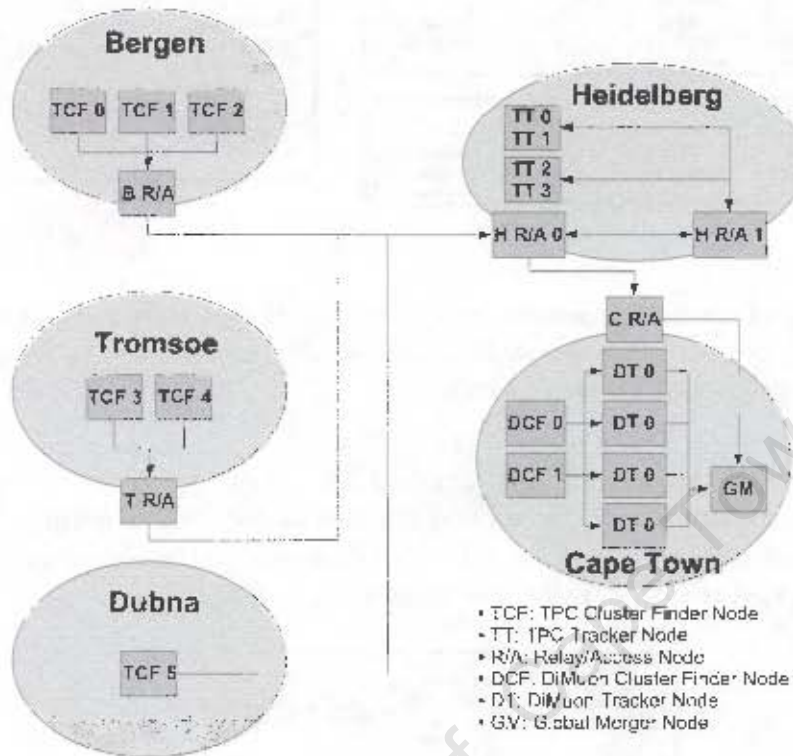


Figure 6.7: Global HLT setup with all involved sites and nodes. The components on each node are not shown.

Northern Hemisphere were set up to run dummy versions of the TPC cluster finder. The TPC data stream is read out hierarchically³ – “slices” are constructed from “patches” on which cluster finding is performed. Dummy processing components for one, two and three patches were set up at Dubna, Tromsø and Bergen respectively, with each patch being processed on a separate node. To illustrate, the processes running on the nodes in Bergen is shown for example in figure 6.8, with the light shaded box representing a node in the cluster and the dark shaded boxes representing HLT components. Except in the case of Dubna, a forwarding node was necessary, since the clusters compute nodes are generally protected from the public network. The forwarding nodes at the various institutes were connected to the internet, and all connections were relayed by them to the compute nodes via their TaskManagers, described below (see also chapter 4). The output data published at these sites was then sent to Heidelberg. Subscribers there merged the 6 published event fragments into a single event and dummy processes simulated the next step in the processing, which was track-finding on the clusters. This was done on two dual-CPU nodes –

³See the ALICE TPC TDR for more details.

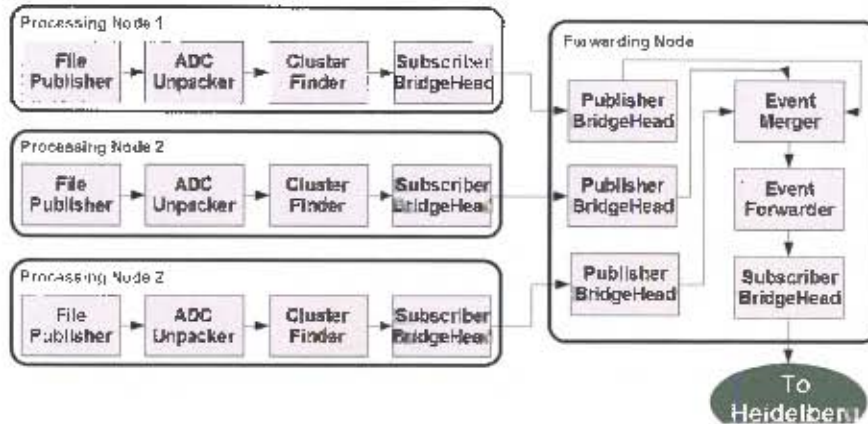


Figure 6.8: Setup of the HLT processing modules at the Bergen cluster. Black ovals denote physical nodes, while grey boxes denote HLT processing components. The network bridges to the relay/forwarding node are also shown.

one of the nodes was responsible for handling the subscription and merging of the incoming data, forwarding it to the second node to perform the tracking. The results of this tracking setup, shown in figure 6.9, were then sent to the cluster in Cape Town. The CPU usage and network traffic was monitored in Cape Town.

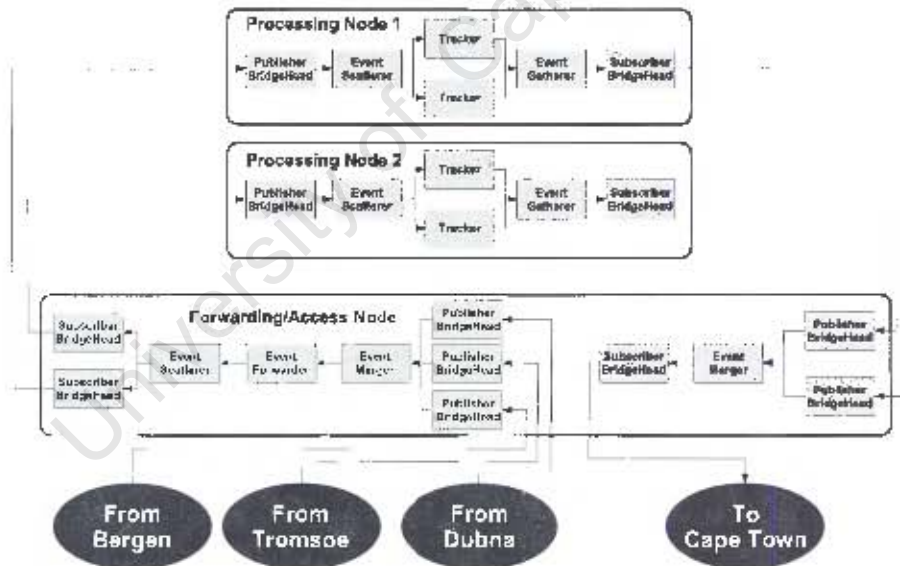


Figure 6.9: Cluster-finding subscription and tracking processes running on the Heidelberg cluster. Black ovals denote physical nodes, with grey boxes representing HLT processes.

In parallel to the processing in the Northern Hemisphere, the cluster in Cape Town was set up to process dimuon data. Two dummy cluster-finding processes, each running independently on a separate node, were scattered to dummy tracking processes. There were four tracking components, each running on a node. The output of these tracking components was set to the global merger node, which also subscribed

to the results from Heidelberg.

6.6.2 HLT *TaskManager* setup

As described in the previous chapter, the HLT uses a module called the *TaskManager* to ensure that the components necessary for the given setup are initiated in the correct order and synchronised. In this case, the system was controlled by a 3-level hierarchical *TaskManager* setup shown in figure 6.10. Overall control of the system

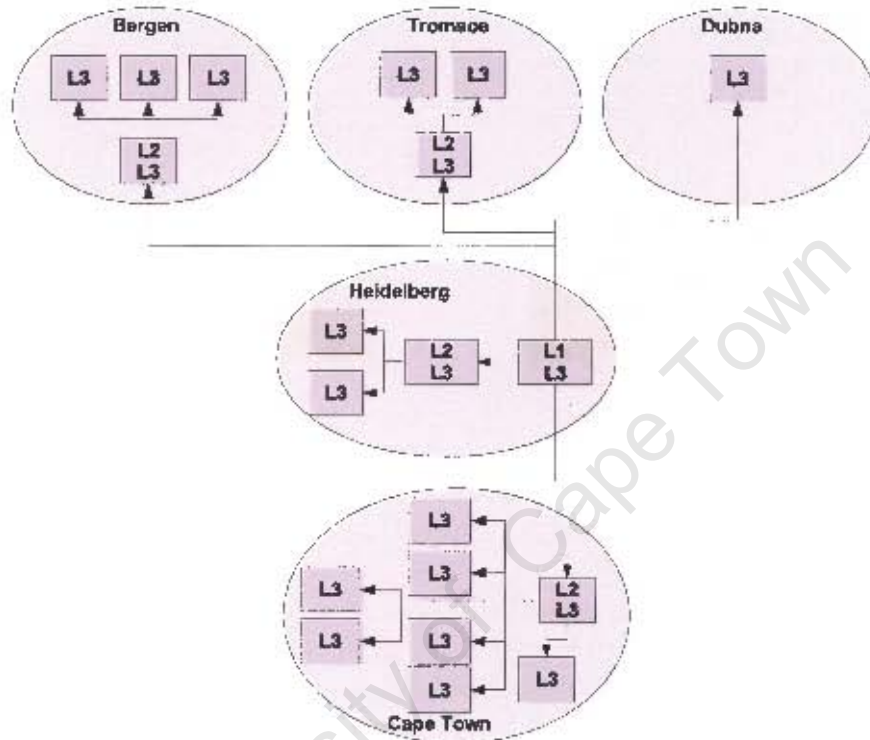


Figure 6.10: The hierarchy of the global test *TaskManagers* at the various sites. Each L3 *TaskManager* slave is controlled by the closest L2 *TaskManager* servants. In the case of Dubna, where there was not a need for an intermediate relay/access node, the L3 slave *TaskManager* is controlled by the L1 Master in Heidelberg.

was given to a top-level (“Master”, or L1) *TaskManager* running in Heidelberg. This top-level *TaskManager* communicated with intermediary (“servant” or L2) *TaskManagers* in Cape Town and Bergen, where relay/access nodes were in place. The third level *TaskManagers* (“Slave” or L3) performed the low-level tasks that the *TaskManager* component was originally designed for, such as initiating and synchronising the HLT’s processing and communication components. In Dubna, where there was direct access to the compute node (no relay/access node), the L1 *TaskManager* in Heidelberg communicated directly with the L3 *TaskManager*.

These hierarchical layers of *TaskManagers* simplified immensely the configuration of the test. It was possible to start and stop the system by issuing a single instruction

to the L1 *TaskManager*, without having to manually initiate every level in the HLT processing chain.

6.6.3 Global test results

At the most basic level, the global HLT test can be termed a success. Referring the aims of the test given above, the primary aim was achieved in that the system was built and run for several hours uninterruptedly. Figure 6.6.3 shows the recorded network activity on the CARMEN head node, showing that there were no lapses in communication or network failures.




Figure 6.11: Network input rates to the CARMEN head node during the global framework test, for the full duration of the run, and during one hour.

The periodic spikes in the network activity represent the transmission of data per event. The event rate was not set by hand as in the dummy stress test, as the aim was to see what rate was achievable. During the global test, a fairly stable rate of about 10-15 Hz was achieved, with the system running for several hours. This very low rate was to be expected due to the bottleneck into UCT, but the system proved nonetheless very stable.

6.7 Next steps

The first integration test of the dHLT components described above had a very limited scope, but achieved its objectives. Many of the components of the full dHLT system were not available and had to be mocked up for the test, but the integration of these is currently under way.

The offline dHLT track-finding code is being integrated into the ALICE offline simulation framework, AliRoot. This is part of the full HLT integration into the offline framework and is necessary in order to test the HLT processing during the ALICE Data Challenge.

As regards the rest of the development of the dHLT, this is somewhat outside of the scope of this thesis, as it is a work in progress. The definitive versions of the track-finding and hit-reconstruction algorithm, as well as the decision regarding the

implementation platform (FPGA vs CPU) will be decided during the following dHLT integration tests, which are planned to occur at least once a year.

We have shown in the last two chapters how the dHLT is built on a larger software framework, the ALICE HLT, in order to improve the physics potential of the spectrometer by improving the accuracy of the trigger information. The design principles and constraints have been discussed, and the specific implementation in both the online and offline case has been presented, with validation of the performance of the system in an artificial physics environment. The simulations used to test the performance of the system were done with muons of flat p_T spectra, leaving any actual physics considerations aside. In the next chapter we will focus making first investigations into the performance of the dHLT in a more realistic physics environment, using offline simulations.

University of Cape Town

Chapter 7

Physics performance of the dHLT

7.1 Introduction

It has been shown in the previous chapters that the ALICE dimuon spectrometer is specifically designed to measure dimuons with good accuracy, and the dimuon High Level Trigger is able to discriminate with great efficiency against the low- p_T background muons compared to L0. In order to perform these studies described in the preceding chapters however, the input spectra of single muons was not that which is expected from actual physics events at ALICE. Rather, the spectra which were passed through the dHLT were intentionally forced to be flat in p_T , allowing to determine the accuracy and efficiency of the trigger across the entire p_T range expected at the LHC, albeit with unrealistic error estimates.

Of course, the spectra of muons from heavy quarkonia decay do not have a constant p_T distribution. Instead, these exhibit peaks around 4-5 GeV and 2 GeV for the Υ and J/ψ respectively (see figure 7.1), and have tails which extend far out into the high- p_T region. The distribution of background muons, on the other hand, is very different – mainly below 2 GeV (see figure 7.3). This chapter aims to investigate the response of the dHLT in a “realistic” physics environment. “Realistic” here refers to the expected phase-space distributions of particles produced at the LHC, but not necessarily to the yields, since such a study would require much larger computing resources than were available.

Recall that, as discussed in chapter 2, a clear signal for the creation of the QGP was predicted to be the suppression of heavy quarkonia, $c\bar{c}$ (J/ψ) and $b\bar{b}$ (Υ), due to Debye screening, measured in the dimuon channel. The p_T spectra of the dimuons from the decay of these mesons is predicted to be heavily influenced by the presence and characteristics of the QGP and one model [174] will be used to predict their subsequent spectra. The NA50 (and consequent NA60) experiment had great success in measuring the anomalous suppression of the J/ψ yield in the dimuon channel [112]. The suppression in the yield relative to “normal” nuclear suppression was accounted for by hadronic absorption models [85], except in the most central Pb-Pb collisions, where an anomalous suppression was claimed to have been observed [112] (see section 2.4.6). At the higher energy LHC experiments, however, theoretical

predictions for the suppression of quarkonia are not as clear as at the SPS. For instance, the gluon shadowing and “normal nuclear” effects (including nuclear and co-mover absorption) on the yield and spectra of the resonances is not completely understood at LHC energies. The NA50 analysis also made use of the medium-independent Drell-Yan dimuon production mechanism, which was used to normalise the quarkonia production rate. However, for the dimuon spectrometer, there exists no single reliable, medium-independent reference for the quarkonia production rate. For these reasons, it is not as simple to identify a single “smoking gun” signal for the QGP by measuring quarkonia yields, as was the case at the SPS.

One way to reduce the amount of uncertainty is to investigate the p_T -dependence of the *ratios* of quarkonia, which may give not only insight into the existence, but also the nature of the produced medium. A detailed study of the performance of the dimuon spectrometer in detecting this signal can be found in [172], where it was shown that the signal can indeed be measured with a realistic description of the detector. Although the aforementioned work represents a relevant case study of interest to the performance of the dHLT, it was performed in a fast-simulation framework and as such is not easily possible to investigate its results with the dHLT, as the latter requires more detailed information (see section 7.2.2). This thesis will focus rather on the more basic investigation of the dHLT’s performance in some well-defined cases, in order to serve as a possible benchmark to further work.

7.1.1 Selection of data sets

Four general sets of data will be considered :

- Pure background in the absence of signal
- Pure signal in the absence of QGP
- “Cocktail” events with mixed signal and background
- QGP-suppressed signal (for Υ suppression scenarios.)

Note that due to the excessive computing time necessary a full-statistics data set, only fractions of the expected number of events were generated, but nevertheless enough to give a reasonable result. The actual fraction of events generated compared to that expected at the experiment is given in table 7.3. The production rates for non-QGP production using parametrised p_T distributions will be shown for the $c\bar{c}$ and $b\bar{b}$ families. The response of the dHLT to this signal will be shown and compared to that of the LO, in order to validate the accuracy and efficiency of the dHLT in a semi-realistic environment. The response of the dHLT to a realistic background distribution, with no signal will also be shown, separately. For this, a full simulation of muonic background sources is done with HIJING and the response of the dimuon High Level Trigger (dHLT) is shown in order to validate the background subtraction. Furthermore, a set of “mixed” events - containing the contributions at the relevant levels of all sources of muons combined - was generated and passed through the dHLT;

this represented a sample of events as close as possible to the actual ALICE dimuon minimum-bias.

Finally, a set of data representing $b\bar{b}$ production in the two QGP scenarios used in reference [172] was generated. The parametrisations of the meson suppression as a function of p_T was adapted from reference [173] and coded into the ALICE offline simulation code (see section 7.5). Finally, note that the data sets simulated here will all be for the most central class of collisions at ALICE, which correspond to an impact parameter of 0-5 fm [141].

7.2 Baseline data production (no QGP)

The first parameter needed for this study is the yield of each of the observables which is to be simulated. The expected yield of an observable in Pb-Pb collisions can be obtained from the same observable in p-p collisions by applying the relevant nuclear modification factors. However, in order to obtain the relevant yield in p-p collisions, a number of basic input parameters are required, the first of which is the direct quarkonia production cross-section, σ_{pp}^{qq} . To this is applied the feed-down factor, F_d , which takes into account the enhancement of lower-lying states due to decay of higher-lying states. This provides a good estimate of the expected yields in p-p collisions.

In order to scale up the p-p yield, the assumption is made that the cross-section for quarkonia production in Pb-Pb collisions scales with the number of binary (hard) collisions, N_{bin} (see equation (2.16)). For the most central collisions which we are investigating here, this is approximately¹ given by A^2 , where A is the number of nucleons in the nuclear projectile – 208 in our case of Pb-Pb. To this, we have to finally apply the shadowing factor C_{sh} , as described in section 2.4.6. This is parametrised as a function of impact parameter b as

$$C_{sh}(b) = C_{sh}(0) + (1 - C_{sh}(0)) \cdot \left(\frac{b}{r_0}\right)^n, \quad (7.1)$$

where r_0 and n are chosen to be $r_0 = 16$ fm and $n = 4$ respectively.

Applying all of the relevant scaling factors, the cross-section for quarkonia production in Pb-Pb collisions is then

$$\sigma_{Pb-Pb}^{qq} = C_{sh} \cdot F_d \cdot A^2 \cdot \sigma_{pp}^{qq} \quad (7.2)$$

The relevant parameter for the study is not the cross-section, however, but the total number of quarkonia $N_{q\bar{q}}$ expected to be produced during an ALICE data-taking run. This is obtained by applying the generic formula used to calculate the number of any observable with cross-section σ , given the expected average luminosity $\langle L \rangle$ (provided by the LHC) and the total data taking time expected², T :

$$N = \sigma \cdot \langle L \rangle T. \quad (7.3)$$

¹This is the case at $b = 0$ fm

²This number takes into account the expected efficiencies of the LHC accelerator, and the ALICE experiment [141].

Upsilon family	Υ	Υ'	Υ''
$F_d \sigma_{pp}^{41\mu\bar{\nu}} (\mu\text{b})$	0.501	0.246	0.0181
$\sigma_{AA}^{bb} = C_{sh} \cdot F_d \cdot A^2 \cdot \sigma_{pp}^{bb} (\text{mb})$	16.47	8.09	0.59
$N_{AA}^{bb} = \sigma_{AA}^{bb} \cdot \langle L \rangle \cdot T$	$8.24 \cdot 10^6$	$4.04 \cdot 10^6$	$2.98 \cdot 10^5$
$N_{AA}^{\mu^+\mu^-} = Br_{(bb \rightarrow \mu^+\mu^-)} \cdot N_{AA}^{bb}$	$2.04 \cdot 10^5$	$5.23 \cdot 10^4$	$5.386 \cdot 10^3$
$N_{\text{spectro}}^{\mu^+\mu^-}$	9792	2510	258

Table 7.1: Cross-sections in proton-proton collisions (σ_{pp}) and nucleus-nucleus collisions (σ_{AA}) for the lowest-lying $b\bar{b}$ family resonances [142] (including feed-down factors). Also shown is the resonance yield and number of dimuons from resonance decay expected in the data taking year, as well as the number of these dimuons expected in the spectrometer acceptance.

J/ψ family	J/ψ	ψ'
$F_d \sigma_{pp}^{\mu^+} (\mu\text{b})$	31.0	4.68
$\sigma_{AA}^{c\bar{c}} = C_{sh} \cdot A^2 \cdot \sigma_{pp}^{c\bar{c}} (\text{mb})$	804.7	121.5
$N_{AA}^{c\bar{c}} = \sigma_{AA}^{c\bar{c}} \cdot \langle L \rangle \cdot T$	$4.0235 \cdot 10^8$	$6.074 \cdot 10^7$
$N_{AA}^{\mu^+\mu^-} = Br_{(c\bar{c} \rightarrow \mu^+\mu^-)} \cdot N_{AA}^{c\bar{c}}$	$2.365 \cdot 10^7$	$6.256 \cdot 10^5$
$N_{\text{spectro}}^{\mu^+\mu^-}$	828 047	21 897

Table 7.2: Cross-sections in proton-proton collisions (σ_{pp}) and nucleus-nucleus collisions (σ_{AA}) for the lowest-lying $c\bar{c}$ family resonances [142] (including feed-down factors). Also shown is the resonance yield and number of dimuons expected from resonance decay expected in the data taking year, as well as the number of these dimuons expected in the spectrometer acceptance.

The nominal values for the LHC luminosity and heavy-ion data taking period per year are $\langle L \rangle = 5 \cdot 10^{26} \text{cm}^{-2} \text{s}^{-1} = 5 \cdot 10^2 \text{barn}^{-1} \text{s}^{-1}$ and $T = 10^6 \text{s}$ [141].

The dimuon spectrometer is tuned to detect dimuons and in order to obtain the number of dimuon pairs from the quarkonia yield, the branching ratio $Br^{q\bar{q} \rightarrow \mu^+\mu^-}$ of the quarkonia decay into dimuons has to be applied to the net number of mesons :

$$N_{P_b P_b}^{q\bar{q} \rightarrow \mu^+\mu^-} = \langle L \rangle \cdot T \cdot Br^{q\bar{q} \rightarrow \mu^+\mu^-} \cdot \sigma_{P_b P_b}^{q\bar{q}} \quad (7.4)$$

This gives for example $N_{J/\psi} \sim 10^6$ and $N_{\Upsilon} \sim 10^4$ for the heavy-ion data taking year. The values expected for the Υ and ψ resonances are given in tables 7.2 and 7.2 respectively, for the full solid angle (4π), as well as the spectrometer acceptance. The number of dimuons actually detected is obtained from the total number of quarkonia by scaling down by the various detector effects of the dimuon spectrometer, due to inefficiencies in track reconstruction, triggering and geometrical acceptance (see section 3.2.3). Table 7.3 shows the number of events generated for each of the data sets described above and the fraction compared to what is expected during data taking. Note that the simulated fraction of the ‘‘signal’’ data set compared to what is expected in true experimental conditions is far higher than that of the underlying background events. This is due only to the fact that since the number of tracks in each

Data set	Yield (events)	Simulated ($b < 5$) (events)	Ratio
Background	$\sim 1 \cdot 10^9$	10^3	$1 \cdot 10^{-6}$
J/ψ	$\sim 2.365 \cdot 10^7$	$5 \cdot 10^5$	$2.1097 \cdot 10^{-2}$
Υ	$\sim 2.5 \cdot 10^5$	$5 \cdot 10^5$	~ 0.5
Υ (SU(N))	185 000	$4 \cdot 10^5$	2.16
Υ (SU(3))	409 000	$5 \cdot 10^5$	1.22

Table 7.3: Statistics of the various data sets used in the dHLT simulation

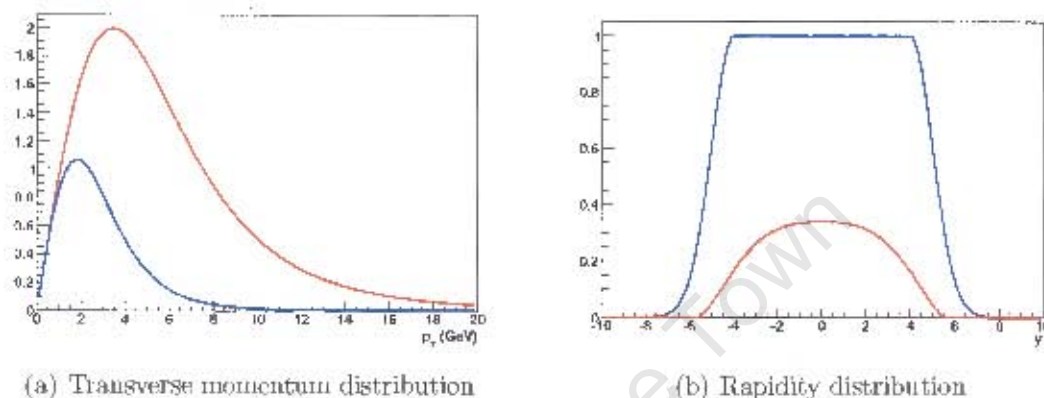


Figure 7.1: Transverse momentum and rapidity parametrisations of for the J/ψ (blue) and Υ (red) mesons, for Pb-Pb collisions at 5.5 TeV. (Curves scaled to match CDF data and are relative to the J/ψ production.)

event of the background data set is of the order of 10^3 , while that of the signal events is only of the order of 10. Thus, the computing time required to process these events is orders of magnitude larger. However, since the dHLT is a track-based algorithm, the sheer number of tracks in the background data set means that an accurate estimate of its performance in this scenario can be given.

For the case of the QGP scenarios, it should be noted that slightly more quarkonia were simulated than are actually expected. This means that the statistical uncertainty on the accuracy of the dHLT performance for these sets will be somewhat over-estimated.

7.2.1 Rapidity and transverse momentum spectra

Although we have defined the total number of quarkonia expected during the ALICE data-taking year, the phase-space distribution of these has not yet been discussed. Knowledge of the p_T and y spectra are necessary in order to simulate their production correctly, and hence the correct response of the detector and trigger.

The p_T distribution for the Υ and J/ψ mesons are parametrised as

$$\frac{d\sigma}{dp_T} = \frac{p_T}{\left(1 + \left(\frac{p_T}{p_T^0}\right)^2\right)^n}, \quad (7.5)$$

The parameters p_T^0 and n are chosen to be $p_T^0 = 4.703$ (7.653) GeV and $n = 3.826$ (3.042). These are obtained by a fit to the MRST HO calculations made by R. Vogt [168], and scaled to match the recent CDF measurements.

MRST HO calculations made by R. Vogt [168], scaled to match the recent CDF measurements. These are shown in figure 7.1 for the J/ψ and Υ states.

7.2.2 Simulation and analysis frameworks

The above parametrisations of particle production are included in the ALICE offline software, AliRoot, which makes it possible to use them to generate physics events to study with the detector. This is done in two ways : a parametrised description of the detector acceptance, response function, efficiencies, etc. is included in the ALICE FASTSIM module in the first case, while in the second case, the full simulation of the detector response at a microscopic level is brought to bear. These two frameworks, denoted *fast* and *full* simulation respectively, are both important to study of physics expectations of the ALICE detector.

The fast simulation framework contains pre-defined look-up tables for the acceptance, response and efficiency of the detector, given a particle momentum. This information is also parametrised as a function of the background level, allowing the user to vary this factor easily, as there is some uncertainty in average event multiplicity at the LHC (see section 2.4.1 for a discussion of this issue). In order to obtain good statistics over the full p_T range, the concept of a “pseudo-particle” is used, which is assigned a random³ p_T , with which the acceptance and tracking or trigger efficiency of the detector is obtained. This is then assigned a weight, which is calculated from a parametrised p_T spectrum of the muon source in question.

In the slow, or full simulation environment, a detailed detector geometry, segmentation and response are used. Primary particles are generated using the desired p_T and y distributions described above, which are allowed to decay with an external decayer, usually provided by PYTHIA [147]. The generated particles are transported using the GEANT3 Monte-Carlo code [148] with the detector description in the MUON module of AliRoot. This method processes data on a particle-by-particle basis, including the reaction mechanisms specified in the initial configuration. This simulation method is thus much slower, but also much more detailed than the fast simulation⁴ containing much the same kind of physics information as would exist in the actual experiment, as well as the input Monte-Carlo. This track-based simulation method is required by the offline simulation of the dHLT, since it processes individual tracks on

³The value is drawn from the desired p_T distribution.

⁴In fact, the parametrisations coded in the fast simulation package were indeed obtained by fits to the results of the full simulation.

an event-by-event basis, in order to investigate the expected efficiency and resolution of its track-finding algorithm.

The two simulation techniques are mutually compatible and give quite the same result for several observables. A comparison between the efficiencies of the fast and slow simulations is given in the ALICE PPR [141], where it is shown, for example, that the overall efficiency (including triggering and tracking efficiency) of the fast simulation for the J/ψ meson at low background levels (20% of the nominal level) is 64%, while that of the full simulation is 67%. At the nominal background level, these numbers drop slightly to 66% and 70% respectively. The Υ efficiencies show a similarly slight discrepancy between the fast and full simulation techniques. In terms of mass resolution (which is directly related to the p_T resolution), the fast simulations tend to over-estimate the resolution of the detector at high background levels, especially for the J/ψ but the disagreement is never more than a few percent.

7.3 Background sources discussion

Since the dimuon spectrometer detects almost exclusively muons, we will now investigate the sources of muons which may contribute to the background expected to contaminate the signal identified by the dHLT. As we have seen in chapter 2, several mesons have decay channels involving muons. The most notable of these are the open charm and beauty mesons (D and B family), since their decay products have a high enough p_T to satisfy the transverse momentum cut. These high- p_T muons contribute to the background dimuon invariant mass spectrum, and can usually be subtracted using fairly standard techniques [170]. However, the dHLT does not perform an invariant mass calculation⁵ and so only the p_T spectra of muons from these sources is shown here. However a short discussion of the characteristics of these sources – separated roughly into *correlated background* and *non-correlated background* – follows below, in order to give a context.

Correlated background

The term “correlated” refers to muons which are physically correlated, but not originating from heavy quarkonia decay. This is particularly the case, as mentioned above, in open charm- and beauty decays. Here, a heavy quark forms a meson by combining with lighter quarks, which then decays into muons :

$$\begin{aligned}
 (c\bar{d}) \quad D^+ &\rightarrow \mu^+ + \nu_\mu + X \\
 (\bar{c}d) \quad D^- &\rightarrow \mu^- + \bar{\nu}_\mu + X \\
 (b\bar{d}) \quad B^+ &\rightarrow \mu^+ + \nu_\mu + X \\
 (\bar{b}d) \quad B^- &\rightarrow \mu^- + \bar{\nu}_\mu + X
 \end{aligned}
 \tag{7.6}$$

The invariant mass spectrum of the correlated background has a structure which reflects the sources of the muons. The D meson is one of the most common products

⁵This is under investigation in later implementations of the dHLT. See [160] for details




Figure 7.2: Transverse momentum spectra of the muons from charm (blue), beauty (pink) and light hadronic (π , K mesons - red) decays in Pb-Pb collisions at 5.5 TeV. The sum of all contributions is also shown in black.

of the B decay chain. The muon from the subsequent D decay which, combined with that of the B decay, contributes to the shoulder at intermediate invariant mass in the correlated background spectrum.

In the high- $M_{\mu^+\mu^-}$ range (from about 6 GeV/ c^2), the spectrum is well-described by an exponentially decreasing function of the form

$$\frac{dN}{dM} \sim A \cdot e^{-bM}, \quad (7.7)$$

where the slope b is about 0.611 GeV $^{-1}/c^2$.

Non-correlated background

Non-correlated background is defined as decay muons which do not originate from a physically correlated source and is usually split into four groups:


charm-charm where distinct, unrelated charmed mesons decay into muons

beauty-beauty where distinct, unrelated beautiful mesons decay into muons

decay-decay where distinct light mesons (pions or kaons) decay into muons

charm-decay where muons from charmed mesons are combined with decay muons from light mesons.

The transverse momentum spectra of these muon background sources, calculated with the AliRoot FASTSIM module, is shown in figure 7.3 (note that the beauty-decay contribution is negligible).



(a) Background p_T spectra of various sources (see text). (b) The resulting dimuon mass spectrum ($M_{\mu^+\mu^-}$) for the background sources discussed in the text.

Figure 7.3: Transverse momentum (left) and invariant mass spectra of background muon sources.

7.4 Detailed simulation of the background muon sources

A set of data corresponding to about 10^{-6} of the expected yield of all muons for events in the most central ($b < 5$ fm) class was generated using the so-called “muon cocktail” - a mixture of the relevant background sources described above. Due to the high multiplicity of minimum-bias events, however, only a small number of events is necessary to get good statistics from the dHLT processing results. The subsequent p_T spectra and invariant mass plot $M_{\mu^+\mu^-}$ are shown in figure 7.3(a) and 7.3(b) respectively.

7.4.1 Result of dHLT processing

Once the background spectra had been generated and the events reconstructed, they were passed through the dHLT filter. The first interesting aspect of the result was the momentum reconstruction accuracy obtained with the dHLT, compared to that of L0, which is shown in figure 7.4. From this figure, it can be seen that the dHLT p_T estimation has a slightly higher resolution than that of the L0 algorithm - a difference of about 100 MeV, or about 30 %.

Also noticeable is a change in the shape of the triggered spectra (both L0 and dHLT) compared to the input spectrum. Since the triggered spectra are built from hits detected on the last two tracking and trigger stations, this change in shape can be attributed to both particle energy loss (mainly due to the presence of the front absorber) and multiple scattering in the detector and muon filter. Indeed, from figure 7.4, we see that there is a systematic shift between both the L0 and the dHLT p_T estimation compared to the input spectra, of about 100 MeV/c. This is a general feature of the triggered spectra and is noticeable in all of the data sets shown here, but affects muons below about 1 GeV more significantly.

Figure 7.4.1 shows the input p_T spectra with a cut at 1 GeV on the L0 and dHLT output. From this figure, it can be seen that although the L0 cut reduces a large

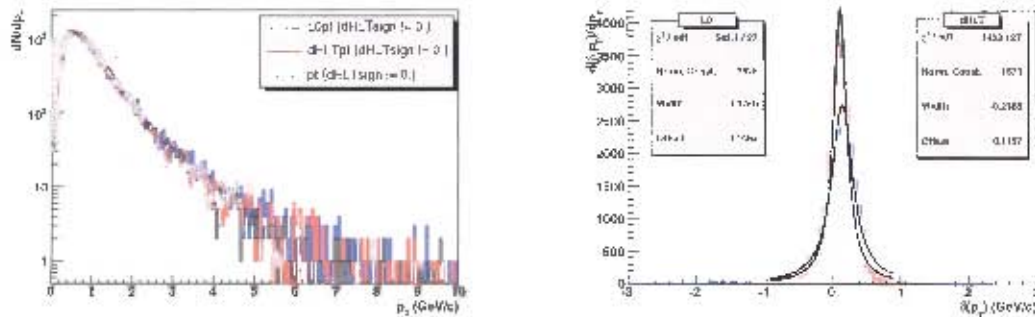


Figure 7.4: **Left:** Transverse momentum spectra for the pure background data sample, as well as the corresponding dHLT and L0 output data samples (charged tracks only). **Right** Reconstructed transverse momentum resolution, for the muon cocktail background sample (see text).

portion of the background, the cut is not as sharp as that of the dHLT. This can be better seen in the plot on the right hand side of figure 7.4.1, where the spectra with a 1 GeV cut on the L0 and dHLT respectively are divided by the raw input spectra. The resulting spectrum, referred to as the “efficiency spectrum” of the trigger algorithm, is parametrised by an error function with a width, a cut value, and a normalisation constant (see equation (5.2)). The smaller the width of the curve, the more accurate the p_T cut is.

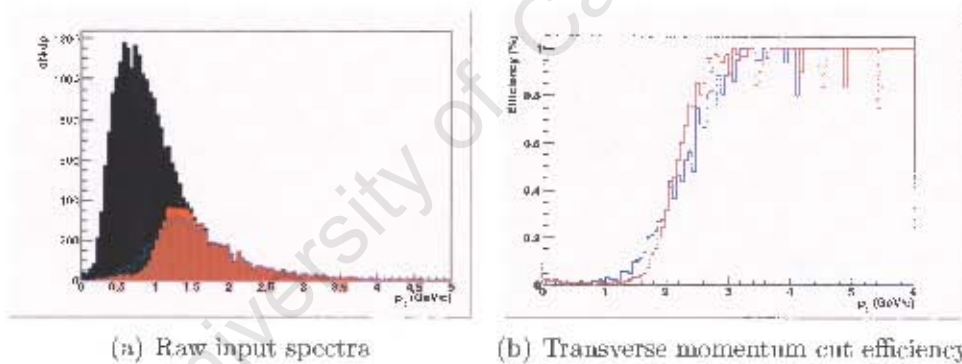


Figure 7.5: Raw input (black), L0 (red) and dHLT (blue) p_T distributions and corresponding efficiency functions for the background data sample with a 1 GeV cut.

Several of the tracks which were generated did not produce hits on the last 2 tracking stations, but of the tracks which did in fact produce hits on station 4 and 5, practically 100 % are found by the dHLT tracking algorithm. The total acceptance of the dHLT is about 79 % of generated tracks (a similar, slightly lower value is found for the L0). Once the 1 GeV cut is applied, the dHLT succeeds in reducing the number of accepted tracks to 21.2 % of the originally reconstructed tracks.

Both of the efficiency functions are fitted with similar fitted cut and normalisation values (“const” and “mean” respectively in the figure), within errors. There is a clear distinction however between the fitted skin widths of the dHLT and L0; a roughly 50 % decrease in cut width is observed. The background rejection factor here is




Figure 7.6: Fitted efficiency functions for a 1 GeV p_T cut on the background event sample. Note that the fitted width of the dHLT function is roughly half that of the L0 width.

encouraging and validates one of the mandates of the dHLT. The further mandate – to faithfully identify the signal – is discussed below.

7.4.2 Result of dHLT processing on signal

The aim of producing signal-only data sets is to verify the performance of the dHLT in terms of signal identification. The results of the dHLT processing of the J/ψ and Υ data sets is shown below. The results for these cases will be shown separately, since they involve different p_T cuts - 1 GeV and 2 GeV respectively.

Performance of the dHLT on the J/ψ sample

The simulated J/ψ data set is summarised in figure 7.7, where the resulting transverse momentum and invariant mass spectra are shown.

As in the case of the background data set, the raw data was passed through the dHLT and L0 algorithms. The resulting transverse momentum spectra of triggered tracks was then divided by the raw input spectrum to give the efficiency plots shown in figure 7.8. The percentage of various trigger conditions is shown in table 7.4.

It can be seen that the width of the dHLT efficiency function is smaller than that of the L0 curve : 0.26 GeV as opposed to 0.38 GeV. However, the error on the fit of the skin width is rather large (of the order of 100 %), so it is not clear with the pure J/ψ data set whether the dHLT has a particular advantage in terms of cut efficiency over the L0 algorithm. This is not entirely surprising, since indeed the L0 trigger is optimised to detect heavy quarkonia such as the J/ψ . The efficiency of the trigger for the detected tracks is about 84 % for the L0 1 GeV/c cut and about 87 % for dHLT

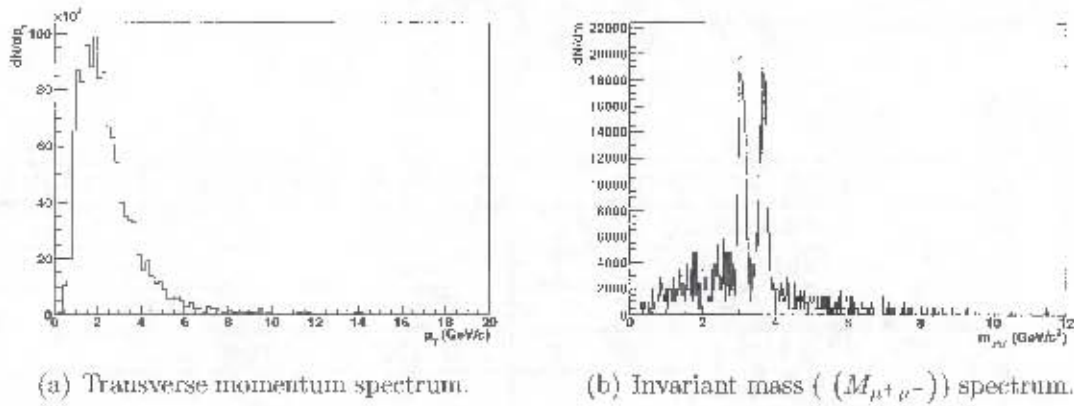


Figure 7.7: Summary plot for the J/ψ data set.

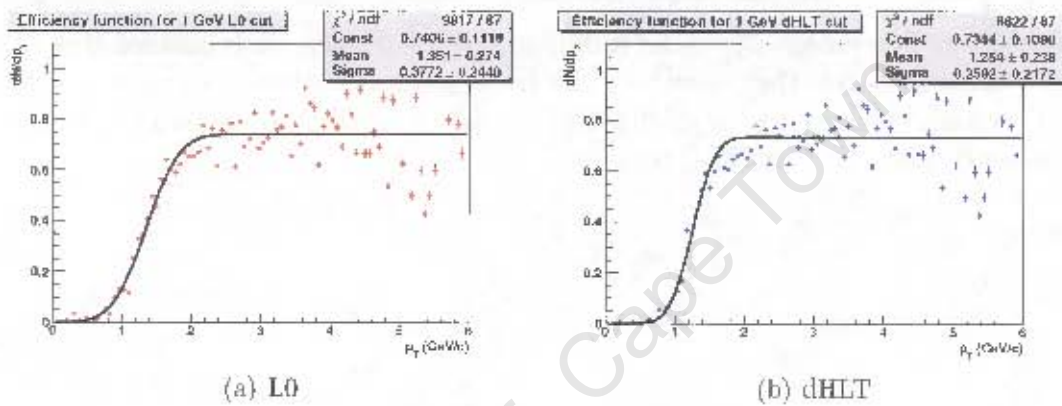


Figure 7.8: The L0 and dHLT efficiency spectra as a function of p_T , for a 1 GeV p_T cut, fitted to an error function for the J/ψ data set.

1 GeV/c cut(with no L0 cut). Here, we show that for this sample, the mandate to select good tracks is shown to be fulfilled, albeit in an averaged way. It should also be noticed that in roughly 5.02 % cases, the dHLT would have validated a J/ψ muon track which was rejected by the L0.

Performance of the dHLT on the Υ sample

Since the transverse momentum spectrum for the Υ family has a peak at a higher p_T value, the trigger cut for this data sample, is at 2 GeV/c, as described in chapter 3. The invariant mass and transverse momentum spectra generated for this data sample are shown in figure 7.9.

The efficiency of the L0 and dHLT trigger algorithms on this set is shown in figure 7.10, where it can be seen that the fitted width of the L0 histogram is wider than that of the dHLT width (about 0.5 GeV for the L0 vs about 0.3 for the dHLT). This represents about a 40-50 % improvement, but since the uncertainty on the fitted parameter is quite large, the significance of this should not be overstated.

Again, the L0 and dHLT both reconstruct a very good percentage of the tracks -

Trigger Condition	No. tracks	Percentage of total (%)
Baseline	1059069	100
L0 low p_T	880881	83.17
dHLT low p_T	909621	85.88
L0 high p_T	521152	49.20
dHLT high p_T	544144	51.86
Asymmetric (L0 low, dHLT high)	540791	51.06
dHLT low, no L0 low	53169	5.02

Table 7.4: Percentage of tracks fulfilling various trigger selections for the J/ψ . The baseline criterion is the case where the track results in a hit on stations 4 and 5 and a hit on each of the trigger stations. In the case of a dHLT trigger with now L0 trigger, the percentage is of total L0 triggers.

Trigger Condition	No. tracks	Percentage of total (%)
Baseline	1216233	100
L0 low p_T	1206208	99.17
dHLT low p_T	1209416	99.43
L0 high p_T	1164504	95.74
dHLT high p_T	1178539	96.90
Asymmetric (L0 low, dHLT high)	1178539	96.90
dHLT low, no L0 low	4812	0.39

Table 7.5: Percentage of tracks fulfilling various trigger selections for the Υ data sample. The baseline criterion is the case where the track results in a hit on stations 4 and 5 and a hit on each of the trigger stations. In the case of a dHLT trigger with now L0 trigger, the percentage is of total L0 triggers.




Figure 7.9: **Left:** The invariant mass spectrum of the Υ family data sample generated. **Right:** The transverse momentum spectrum of the Υ data sample.




Figure 7.10: **Left:** The L0 efficiency function as a function of p_T with a 2 GeV cut for the Upsilon data sample. **Right:** The dHLT efficiency function for the same cut on the Upsilon data sample. Both have been fitted to the error function (see text). Note that the fitted width of the L0 histogram is larger than that of the dHLT width.

about 95 % for both triggers. It should also be noted that in a very small percentage (roughly 0.4 %) of the tracks, a the dHLT would have validated a track which the L0 rejected. In this case, it is considered a negligible effect.

Clearly, in the signal cases above, we are using a scenario which is entirely ideal, since it is extremely improbable to experimentally produce purely quarkonia in heavy-ion collisions. What we have shown is that the dHLT does not reject significant amount of good tracks, and increases the efficiency of the p_T cut over that of L0. In this sense, we show that the dHLT validates in part its second mandate, which is to faithfully reconstruct relevant physics tracks. In the next section, we will investigate a more realistic scenario of physics events with a mixture of all sources of muons, both from signal and background sources.

7.4.3 Mixed events

We have seen that the dHLT can perform well in either a purely background or pure signal scenario. The next data set to be investigated was generated using the same sources of background muons described above, however contributions from open charm, open beauty and the heavy quarkonia families were also added. The centrality

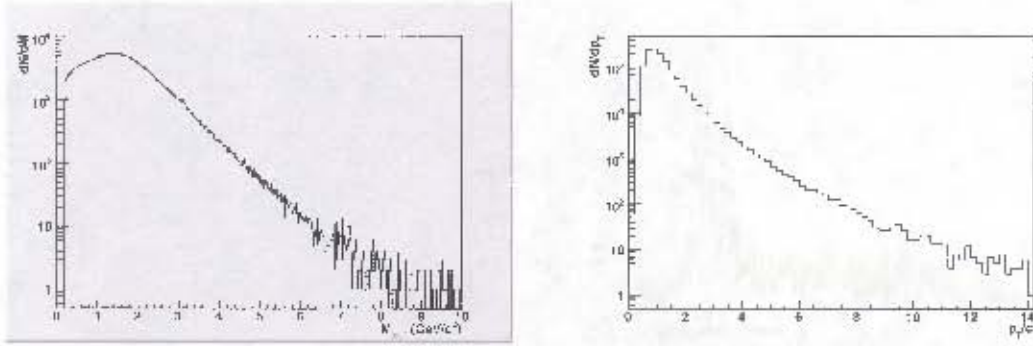


Figure 7.11: Invariant mass (left) and transverse momentum spectra (right) of the full muon cocktail data sample.

dependence of the open charm and beauty shadowing was taken into account using a Glauber model to obtain the average impact parameter. Using the same Glauber model, the production rate of ψ -onium species was calculated from the average number of collisions, while the production of low- p_T background muon sources (π and K mesons) was calculated with the average number of participants. Again, the invariant mass and transverse momentum spectra of this mixed event sample are shown in 7.11.

When this data sample is passed through the dHLT filter, we expect to suppress a significant amount of the background at low- p_T , while retaining good statistics for the quarkonia signals. This is a validation of the preceding two data sets, in the case of signal only and the first data set, in the case of background only studies. The results of the dHLT and L0 triggering are shown in figure 7.12 and in table 7.6

Trigger Condition	No. tracks	Percentage of total (%)
Baseline	24504	100
L0 low p_T	8040	32.81
dHLT low p_T	7982	32.57
L0 high p_T	2057	8.39
dHLT high p_T	1810	7.38
Asymmetric (L0 low, dHLT high)	1781	7.26
dHLT low, no L0 low	1481	18.42

Table 7.6: Percentage of tracks fulfilling various trigger selections for the “minimum-bias” data sample. The baseline criterion is the case where the track results in a hit on stations 4 and 5 and a hit on each of the trigger stations. In the case of a dHLT trigger with now L0 trigger, the percentage is of total L0 triggers.

From these, we can see that the shape of the efficiency spectrum is not much changed from the previous two cases. Although this data set is a combination of the previous signal (J/ψ and Υ) and background data sets, clearly in such a minimum-bias scenario, background contributions below $p_T = 1$ GeV/c dominate to a large extent. It should also be noted that the dHLT would trigger roughly 18.4 % of the tracks which are not triggered by the L0.




Figure 7.12: **Left:** Efficiency as a function of transverse momentum of the 1 GeV L0 cut for the mixed data sample. **Right:** Efficiency as a function of transverse momentum of the 1 GeV dHLT cut for the mixed data sample. Efficiency functions have been fitted to the error function described in chapter 5.

We have now investigated the three main control scenarios - pure background, pure signal and mixed background and signal data sets - and the dHLT has been found to perform its required task to within specifications. The standard, hadronic production only, case has been verified and we now move on to investigate the performance of the dHLT in a speculative scenario of QGP production.

7.5 The dHLT with a quarkonia suppression model

Having now investigated the validity of the dHLT in several somewhat realistic, but still fairly artificial scenarios, we would like to extend our physics beyond purely hadronic scenarios. A simple QGP model will be used to investigate what effect modifications on the p_T spectrum due to this model would have on the dHLT response. The effects of the model in question have already been shown with fast simulations to be measurable by the dimuon spectrometer [173], where it was shown that by using ratios of quarkonia yields as a function of p_T , it was possible to disentangle various scenarios of QGP formation. A short theoretical background of this model will be given to derive the expressions of the modified p_T spectrum for quarkonia, and thereafter the response of the dHLT to a set of data generated with these p_T distributions will be shown.

Context and background

In the late 1980's, Blaizot and Ollitrault [174] published a calculation of the effect of the temporal and spatial evolution of a produced QGP on the survival of the produced quarkonia, which allowed them to relate the resonance spectra with the lifetime of the QGP. Kharzeev and Satz used this model to predict the p_T dependence of the J/ψ and ψ' suppression at RHIC and LHC energies [175]. Using updated gluon distribution function calculations, this study was extended by Gunion and Vogt, who calculated the p_T dependence of the ratio of resonance families, $\frac{\psi'}{J/\psi}$ and $\frac{\Upsilon'}{\Upsilon}$. As observables, these resonances provide a number of benefits:

- Gluon shadowing would have no effect on the ratio $\frac{\psi'}{J/\psi}$ or $\frac{\Upsilon'}{\Upsilon}$, since it would affect the production of each identically.
- It is safe to assume that the p_T variation of the hadronic absorption cross-sections of Υ' and Υ are the same, and therefore the ratio of these would cancel out the contribution from the poorly-understood (at LHC energies) hadronic absorption. While the global value of the ratio would be affected, its variation with p_T would not.
- No p_T variation in the ratio of $\frac{\Upsilon'}{\Upsilon}$ has yet been observed in $p - \bar{p}$ collisions. The CDF collaboration at Fermilab has measured this ratio out to about 10 GeV and their results [176] are consistent with no variation. If any variation is seen in Pb-Pb collisions, this may be attributed to medium effects.

This model will be used here to simulate a set of data for QGP suppression scenarios.

7.5.1 Model description of the nature of the produced QGP

Recall that in zero-temperature field theory, equation (7.8) gives the non-relativistic inter-quark potential:

$$V(r, 0) = \sigma(r) - \frac{\alpha}{r}, \quad (7.8)$$

where r is the inter-quark separation. The screening of the inter-quark potential which is responsible for the dissociation of bound states in a deconfined medium modifies this potential which then becomes

$$V(r, T) = \frac{\sigma}{\mu(T)} (1 - e^{-\mu(T)r}) - \frac{\alpha}{r} e^{-\mu(T)r}. \quad (7.9)$$

The screening mass here, $\mu(T)$, is an increasing function of temperature and contains the information on the description of the QGP in question.

Two forms of $\mu(T)$ are considered by Gunion and Vogt [171]. The first is obtained from a parametrisation of N-flavour QCD on the lattice [177], which is designated as the $SU(N)$ QGP and is written in terms of the critical temperature T_c as

$$\frac{\mu(T)}{T_c} \sim 4 \frac{T}{T_c}. \quad (7.10)$$

The second form is an estimation from 3 flavours perturbative QCD, applied to the group $SU(3)$, which is consequently designated the *3 flavours QGP* or *3FP*. The screening mass-dependence in this description is written as

$$\frac{\mu(T)}{T_c} = \sqrt{1 + \frac{n_f}{6} \cdot g \left(\frac{T}{T_c} \right) \cdot \frac{T}{T_c}}, \quad (7.11)$$

where the temperature-dependent running coupling constant is

$$g^2 = g \left(\frac{\mu(T)}{T_c} \right) = \frac{24\pi^2}{(33 - 2n_f) \ln \left[\left(\frac{19T_c}{\Lambda_{MS}} \right) \left(\frac{T}{T_c} \right) \right]}. \quad (7.12)$$

These two forms of $\mu(T)$ will be used to predict the modification of the transverse momentum spectra of quarkonia.

	J/ψ	ψ'	χ_c
τ_F (fm)	0.89	2.0	1.5
T_D (3FP) (MeV)	459	211	204
T_D (SUN) (MeV)	190	128	26

Table 7.8: Formation time and dissociation temperatures for the $c\bar{c}$ system in the 3-flavours plasma [171] and SU(N) [111] plasma case.

are chosen following [172] for this study, to be

$$\begin{aligned} t_0 &= 0.1\text{fm} & T_0 &= 1.14\text{GeV} & (\text{SU}(N)) \\ t_0 &= 0.5\text{fm} & T_0 &= 0.82\text{GeV} & (\text{SU}(3)). \end{aligned} \quad (7.14)$$

The entropy density is proportional to T^3 , thus time at which the temperature drops to below T_D is

$$t_D = t_0 \left(\frac{T_0}{T_D} \right)^3. \quad (7.15)$$

This means that as long as $t_D/t_f > 1$, quarkonium will be suppressed; it follows that the maximum p_T at which suppression occurs is then given by

$$p_T^{max} = M \sqrt{\left(\frac{t_D}{\tau_F} \right)^2 - 1}. \quad (7.16)$$

This upper bound on the transverse momentum is set purely by the finite lifetime of the plasma, however, as stated above, the finite spatial extent of the plasma also imposes a limit on the suppression of the quarkonia.

If the $q\bar{q}$ is created at transverse position $x^\mu = (0, \vec{r}, 0)$ and with transverse momentum $p^\mu = (\sqrt{M^2 + p_T^2}, \vec{p}_T, 0)$ then the bound state is formed after formation time τ_F at transverse position $x^\mu = \left(\sqrt{1 + \left(\frac{p_T}{M} \right)^2}, \vec{p}_T, 0 \right)$. The quarkonia will therefore not be suppressed if $|\vec{r} + \frac{\tau_F}{M} \vec{p}_T| > r_S$, where r_S is the evolving screening radius of the plasma, with $R = r_S(0)$ with fixed initial size of the plasma at $t = 0$. In this model, there is a clear competition between the spatial and temporal extent of the QGP and these two competing effects contribute accordingly to the final suppression pattern.

The survival probability of $b\bar{b}$ pairs is calculated as a function of p_T of quarkonia in the two types of plasma (3FP and SU(N)) are shown in figures 7.13 and 7.14 The p_T -dependent expression of the survival probability of the quarkonia is defined the ratio of the number of those that escape the plasma intact, to the total number that would have been formed in the absence of plasma at that p_T . This is obtained by integrating over the spatial co-ordinate at fixed p_T :

$$S(p_T) = \frac{\int_0^R dr \cdot r \rho(r) \theta(r, p_T)}{\pi \int_0^R (dr \cdot r \rho(r))}, \quad (7.17)$$




Figure 7.13: Survival probability of $b\bar{b}$ quarkonia family for a SU(3) plasma radius of 1 fm (7.13(a)) and $R = R_{p_b}$ fm (7.13(b)). The p_T dependence of the suppression of Y (solid blue), Y' (dashed red) and χ_b (dot-dashed magenta) is shown. Figure adapted from [171]

where the density $\rho(r)$ is parametrised as

$$\rho(r) = \left(1 - \left(\frac{r}{R}\right)^2\right)^\alpha \quad (7.18)$$

and $\theta(r, p_T)$ is given by [171]

$$\theta(r, p_T) = \begin{cases} \pi & z \leq -1 \\ \cos^{-1} z & |z| < 1 \\ 0 & z \geq 1 \end{cases}, \quad (7.19)$$

Here, we make use of the substitution

$$z = \frac{r_S^2 - r^2 - \left(\frac{r_{FP} p_T}{M}\right)^2}{2r_F r \frac{p_T}{M}}, \quad (7.20)$$

which is convenient in describing the fate of the quarkonium pair: when $z \geq 1$, the pair never escapes, while when $z \leq -1$, the pair always escapes. It is then clear that for small values of p_T , the suppression is large, and for values of $p_T > p_T^{\max}$ the suppression is zero.

Regarding the shape of the suppression distributions, it is worthwhile to note right away certain distinct characteristics of the four cases under investigation (the two models and the two plasma sizes each considered). In the SU(3) case, for both the large and the smaller plasma sizes ($R = R_{p_b}$ and $R = 1$ fm respectively), it is seen that there is no suppression of the Y , while both the Y' and χ_b are suppressed out to fairly large transverse momentum values. On the other hand, in the SU(N) case, not only is there significant suppression for all species even for $R = 1$ fm, but the Y is suppressed more significantly than the Y' and χ_b . The situation is different again when the larger plasma case of $R = 7.1$ fm is considered. In this case, all species are heavily suppressed, but there is a slightly more complex evolution. While at low p_T




Figure 7.14: Survival probability of $b\bar{b}$ quarkonia family for a 3 flavours plasma radius of different plasma radii. The p_T dependence of the suppression of Υ (solid blue), Υ' (dashed red) and χ_b (dot-dashed magenta) is shown. Note the extended scale on the right. (Figure adapted from [171]).

(< 25 GeV/c), the resonances are progressively suppressed, with Υ suffering more than Υ' , which in turn is suppressed more than the χ_b . However, after about 25 GeV/c, the Υ is less suppressed than the Υ' .

These distinct suppression patterns predicted by the two models and the two sizes considered should assist in the unravelling of the true nature of the produced state. However, the suppression can not be measured directly and has to be folded into the production spectra discussed previously in order to produce the expected suppressed spectra.

7.6 Suppressed quarkonia spectra

An expression of the suppression of quarkonia states as a function of p_T has been derived, allowing to convolute it with the distribution in the case of no plasma (hadron gas scenario) in order to obtain a parametrisation for the suppressed p_T distributions :

$$\left. \frac{d\sigma}{dp_T} \right|_{QGP} = S(p_T) \otimes \frac{d\sigma}{dp_T}. \quad (7.21)$$

These convolutions are shown in figure 7.15 and 7.16 for the $SU(N)$ and $SU(3)$ cases respectively. Also shown is the comparison of a large plasma radius and small plasma radius for a given model. We note that for the suppression of directly-produced quarkonia, in the $SU(N)$ case there is an increased suppression of the lower-lying states compared to those of higher-states. The Υ' is more suppressed with respect to the χ_b and similarly, the Υ is more suppressed with respect to the Υ' . This situation is already visible in the case of the small plasma radius of $R = 1$ fm, but is very pronounced in the extreme case of the larger plasma radius of $R = 7.1$ fm. This scenario is reversed in the case of the $SU(3)$ model of the plasma, where the higher-lying Υ' is more suppressed with respect to the lower-lying Υ . It is noted that in the

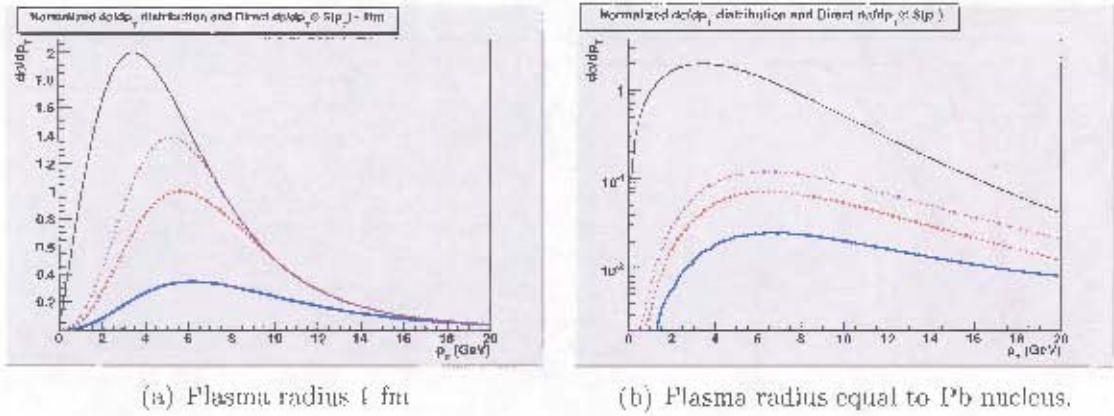


Figure 7.15: Convolution of the SU(N)-predicted QGP suppression scenario compared to the standard production distribution (black) of Υ (solid blue), Υ' (dashed-red) and χ_b (dot-dashed magenta), for two extreme cases of the plasma radius – 1 fm and 7.1 fm.

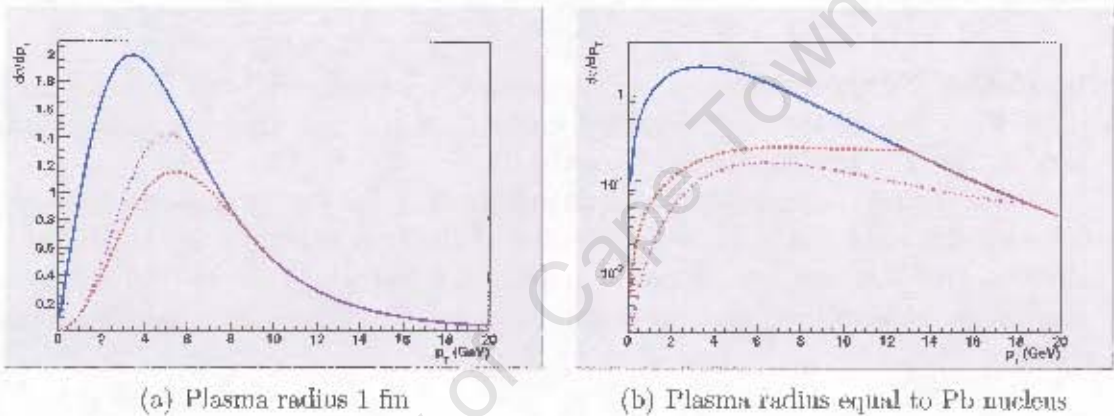


Figure 7.16: Convolution of the SU(3)-predicted QGP suppression scenario, of Υ (solid blue), Υ' (dashed red) and χ_b (dot-dashed magenta), for two extreme cases of the plasma radius – 1 fm and 7.1 fm (log scale shown for the $R = 7.1$ fm case in order to make the differences clearer).

case of the small plasma radius, the χ_b suppression is less than that of the Υ' , while when the radius is increased to $R = R_{\text{Pb}}$, this is no longer the case and the χ_b is suppressed more heavily.

It remains, however, to include the feed-down contributions as was done previously for the hadronic case in section 7.2, to determine the actual observable in the experiment, since it will not be possible to make exclusive measurements. For this reason, equation (7.21) should be modified by adding the relevant feed-down sources for for each resonance.

7.6.1 Yields including feed-down decays

We shall consider only the two lowest-lying states - Υ and Υ' , as the Υ'' and χ_b states will not be produced abundantly enough to permit a statistically relevant study of it,

nor does the spectrometer have sufficient resolution in their mass régime to permit a clean identification. These excited states may, however, contribute to the lower Υ and Υ' states.

For the Υ' , the sources are Υ'' and χ_b family, while for the Υ , the sources include the Υ' too. The indirect production spectrum is built up by using the respective direct production of the resonance states and the branching ratios of the higher states to the lower ones, convoluted with their suppression factor. We write first the integrated indirect production cross-section for the Υ and Υ' :

$$\begin{aligned} \sigma_{\Upsilon}^{Ind} = & \sigma_{\Upsilon}^d \cdot S_{\Upsilon}(p_T) + \\ & Br(\chi_b(1P, 2P) \rightarrow \Upsilon) \\ & \sigma_{\chi_b(1P, 2P)}^d \cdot S_{\chi_b(1P, 2P)}(p_T) \\ & + Br(\Upsilon' \rightarrow \Upsilon) \cdot \sigma^d \Upsilon' \cdot S_{\Upsilon'}(p_T) \end{aligned} \quad (7.22)$$

and

$$\begin{aligned} \sigma_{\Upsilon'}^{Ind} = & \sigma_{\Upsilon'}^d \cdot S_{\Upsilon'}(p_T) + \\ & Br(\chi_b(2P) \rightarrow \Upsilon') \cdot \sigma_{\chi_b(2P)}^d \cdot S_{\chi_b(2P)}(p_T), \end{aligned} \quad (7.23)$$

where we have introduced the branching ratios of the quarkonia states. The values of these branching ratios are taken from the ALICE Physics Performance Report (PPR), Volume 1 [141].

The convolutions of the p_T distribution of quarkonia production with the indirect suppression for the Υ and Υ' resonances for the different QGP scenarios are shown in figures 7.17 and 7.18.




Figure 7.17: Convolution of indirect production spectra of Υ (solid blue) and Υ' (dashed red) and the SU(N)-predicted QGP suppression distribution, with the standard production distribution (black), for two extreme cases – plasma radius of 1 fm and of 7.1 fm.

With the feed-down decays incorporated into the convolution of the spectra predicted by the model, we see a clear difference in the shape of the quarkonia spectra for the two models. For the SU(N) case, the Υ and Υ' spectra are quite similar, both for the small and large plasma sizes. For the SU(3) plasma scenario, however, this is

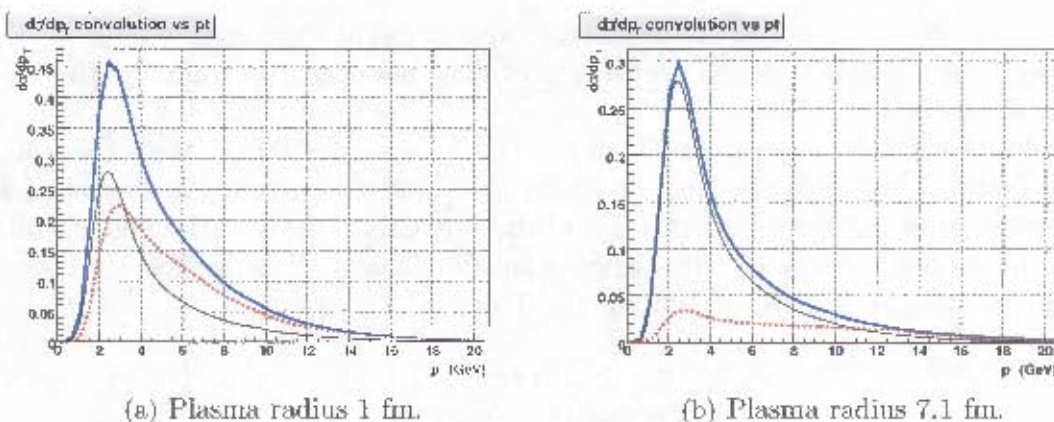


Figure 7.18: Convolution indirect production spectra of Υ (solid blue) and Υ' (dashed red) and the SU(3)-predicted QGP suppression distribution, with the standard production distribution (black).

no longer the case and there is a marked difference in the Υ and Υ' transverse momentum spectra. The Υ' is significantly more suppressed compared to the Υ , which is produced in fact in quite large numbers, similar to the case of no suppression.

7.6.2 Modifications to the ALICE simulation framework

The ALICE offline framework, AliRoot, contains a set of Monte-Carlo event generators which are used to simulate the expected physics during the experiment. This is contained in the `EVGEN` module of AliRoot. Among these generators is a library of p_T and y distributions of various sources of muons, `AliGenMUONlib`, which contains the parametrisations of heavy quarkonia spectra. In order to perform Monte-Carlo simulations with this parametrisation in the ALICE simulation framework, a few modifications had to be made to this library, however, which included the addition of the suppression parametrisations. Equations (7.17 - 7.21) were coded into `AliGenMUONlib`, taking into account the size of the system and the type of QGP under study - either the SU(3) form or the SU(N) form of $\mu(T)$ (equation (7.10) or (7.11)). With this addition, the parametrised p_T distribution of QGP-suppressed quarkonia in this model were available to the generator `AliGenParam`. Once this was done, the full offline simulation and reconstruction framework, including the simulation of the dHLT could be brought to bear to simulate a data set for these QGP models.

7.7 Analysis of Monte-Carlo suppression data

The resulting data generated using the parametrised spectra were processed with the offline dHLT simulation code, in order to obtain the expected response using this trigger. The resulting p_T spectra in the two QGP cases are shown in figure 7.19.




Figure 7.19: Transverse momentum distributions for the SUN and SU3 QGP cases, for an event sample of pure signal - $1.85 \cdot 10^5$ and $4.09 \cdot 10^5$ events respectively.

7.7.1 Results of dHLT processing

These two event samples were passed through the trigger filters, as with the hadronic samples. The relative number of dHLT and L0 triggers validated (high cut only) are given in table 7.9. The efficiency of the L0 and dHLT cuts are shown in figure 7.20(a)

Scenario	% L0 trigger passed	% dHLT trigger passed
No QGP	95.43	96.90
SU(N)	96.57	98.06
SU(3)	98.12	99.17

Table 7.9: The percentage of L0 high p_T and asymmetric (L0 low p_T and dHLT high p_T) triggers passed on average for the scenarios of no QGP production and the SU(N) and SU(3) cases.

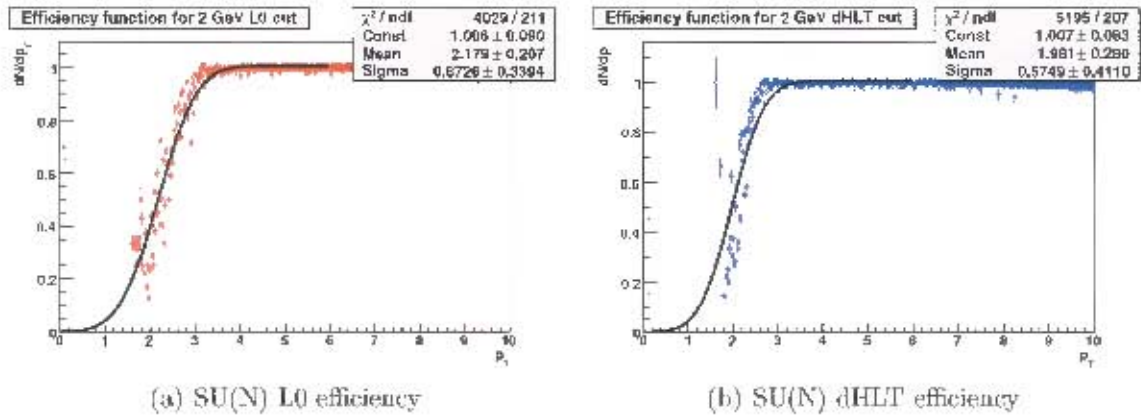
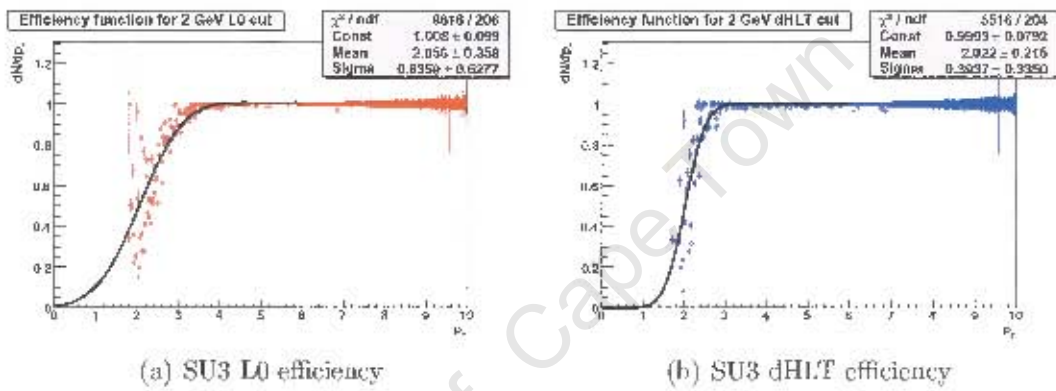
and 7.20(b) respectively.

The shape of the dHLT efficiency curves suggests that the cut could be made slightly higher, without sacrificing significant signal acceptance, in the scenario of QGP-suppressed Υ production. This is corroborated by the exceedingly high acceptance (near 100 %) of the Υ decay muons, shown in table 7.9. This is due to the fact that the p_T spectra in figure 7.19 show that the peak lies somewhat higher in p_T than in the un-suppressed case. This in turn can be explained by the higher suppression of low- p_T quarkonia, which has the effect of shifting the peak out to higher transverse momenta.

7.8 Discussion

7.8.1 Interpretation of dHLT results

We have here demonstrated in several test cases, the dHLT performs well on simulated data. The reduction in the background level is consistent with simulations

Figure 7.20: L0 and dHLT efficiency spectra for the SU(N) Υ data sample.

on flat spectra data. The increased accuracy of the track-finding algorithm shows some enhancement of the dHLT efficiency over that of the L0, while it is shown also that an insignificant level of signal is lost. However, it should be pointed out the above conclusions of background rejection and signal retention rely on the assumption that the signal is characterised exclusively by a track with p_T above the cut. A better variable to make a cut on is that of the di-muon invariant mass, which is calculated on muon pair candidate. This strategy for signal identification has been identified as an avenue of extension of the dHLT algorithm.

It should also be noted that the current estimates have been obtained in an averaged way. This means that the spectra are built up by considering the dHLT applied to several tracks from several events, but with no timing constraints, which represents a kind of upper limit on the capability of the track-finding algorithm. The online implementation of the tracking algorithm will have timing constraints, which will be exacerbated in more central events due to the larger number of tracks per event. This effect is currently under study in two future works [165, 160].

7.8.2 Description and choice of QGP model

It should also be mentioned that the QGP model used here to generate modified p_T distributions of quarkonia is accompanied by certain conditions and assumptions. The first of these is the choice of the initial conditions. We have chosen the initial conditions which correspond to the most suppression of the Υ , as these result in the greatest modification of the resulting muon p_T spectrum. It should be mentioned that certain initial conditions, such as those obtained by using the uncertainty principle as a guiding principle leads to no suppression of the Υ , for example. See [171] for a full discussion on the effects of various choices of initial conditions on the final suppressed quarkonia spectra.

The model chosen is not unique in predicting suppression, and there are also cases in which an enhancement in the yield of quarkonia, due to statistical recombination, for example see reference [179]. The specific model chosen here predicted a significant modification in the transverse momentum distribution as well as an overall suppression. Since this modification would impact the dHLT efficiency, the model was chosen as an example. However, it would be interesting to see how the dHLT efficiency would behave in different conditions, for example, with enhanced quarkonia yield around the cut value.

Chapter 8

Summary and Outlook

8.1 Summary

As the RHIC and SPS data is re-analysed and extended, a deeper understanding of the produced state of matter is attained. It is becoming possible, furthermore for theoretical lattice QCD calculations to make accurate predictions at non-zero baryon density; the picture emerging is that of a dense, tightly coupled and well-thermalised system which lives for a few fermi/c and expands with a velocity of a large fraction of the speed of light. Heavy quarks are seen - indirectly, via their decay into di-electrons - to lose most of their energy, as heavy-flavour mesons are suppressed out to large p_T . The disappearance of back-to-back jets in central Au-Au collisions at high centre of mass (cms) energy and the re-appearance of these in d-Au collisions indicates that the suppression is indeed a final-state effect on the produced particles, rather than an initial-state modification of the production cross-sections. These heavy quarks show a large degree of thermalisation too, as they display a significant elliptic flow component.

As progress is made in understanding the physics in this régime, a new scale is opened up with the LHC experiments - of these, ALICE will be the only dedicated heavy-ion experiment. As discussed in chapter 2, one of the components of ALICE, the dimuon spectrometer, is dedicated to detecting muons from muonic decay of heavy-flavour mesons, with good accuracy and statistics. The principle raison d'être of this detector is the detection of dimuons from heavy quarkonia decay (J/ψ and Υ), since the anomalous suppression of these is considered as a signal for the presence of a QGP.

The high resolution which can in principle be attained with the dimuon spectrometer is degraded however, by the presence of a muonic background¹, which requires the presence of a filtering trigger layer in the DAQ, in order to reduce the background levels. This layer, the dHLT, has been discussed in the fourth and fifth chapters. The design of the system, discussed in chapter 5, builds on the data transport software already used in the overall ALICE HLT, which uses a publisher-subscriber model for

¹This is over and above the large soft background from beam-fragment interactions in the beam-shield, consisting mostly of electrons and photons.

data flow from the front-end processors to the eventual decision-making components.

For the dimuon High Level Trigger (dHLT), an initial study performed by collaborators in Clermont-Ferrand was extended upon and the two-step track-finding algorithm was integrated into the ALICE offline software. Using this algorithm and Monte-Carlo simulations of the spectrometer, the performance of this algorithm was tested using first single muons with flat p_T distributions. This confirms the capability of the dHLT to correctly identify the presence of high- p_T muons and reject background consisting of low- p_T muons, by using a p_T cut. This cut is shown to be sharper than the L0 cut, due to the increased resolution of the p_T calculation by the dimuon High Level Trigger (dHLT) tracking algorithm. With the dHLT in place, it should therefore be possible to improve significance of the data set.

Apart from validation of the dHLT using offline simulations, the code was also integrated into the online software framework. The dHLT components, using the data transport framework mentioned above, were tested with dummy data, in order to check the data flow. By increasing the rate of the data input, we were able to check whether the data transport framework could handle the foreseen data readout rate at the experiment - 1 kHz. As a further exercise, the framework was tested on a global scale with (dummy) data sources in Europe and a decision component in Cape Town in a topology similar to that used in the ALICE TPC. This test - more of a proof-of-principle - also succeeded, demonstrating the inherent robustness of the data transport framework.

Having been shown to work on benchmark flat spectra of muons, this thesis attempted to apply the dHLT to more realistic data sets of pure quarkonia, pure background spectra and mixtures of the two. Spectra and yields expected by binary-scaling of pp data to Pb-Pb collision was generated, as well as a data set which included a specific parametrisation of expected spectra from a QGP model. These results confirm that the system fulfils the requirements made of it to suppress background efficiently while still identifying signal tracks.

8.2 Outlook

8.2.1 The future of the dHLT

The dHLT has been simulated in a small-scale study on Monte-Carlo data and has been shown to perform its assigned task of reducing background and correctly identifying candidate muons from quarkonia decay. While this thesis is limited to the initial work and background of the dHLT, it is worth mentioning the next steps in the further development of the dHLT, which rely on these initial results.

Stress testing of the dHLT

One of the most important tests of the system before being put into full use at the experiment is to have the software and hardware fully integrated. This means that the components of the entire processing chain, from raw data input, to dHLT decision, have to be integrated and properly tested. This testing includes stress testing on the

given hardware, in order to determine the limits, bottlenecks and points of failure in the system, as well as quality-assurance testing to be assured of the validity of the system's performance. Depending on the outcome of the hit-reconstruction algorithm development on the FPGAs, the benchmarking system will include these as a data input source. If this is the case, then a detailed and rigorous comparison of the software hit reconstruction running on the node CPUs and the FPGA. The results of this stress-testing will be in an upcoming thesis [165].

Evolution of the tracking algorithm

The now-standard two-step tracking algorithm tested so far with the dimuon High Level Trigger (dHLT) is not the only possible one. Indeed, although it is possibly the simplest tracking algorithm needed in order to identify high- p_T tracks, it does have drawbacks. These include the fact that it uses the L0 output (which does not have 100 % efficiency) as tracking seeds, as well as the inherent resolution of an algorithm which only uses 2 points to calculate the p_T of a track. In principle, if the positions from the other tracking stations could be used, a better estimation of the p_T could be made and as a result, the cut could be made sharper. The development of alternative tracking algorithms is contained in an upcoming thesis [160].

8.2.2 Evolution of quarkonia models and implications for the dHLT

Recent quarkonia measurements at RHIC have driven theoretical developments attempting to better describe the evolution and effect of the QGP on quarkonia, in order to be able to make accurate predictions. Although only a static suppression model is presented here, dynamic suppression models, including regeneration factors are being developed. In the near future, the application of these models should lead to a deeper understanding of the relevant importance of the various suppression and regeneration mechanisms present in the produced state, at least at RHIC energies.

8.2.3 Conclusion

In conclusion, this thesis has demonstrated the first step in the development of the dHLT for the ALICE experiment. The hit reconstruction and track reconstruction algorithms have been implemented as processing components in the ALICE HLT data transport framework and have been benchmarked. The efficiency tracker is roughly constant near 98 % as a function of the number of tracks in the spectrometer. For the worst-case scenario of several hundred tracks in the spectrometer, we expect only a 10-15 % fake track level which can be accommodated by the processing infrastructure. Hits can be reconstructed online to an accuracy of around $100\mu\text{m}$ and the p_T calculation of reconstructed tracks is less than 200 MeV/c at both 1 GeV/c and 2 GeV/c cuts. The estimated time required to perform hit reconstruction ($\sim 0.5\text{ms}$) and track reconstruction (0.3 ms) is within the limit set by the 1 kHz L0 trigger rate

expected. The background rejection factor is estimated at around 21 %, integrated over p_T .

While the basic requirements of background reduction and tracking efficiency are met, there are several areas of improvement. Besides extending the functionality and performance of the system, it is being stress-tested in order to be ready for full integration into the experiment when the time comes. More precise estimates of the exact reduction in the rate of specific dHLT trigger levels (high- p_T , low- p_T) await a full integration of the system into the HLT framework, which is currently under way.

University of Cape Town

Appendix A

The UCT-CERN Computing Facilities.

A.1 Introduction

When the UCT-CERN Research Centre was accepted into the ALICE experiment to work on the dHLT, it was clear that some kind of computing facility would be essential. After some internal discussion and consultation with other groups at UCT working in fields of science where computationally intensive simulations were also used, the COTS cluster model was chosen as the design model. This meant that initially, cheap, commercially available desktop PCs could be bought and with the appropriate software, they could be quickly and cheaply integrated into a working cluster. Since many software tools existed at the time to do this - both open-source and proprietary - it was not clear at first in which direction to proceed. After some investigation into the “best practices” of Linux clustering, the Open Source Clustering Applications Resources (OSCAR) toolkit was chosen. Reasons for this choice are detailed below, but the main motivating factors were cost and freedom to expand or modify the system later.

This appendix describes the motivation, design, implementation and experience of the UCT-CERN computing facility. It represents over 3 years of work and evolution and has formed an important backdrop to the work presented in this thesis.

A.2 Motivation and goals

The main purpose of the UCT-CERN Research Centre was to develop and test the dimuon High Level Trigger (dHLT) for the ALICE experiment at CERN. However, this was not the only activity which necessitated the development of a computing

facility. The physics studies being undertaken at the Centre required the simulation and reconstruction of heavy ion collisions, which – in the case of ALICE physics – had to be passed through the detector with a transport model¹. The resulting interactions of the produced particles with the detectors would form part of the detailed analysis. These particle transport simulations are highly computationally intensive and can take up to 10 hours on a Pentium III 1 GHz desktop. For physics studies not to be limited by statistical noise, several hundred thousand of these events may be needed, depending on the particular study. Also, in order to realistically simulate what may be measured by the ALICE detectors during experimental conditions, the number of events generated should correspond to that which is expected from physics and this can reach very high numbers too, from 10^2 to 10^8 events, depending on the signal in question. Apart from the time issue, the size of the output of these simulation runs, which can reach several Gigabytes, is an important factor to consider.

It was thus clear that in order to do any physics studies at all, a processing farm would be needed, since these requirements are clearly far beyond the capability of the average desktop computer. What is more, the fact that there were several students in the group meant that some sort of multi-user system would be needed, with load-balancing and proper dynamic resource allocation capabilities.

Apart from the near-term requirement to provide a simulation, development and analysis platform for the members of the centre, a strategic goal of the centre was to join the ALICE data grid as a Tier 2 centre². This would require providing computing resources (CPU, disk storage, etc) to be attached to the grid and made available to the other users in the ALICE experiment. Again, fulfilling this goal clearly meant that a computing facility was needed and the first Linux cluster built served as a useful example of a small-scale prototype of a Tier 2 centre.

An important concern was the time it that would be required to build the system and the human resources needed to administer the system. UCT-CERN Research Centre consisted of physicists and post-graduate physics students, with no dedicated IT staff. Two options presented themselves : outsource the installation, configuration and administration to capable professionals, or do everything in house, developing expertise along the way. The second option was chosen, partly due to the lack of funds needed for outsourcing, but also since it also allowed to build internal competence in the field of high performance computing, and furthermore led to fruitful short-term collaborations with other groups working in different fields of science. The most important factor leading to the success of this choice was the availability of sophisticated administration, resource management and monitoring software bundled with OSCAR. This allowed to build in a high level of automation, which took a significant portion of the load off of the human system administration, leaving more time being dedicated to actual physics applications.

¹Mainly GEANT3 was used, but the Virtual Monte Carlo of ROOT allowed the user to choose from a set of supported Monte-Carlo codes.

²This has now become the LHC Computing Grid, or LHC Computing Grid (LCG). For the definition of the ALICE distributed computing Tiers, see the ALICE Computing TDR [178]

A.3 Design considerations

Apart from financial constraints, the design of the facility was guided by the short- and long-term needs of the centre, described above. Although these constraints remained rather general, a few specific technical design considerations were identified. The system required :

- network-mountable shared disk partitions for permanent data storage
- fairly large storage capacity, of the order of a Terabyte.
- writable disk space on the processing nodes, on the same PCI bus as the processor.
- reasonably fast network connection between the nodes, as well as exterior network connectivity to the rest of the internet.

The third consideration mentioned here stems from the nature of our simulations, which have disk-I/O rates of the order of a few MB/s. In order not to saturate network bandwidth and throttling the performance, it was required to be able to write simulation output directly to the local disk. The results would be staged and transferred to a Network-Attached Storage (NAS) later, in a more controlled manner. These requirements may be termed “hard” requirements, since if they were not satisfied, it is likely that the computing facility would be so severely limited in performance of our application as to be useless.

In addition to these “hard” requirements, a set of “soft” requirements were identified. In contrast to the former, the soft requirements were not essential to the functioning of the system, but would make it more efficient and easier to use. Examples of soft requirements would be a set of software for scheduling, distribution and execution of jobs, monitoring and reporting performance and security of the systems, as well as automation of administrative tasks. Just such a set of software existed in the OSCAR toolkit. This toolkit contains ready-made solutions to the problem of compute cluster creation and maintenance, as well as tools for semi-automated server configuration, node imaging [180], networking, administration and monitoring [181]. In addition to the set of administrative software, interfaces to common high-performance computing (HPC) tools for batch-queue (TORQUE and MAUI) and distributed-memory (Message-Passing Interface (MPI)) [182] software are bundled with OSCAR. In short, the OSCAR toolkit was found to be comprehensive, extensible, flexible and the development team showed very high levels of support and commitment. These reasons led us to choose the OSCAR toolkit to build the system.

A.4 Implementation

Once the general design of the system was finalised, the hardware was acquired and the implementation - hardware and software installation, configuration and testing




Figure A.1: A view of a CARMEN compute node, showing the PC components. Each node was equipped with 80 GB hard drive and 512 MB of RAM

- was begun. The hardware consisted of 20 compute nodes, a single head-node, a 10/100 Mb/s Fast Ethernet switch and a rack with the required cabling for network and power.

A.4.1 Physical setup

A 4-level rack was capable of holding 6 PC cases per level, as well as providing space for at most 2 1U ethernet switches was built by the workshop in the Physics Department at UCT. The compute nodes generic desktop PCs (see figure A.1) were stacked onto the rack, as shown in A.3. The original switch³ had 24 10/100 Mb/s ports and 2 Gb/s daisy-chain ports. The compute nodes were each connected to a Fast Ethernet port, while the head node was connected to the GbE port, since most traffic would be between the compute nodes and the head node, and not intra-node.

The `/home` directories of the head node were mounted via NFS on the compute nodes, as well as a 500 GB partition of the external hardware RAID array. This array consisted of a rack of 8 integrated digital electronics (IDE) disks, connected via IDE-SCSI interfaces to a small computer systems interface (SCSI) bus and directly

³This was later replaced by a GbE switch for higher performance and reliability. See section A.5

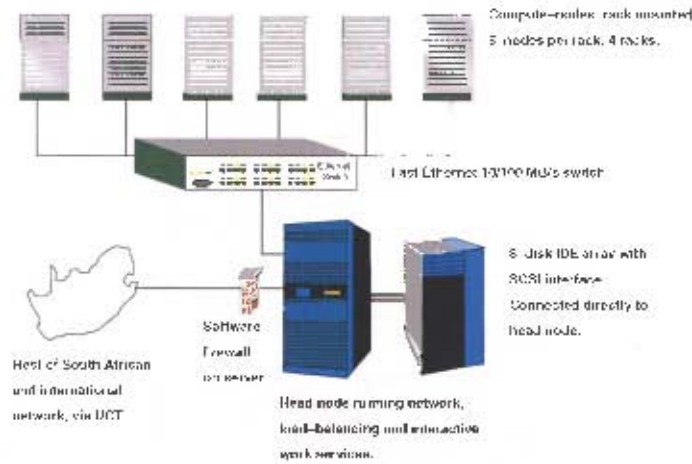


Figure A.2: Schematic diagram of the first implementation of the UCT-CERN computing facility CARMEN, showing network connections amongst nodes and data storage.

attached to the head node. This array of disks served as redundant mass-storage to which data sets generated on the compute nodes were mirrored. Thus, in the case of a local disk failure on one of the nodes, the data would be safely available on the RAID array. This problem of disk failure was to be a regular, but not entirely surprising, occurrence. With the disks in near-continuous use, the probability for a disk failure was highly increased. However, it was noticed that the the disks of certain machines failed more often than others. These would fail fairly regularly, with a frequency of a few months, and have to be checked or re-imaged (or at worst, replaced), while others sustained vigorous and continuous use, both read and write. The root cause of the disk failures was never reliably isolated, although faulty power supplies were the primary suspect.

A.4.2 Data storage capacity

Data storage capacity was a crucial requirement of the cluster: not only did it need to be sufficiently large for the group members to perform large simulation runs and store the output data, but it should be organised in a way that made it simple to access the data once it was generated. Thus, a set of machines with a disk each would not suffice, even if that disk were fairly large. What is more, the storage should be reliable and redundant, so that in the rare case of failure, the disks could be rebuilt. For this reason, it was chosen to make a distinction between “scratch” space and semi-permanent storage.

The scratch space took the form of a `/data` partition on each node which was a very large fraction of the available space, since only about 1.5 GB was required for the operating system. The `/data` partitions on each node were thus about 60 GB each, giving 1.2 TB of scratch space.

The disk array initially contained 8 disks of 100 GB capacity in a RAID-5 array, but this configuration allowed neither hot-swapping nor fail-over. Since both these




Figure A.3: Photo of the current CARMEN cluster at UCT-CERN, showing 3 levels of the 4-level rack, compute nodes, RAID array and head node

functionalities were deemed very useful and were supported by the crate, it was decided to move to a 6-disk RAID-5 array with one disk as a hot spare, and the other as a fail-over disk (see figure A.5). The choice of RAID level meant that the available space for data storage was reduced from 660 GB to about 550 GB, but the reliability that this configuration offered far outweighed the reduction in storage capacity.

The sum of these two figures gives an idea of the total storage available at any given time, roughly 1.7 TB - 2 TB.

A.4.3 Operating system and resource management

Once the hardware installation was completed, the software installation was done. This was essentially split into 3 steps :

Server installation : The head node installed with an operating system.

OSCAR installation : an OSCAR distribution installed on the head node, with the services and software necessary.

Compute node installation : the images created on the head node pushed out to the compute nodes and post installation scripts run.

The operating system (OS) chosen for the first implementation was RedHat Linux 9, for the reasons that it was in common use at UCT-CERN as well as at CERN itself and many other labs, as well as being supported by OSCAR⁴. This choice would change with the evolution of the computing facilities, for reasons given below.

⁴The OSCAR distribution used in the initial implementation of CARMEN was a modified version 1 beta release, v1b4.

Resource Management

Due to the nature of the work done at UCT-CERN, which involved mostly long batch jobs of Monte-Carlo simulations, the software of interest to us was the Portable Batch System (PBS). After the default installation of OpenPBS, which proved to be unstable, it was decided to move to PBSPRO, from Altair Corporation, mainly because this provided the same functionality with higher stability, and because they offered an educational grant which waived license fees. This was in turn retired in favour of the more open and flexible Terascale Open-Source Resource and QUEue Manager (TORQUE) from Cluster Resources.

All processing jobs were managed by the scheduler, which had a policy for determining when and on which node to run them. Since the number of users was quite small and there was no great disparity between the nature of the jobs (in terms of execution time or computing resources required), the policy was thus kept fairly simple to a “first-come-first-serve” basis. The scheduler did however provide the possibility of more sophisticated resource allocation, including backfilling the queue if smaller jobs were available.

A.4.4 Monitoring

Monitoring the performance and state of the compute nodes is of primary importance in ensuring that the system is maximally efficient. To this end, a very widely-used monitoring and reporting package – ganglia – is included in the OSCAR distribution. This provides a system data polling and reporting daemon (`gmond`) on the client side and a collection daemon on the server side (`gmetad`). The data is collected and stored in eXtensible Markup Language (XML) and displayed via a web interface using the round-robin database (RRD) tool. This information was very useful in knowing at a glance what the status of the nodes was, whether any of them had suffered a failure or of the jobs were being properly distributed properly. It did not, however provide access to detailed job information or the process tree.

This additional functionality was provided by CLUSTER MONITOR (CLUMON), an OSCAR package developed at the National Centre for Supercomputing Applications (NCSA), with an interface to the operating system and other software processes provided by the Performance Co-Pilot (PCP) package from Silicon Graphics Inc. (SGI). With the PCP reporting, CLUMON displayed, again via a web page, the scheduler job list, status and the process tree of every reporting node. These two monitoring systems provided access to the nodes’ ‘vital statistics’, without having to actually log into the node⁵

⁵This information is, however “read-only” and no executions or changes can be made.

A.5 Evolution

Once the initial implementation had been finished and the system proved stable over a long time, the next stage in the evolution was discussed. Several areas of improvement were identified by the users and administrators and some of these were slowly implemented during 2004 and 2005. A discussion of the problem areas and the solutions are given below.

A.5.1 Node software and hardware upgrade

After the first year of running, a funding application for a hardware upgrade was made, which was approved. The upgrade consisted of new CPUs and motherboards for the compute nodes and a more robust, faster head node. The compute nodes were upgraded from 1.7 GHz Pentium IVs to 2.6 GHz Hyper-Threading Technology (HTT) Pentium IV's, from Intel. The existing dual-Pentium III 1 GHz head node was replaced for a 3.0 GHz HTT Pentium IV with double the original RAM - 1024 MB.

Since the machines had to be effectively de-commissioned and then re-commissioned, it was decided to profit from the disruption by performing a software upgrade at the same time, since support for RedHat Linux 9 had expired. This meant moving from Red Hat Linux 9, (which was the last free and open distribution from Red Hat corporation) to another Linux distribution. The obvious choice for cost reasons was the community-driven Fedora project, over the for-fee RedHat Enterprise Linux (RHEL), since it was a natural extension of the operating system which had been in use at the Centre. Other choices, however, such as RHEL "clones"⁶ such as Scientific Linux and Community Enterprise Operating System (CEntOS) were becoming viable too (see section A.6). The final choice of OS was Fedora Core 3, due in part to the fact that it was one of the first flavours of Linux to be shipped with a native 2.6 kernel.

As well as upgrading the OS, an upgrade of the OSCAR software was also effected, moving from the version 1 beta to a stable version 4 release. Several updates and enhancements to the software packages contained in OSCAR had been made, including a change in batch-queue system⁷, better scheduling software and much improved user interface and hardware support.

A.5.2 Network upgrade

One of the major sources of inefficiency in the cluster arose from the original switch. This was a Surecom 10/100 Mb/s Fast Ethernet switch. Since the `/home` partition of the head node and a partition of the RAID array were mounted over NFS on the nodes via this switch, the network bandwidth was only barely sufficient. This became very much apparent when several jobs wrote simultaneously to NFS file partitions and

⁶These distributions were re-built from the source RedHat Enterprise Linux (RHEL) RPMs, and released under the GPL, with all reference to the company Red Hat removed.

⁷The more flexible and stable TORQUE resource manager with the MAUI scheduler was used.

when a high rate of I/O errors was experienced. It was decided to eliminate this inefficiency by upgrading the switch to a Gigabit Ethernet switch.

The new HP Procurve Gigabit Ethernet switch was installed in 2004 and immediately the network-related issues were resolved. The NFS-mounted file systems were much more stable and the time taken for node imaging was reduced drastically, from about an hour to just under 10 minutes.

A.5.3 Separation of roles

The upgrade cycle resulted in several unused motherboards, and it was decided to put them to use. From the outset, one of the design flaws often pointed to was the fact that the head node, running all the cluster services, as well as acting as an interactive compute node, was under severe load. Apart from presenting an obvious single point of failure, this also generally made the cluster less efficient. It was decided to lighten the load on the head node by moving some of the services it ran onto dedicated machines. Obvious candidate services for this separation were the file serving and firewall roles, which were transferred to dedicated machines (see figure A.4). For the other services, like the `http` server, batch-queue server and monitoring services it was decided to keep them running locally on the head node, since the amount of re-configuration required to move them was substantial and outweighed any potential performance benefits.

Dedicated firewall

One of the most important, but most overlooked aspects of the design of the computing facility was the network security. Following a malicious exploit in the `ssh` daemon previously, the system was cracked once already during the first year of running, resulting in significant downtime. With the hardware and software upgrade, it was decided to finally address this point. A stand-by box was re-tasked as a dedicated firewall, running a completely different OS than the cluster head node⁸ and acting as a transparent (in the sense that its presence is not noticed by the user or the university network administrators) port-forwarding node. This node, `fire.phy`, was placed on the cannibalised 3Com Fast Ethernet switch, and connected to the local UCT network. Several servers, including the group web server, were also connected to this switch and all traffic to these servers and cluster nodes was henceforth via this node. Special care was then taken to ensure its network and physical security.

Dedicated file server

In keeping with the theme of optimising the performance and security of the computing facility, the next step was to construct a dedicated file server for the RAID array. This would take significant load off of the head node, and keep disk errors to

⁸This node was installed with a Debian Linux distribution.

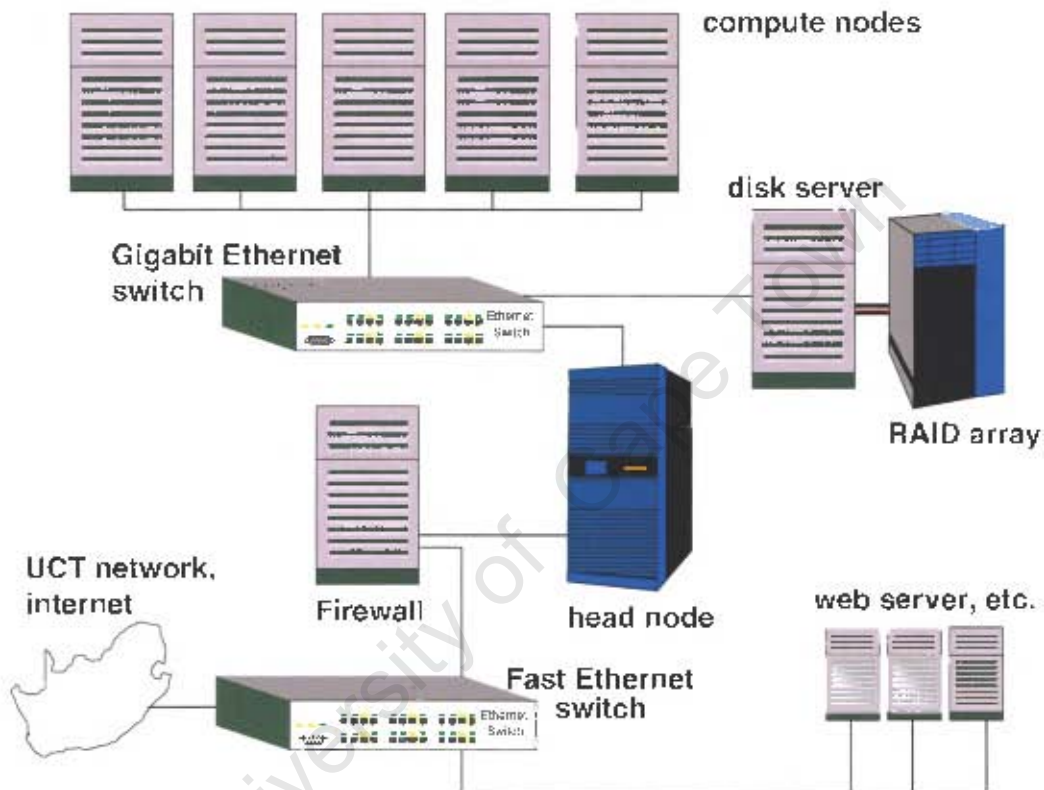


Figure A.4: Schematic diagram of the upgraded network and hardware configuration of CARMEN. Note the separation of file server and firewall roles from the head node to dedicated nodes.



8-disk IDE array
with SCSI interface
to file server host

6 disks in
RAID-5 array

1 hot spare

1 fail-over disk

Figure A.5: Photo showing the dedicated file server and RAID array with bay-doors. Red tags denote spare disk and hot-swap disk, for fail-over capability.

a minimum. The dual-PIII that served as the previous head node of CARMEN (see figure A.5) was chosen for this role and the re-imaged machine with the RAID array directly attached to it was integrated onto the cluster subnet, as shown in figure A.4.

A.6 Sandiego - 64 bit extension to UCT-CERN computing facilities.

The first two implementations of CARMEN were both based on 32 bit CPUs, and proved very useful and reliable for the purposes that had been envisaged for them. The next step in the progress of the computing facility was to add a cluster of 64 bit-based machines. These were acquired in early 2005 from Sun Microsystems, when 16 rack-mounted Sun "Fire" v20z Opteron-based servers were bought.

A.6.1 Challenges set by the 64 bit platform

There were essentially 3 challenges presented by the new 64 bit machines:

- Operating System
- Applications
- User experience

Choice of operating system

The first challenge was to decide on the right OS. This choice merited some discussion, because two avenues presented themselves : continue with a Red Hat-based or derived distribution of Linux, or use the recently open-sourced OS from Sun Microsystems itself, OpenSolaris. Solaris has been well-known for many years to be a very reliable and stable choice on the 64 bit platform and all indications were that OpenSolaris would continue in that tradition. Besides, the V20z was at the time of writing a very popular choice for high-performance computing and there was no doubt that there would be a significant drive to support OpenSolaris for the Opteron on that hardware.

On the other hand, the Scientific Linux distribution had reached maturity by this stage and was well-supported by our collaborating laboratory, European Centre for Nuclear Research (CERN), as well as almost every other large HEP laboratories. What is more, the physics applications which in use were known to be supported and actively developed on this platform. Therefore, for reasons mainly of consistency with other laboratories in the ALICE collaboration, as well as to allow development on the dHLT to remain coherent, it was decided to install Sandiego with Scientific Linux, version 4. This distribution was built for x86_64 architecture and based on the Linux 2.6 series kernel.

The second challenge alluded to above was that of requiring the applications in use to be compatible with the new architecture. The applications used by the UCT-CERN Research Centre are divided into two broad categories, namely offline and online. The source distributions of both of these are available since the development is still underway, however the offline code has less platform-specific restrictions than the online (dHLT) code.

Offline simulations : The code for offline simulations, AliRoot, is maintained and tested for several different platforms and operating systems. As a result, this did not place much constraint on our choice of OS. The official distribution of Linux used at CERN was chosen, in order to reduce the overhead of maintaining the ROOT and AliRoot distributions. The offline code has been tested to compile on the 64 bit platform, and simulation and reconstruction benchmarks have shown that there are no incompatibilities.

dHLT work : The continued dHLT development was the main reason for choosing Linux over any other operating system, since it had been developed on Linux until that point, but more importantly, required several drivers (for shared memory management and for communication with the FPGA over the PCI bus) which are specific to the Linux kernel. The dHLT code was successfully compiled and tested on the SL platform in November 2005, but the 'bigphys' patch for the 2.6 kernel was not available, which means that the dHLT will be developed primarily on the existing 32 bit cluster, CARMEN.

By physically separating the clusters, we forced users to choose whether to use 32-bit or 64-bit applications and libraries. This implicit choice resulted in greater ease of use for the users, who were not forced to explicitly request the one or the other platform when submitting their jobs

A.6.2 Sandiego implementation

Having taken care of software considerations and made the choice of OS, the implementation of the cluster could be discussed. The major issue was that of integration with existing systems, due to their disparate nature, both in terms of computing platform as well as physically. The hardware for Sandiego is housed in a completely separate rack (see figure A.6 and is on a separate subnet. Apart from the physical separation, Sandiego is a native 64 bit system, and other software issues aside, even the compilers were completely different to CARMEN were different as were most of the system libraries.

Proposed integration schemes

On the one hand, from the administration point of view, it was desirable to have Sandiego tightly-integrated with CARMEN in order to present a single resource to the users and administrators. In this scenario, a single head-node on CARMEN would manage the monitoring and reporting of both systems and the file server on `file.phy` would be visible globally as well. The physically separate subnets could be daisy-chained and placed behind the existing firewall. A single job queue would feed jobs to the schedulers running on the head-nodes of the 32 bit and the 64 bit sub-clusters, which would allow users to continue to work transparently.

On the other hand, the systems were disparate enough for the reasons given above for the integration to pose several practical complications. The most important technical issue encountered was that of disk imaging to nodes of different architectures, from a single head node. This means that two different repositories would be required to create images from. Not only would these repositories be for different Linux distributions⁹ but more importantly for different processor architectures. Although 32- and 64-bit libraries and binaries can co-inhabit an x86_64 platform, this is not possible on an x86 platform. Since this is the architecture of the current image server (which is also the CARMEN head node), the image server would have to be moved to an x86_64 node. This process was deemed too complicated and provided no significant benefit, to be implemented.

Besides image serving, file serving between 32 bit and 64 bit operating systems was not guaranteed to be simple, neither were the networking changes required entirely obvious. Also, the users would have to explicitly request a given platform and environment variables be set accordingly, depending on which platform they wanted their jobs submitted to. This was only an extra complication in already complicated scenario.

Eventually, it was decided to simplify matters and keep the clusters physically and logically separate, with their own distinct subnets, head nodes and job queues. Further integration would be left to the next iteration evolution.

⁹The problem of hosting different Linux distributions on a single head node has since been addressed with the `yum` tool of Erich Focht, now included in OSCAR. See <http://home.arcor.de/efocht/oscar/>

Hardware

Contrary to CARMEN, which was built from the ground-up from cheap PCs and a home-made rack, Sandiego was delivered by Sun-certified hardware vendors. This included the rack, which was provided with its own power distributor. The nodes themselves were 1U racks and more information can be found on the web. For illustration, the system is shown in figure A.6

Remote server management

A problem identified with CARMEN which was never overcome was the issue of how to deal with nodes in the event of a unexpected power failure, or when there was no network access to them due to failures in the boot sequence, etc. At each such event, it was necessary for the administrator to physically enter the room where the clusters were housed, attach a monitor and debug the problem locally. It was desirable to circumvent the necessity of physical presence in the room, for security as well as practical reasons¹⁰, as much as possible.

This issue was solved by having a baseboard management controller (BMC) on the servers, interfaced with an ethernet cable, which allowed to communicate remotely with the computers even if the platform power was disconnected. The BMCs and the various sensors (fan power, temperature, disk revolutions, cover state, etc) allowed to monitor in great detail the health of the server, and communication via the intelligent platform management interface (IPMI) [184] protocol made this information available on the head node for monitoring and alerting purposes. IPMI is a message-based protocol for communicating with the BMC - the data path is shown in figure A.6.2. Apart from passive monitoring (and active reporting), IPMI allows the administrator to perform platform operations on the server as well, such as rebooting the node, checking the OS status, monitoring the console messages powering on from a powered-off state, etc.

The BMCs were connected in so-called "out-of-band" mode (by not using the platform network interface card (NIC)s) by an ethernet adapter to the local subnet. To avoid an excess of cabling, these NICs were daisy-chained to each other via cross-over ethernet cables. This setup afforded full control over all aspects of the cluster nodes from the head node, using secure IPMI commands¹¹.

A.7 Outlook

The UCT-CERN Research Centre computing facilities at this stage comprise a significant and very flexible resource for the members of the group. The requirements laid out above in section A.1 have been fulfilled on the local scale. Work on the

¹⁰For example, the administrator may be physically unable to be present, attending a conference, etc.

¹¹These were issued via SSH.



Figure A.6: Photo of the Sandiego cluster rack from Sun Microsystems

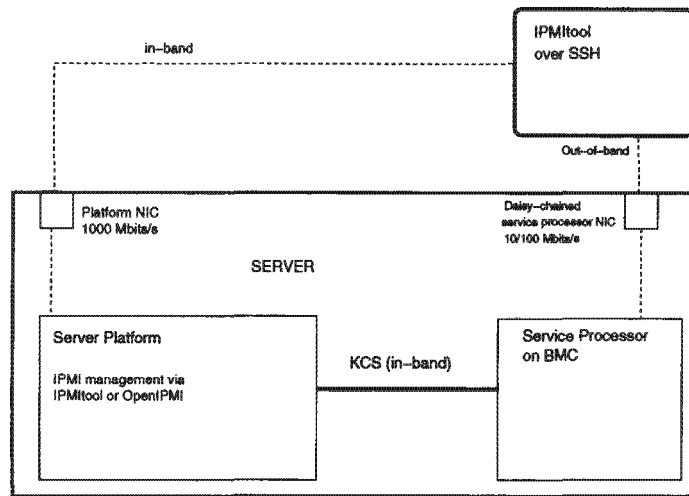


Figure A.7: Diagram showing how the IPMItool application manages the Sandiego node platforms through out-of-band communications with the BMC and in-band keyboard controller style (KCS) communications with the server platform.

dHLT has undergone several integration tests on CARMEN during 2004 and 2005 and it is envisaged to perform a fairly long stress test during 2006. The system has been shown to be flexible and simple enough for external users to use, due to a large part to the fact that industry-standard open-source software was used to build it. The short-term collaborations that have been undertaken with scientists from fields such as molecular dynamics [185], radiotherapy and materials science and have contributed enormously to the understanding of the limits of the system. It is hoped that, given the financial and human resources, the system can continue to evolve so keep up with the demands placed on it by the UCT-CERN Research Centre, its commitment to the ALICE experiment in terms of the dHLT and the LCG, as well as provide opportunities for local external scientists.

Glossary

Institute of Electrical and Electronics Engineers (IEEE)1394

An Institute of Electrical and Electronics Engineers (IEEE)high-speed serial bus standard, commonly known as FireWire. 4.3

1U

1U refers to the standard height of commodity computing components in racks A.4

ALICE Data Challenge

The ALICE data challenges happen about once per year and are done in order to test the readiness of ALICE for the data it plans to take. Part of the Data Challenges is generating and processing a fraction of the data expected, in order to benchmark the computing facilities and included in this processing is the HLT processing. 6.7

AliRoot

Offline code base of the ALICE experiment. Based on ROOT, implemented mostly in C++ with interfaces to geometrical modelling, simulation, reconstruction, monitoring and analysis modules 4.2

backfilling

Backfilling is a process by which smaller jobs which require less resources than large, long jobs may be scheduled to run before the longer jobs, even if the longer jobs were submitted before the shorter ones. The backfilling policy defines the requirements for this process. A.4

Bevelac

Accelerator used at Berkeley during the 1970's capable of accelerating particles to 1 GeV 1.2

Fast Ethernet

Fast Ethernet is a collective term for a number of Ethernet standards that carry traffic at the nominal rate of 100 Mbit/s, against the original Ethernet speed of 10 Mbit/s. (Wikipedia) 6.5

fault-tolerance

A system is fault-tolerant if it can suffer a fault and continue to perform its task. This is usually achieved by re-assigning tasks to other parts of the system
4.1

hot-swapping

A hot swap is the replacement of a device with a similar device while the computer system using it remains in operation. The replacement can be because of a device failure or for storage devices to substitute other data. A.4

Hough transform

The Hough transform is a feature extraction technique used in digital image processing. The rho-theta parametrisation universally used today was first described in Duda, R. O. and P. E. Hart, Use of the Hough Transformation to Detect Lines and Curves in Pictures, Comm. ACM, Vol. 15, pp. 11-15 (January, 1972). 4.2

Hough transform

The Hough transform is a feature extraction technique used in digital image processing. 4.1

L1

Level 1 of the ALICE trigger system. 3.4

L2

Level 2 of the ALICE trigger system 3.4

L2A

Dimuon trigger level 2 "accept" signal 6.2

load-balancing

The process of dynamically assigning more resources to a task when required
4.1

massively-parallel

In the sphere of parallel computing, the system is referred to as "massively parallel" or "embarrassingly parallel" depending on its latency and inter-node requirements. Systems which require a high degree of (synchronous) inter-node communication are called "massively-parallel", while those which, like the HLT, are more hierarchical and granular in nature and do not require much inter-node communication are referred to as "embarrassingly parallel". 4.3

Moore's Law

Moore's law is the empirical observation that at our rate of technological development, the complexity of an integrated circuit, with respect to minimum component cost will double in about 24 months. 4.1

Myrinet

Myrinet, ANSI/VITA 26-1998, is a high-speed local area networking system designed by Myricom. See [183]. 4.3

NA50

North Area 50 experiment at CERN's SPS. The NA 60 experiment evolved from the results of the NA 50 collaboration 2.3

NA57

Heavy-Ion fixed target experiment at CERN's SPS 3.6

non-Abelian

Non-Abelian refers to Lie groups whose elements do not commute. 1.1

pileup

Refers to a situation when reading data from beam collisions, where several collision events have been read-out in a single read process. Thus, multiple events have been "piled up" or superimposed on each other. 4.1

radiation length

A radiation length, χ_0 of a particle is the distance it can travel within a medium before losing $1/e$ of its energy 3.5

ROOT

ROOT is an object-oriented analysis framework developed at European Centre for Nuclear Research (CERN), which is the de-facto data analysis framework in heavy-ion experiments (see <http://root.cern.ch> and [146]) 5.3

TCP/IP

Transmission Control Protocol / Internet Protocol. A network communication protocol 3.6

WA98

A heavy ion experiment at CERN's Western Area experimental facility 2.4

List of Acronyms

ADC: Analog to Digital Converter

51

ADULT: A Dual-Level Threshold

51

AGS: Alternating Gradient Synchrotron

3, 12

ALICE: A Large Ion Collider Experiment

The dedicated heavy ion experiment at the CERN LHC 5

ASIC: Application-specific integrated circuit

48, 51

ATLAS: A Toroidal Large Acceptance Spectrometer

Experiment at CERN 5

BMC: Baseboard management controller

144

BNL: Brookhaven National Laboratory

US laboratory in Long Island New York Home to AGS and RHIC 3

BRAHMS: Broad RAnge Hadron Magnetic Spectrometers

An experiment at RHIC 4

CARMEN: Cluster for African Research in Massive Energetic Nuclei

Cluster of Linux PC's inaugurated October 2003 which was used in the dHLT integration test. See Appendix A 92

CDF: Collider Detector facility at Fermilab

One of the facilities for high-energy experiments at Fermi National Laboratory.

108

CEM: Colour Evaporation Model

28

CEntOS: Community Enterprise Operating System

138

CERN: European Centre for Nuclear Research

3, 12, 142, 149

CLUMON: CLUster MONitor

An OSCAR package based on PCP and MySQL to provide system metrics and cluster job state monitoring and reporting 137

CMS: Compact Muon Solenoid

Experiment at CERN's LHC 5

cms: Centre of mass

A convenient reference frame to to kinematic calculations in 3, 9, 128

COM: Colour Octet Model

28

COTS: Commodity off-the-shelf

78

CPU: Central Processing Unit

92

CROCUS: Cluster Read Out Concentrator Unit

48

CSM: Colour Singlet Model

28

CTP: Central Trigger Processor

52, 53

DAQ: Data Acquisition

x, 54

DARC: Dimuon Arm Readout Card

51

DATE: Data Acquisition and Test Environment

ALICE software for testing data acquisition and online monitoring programs 55

DCS: Detector Control System

60

DDL: Detector Data Link

ALICE standard digital data transport protocol 48

dHLT: Dimuon High Level Trigger

High-level trigger dedicated to the muon spectrometer of the ALICE experiment
v, 57, 60, 70, 104, 128–131

DIU: Destination-Interface Unit

54

DMA: Direct Memory Access

55

DPM: Dual-Parton Model

30

DRORC: Destination Read-out Receiver Card

55

DSP: Digital Signal Processor

48

FEE: Front-End Electronics

Electronic cards usually containing pre-amplifier amplifier and zero-suppression
chips amongst others 48

FEP: Front-End Processor

Raw data processing nodes of the ALICE HLT farm. 60, 61

FNAL: Fermilab National Accelerator Laboratory

Synchrotron accelerator facility near Chicago IL (USA) 30

GbE: Gigabit Ethernet

Gigabit Ethernet is a term describing various technologies for implementing
Ethernet networking at a nominal speed of one gigabit per second. 64

GDC: Global Data Concentrator

55

GMS: Geometrical Monitoring System

System of optical monitors for the ALICE dimuon spectrometer which has the responsibility of monitoring changes in position of the chambers 82

HBT: Hanbury-Brown Twiss

An analysis method for determining the size of objects emitting bosons. 43

HIJING: Heavy-Ion Jet Interaction Generator

A software package which simulates the production and interaction of jets produced in heavy-ion collisions 45

HLT: High Level Trigger

51

HMPID: High Momentum Particle IDentification

Detector subsystem in ALICE which uses RICH detectors in order to identify particles with very high momentum 38, 41

HPC: High-performance computing

A generic term given to intensive computing 133

HTT: Hyper-Threading Technology

Intel's trademark for their implementation of the simultaneous multi-threading technology on the Pentium 4 micro-architecture (see www.intel.com/technology/hyperthreading) 138

IDE: Integrated digital electronics

134

IEEE: Institute of Electrical and Electronics Engineers

147

I/O: Input / Output

79, 92

IP: Interaction Point

The point in ALICE where the two counter-rotating beams intersect 81

IPMI: Intelligent platform management interface

144

ITS: Inner Tracking Chamber

The innermost layer of detectors in ALICE based on silicon technology 38, 60

KCS: Keyboard controller style

One of the System Interfaces used to connect a BMC to a system CPU. It is based on the protocol of the legacy Intel 805x keyboard controller chips (hence the name) xiv, 146

LBNL: Lawrence Berkeley National Laboratory

3

LCG: LHC Computing Grid

132

LDC: Local Data Concentrator

55

LEP: Large Electron-Positron

Accelerator complex built for accelerating and colliding electrons and positrons which was upgraded in order to accelerate hadrons too 35

LHC: Large Hadronic Collider

The next-generation accelerator at CERN which was built as an upgrade to the pre-existing LEP accelerator (See also LEP) iv, 4, 35

LHCb: LHC b factory

Experiment at CERN LHC 5

LTU: Local Trigger Unit

52

MANAS: Multiplexed ANAlog Signal processor

The front-end signal processor chip used in the dimuon spectrometer tracking chambers which performs amongst other functions the pre-amplification of the detected signal. 51

MANU: MANAS Numérique

The sequencing and digitising chip which processes the output of the MANAS 48

MARC: MUON Arm Readout Card

48, 51

MBZ: Must Be Zero

74

MPI: Message-Passing Interface

A standard interface for message-passing between coupled processes usually in a shared memory environment but also used with distributed memory over a LAN in a cluster. 133

NAS: Network-Attached Storage

Mass storage attached to a local network as opposed to local storage connected to the compute node directly. 133

NCSA: National Centre for Supercomputing Applications

A Supercomputing centre at the University of Illinois at Urbana-Champaign 137

NFS: Network File System

92

NIC: Network interface card

144

NLO: Next-to-leading order

A perturbative calculation which has terms of the order of α^3 ix, 24, 28

nPDF: Nuclear parton distribution function

Functions describing the distribution of the momentum of partons in nucleons within a nucleus as opposed to single nucleons. This distribution includes nuclear effects such as shadowing. 31

OS: Operating system

136

OSCAR: Open Source Clustering Applications Resources

set of "best-practices" software and cluster integration procedures. 131

PATCH: Protocol for ALICE Tracking Chamber

48

PBS: Portable Batch System

A workload management package originally developed at NASA with several derivatives such as TORQUE PBSPro and OpenPBS. 137

PC: Personal Computer

63

PCP: Performance Co-Pilot

A framework developed by SGI corporation to support system-level performance monitoring and performance management 137

PHENIX: Pioneering High Energy Nuclear Interaction eXperiment

An experiment at RHIC 4

PHOS: PHOton Spectrometer

38, 41

PID: Particle identification

38

PMD: Photon Multiplicity Detector

x, 43

PPR: Physics Performance Report

124

pQCD: Perturbative QCD

ix, 24

QCD: Quantum Chromodynamics

The current standard theory describing the fundamental interactions of quarks and gluons 2

QGP: Quark-Gluon Plasma

A state of nuclear matter where quarks and gluons are no longer bound within hadrons but free to explore a wider volume 3

RAID: Redundant Array of Inexpensive Disks

92

RAM: Random-Access Memory

92

RHEL: RedHat Enterprise Linux

138

RHIC: Relativistic Heavy Ion Collider

12

RICH: Ring-Imaging Čerenkov detector

41

RoI: Region of interest

xv, 56, 61, 81

RORC: Read-out Receiver Card

The input terminus of the ALICE HLT. 61

RPC: Resistive Plate Chamber

47

RRD: Round-robin database

An “industry standard” software package for displaying collected system data visually in graphs usually via a web page. 137

SCSI: Small computer systems interface

interface consisting of a standard port between a computer and its peripherals
134

SDD: Silicon Drift Detector

38

SGI: Silicon Graphics Inc.

137

SIU: Source-Interface Unit

Electronic component which acts as the source interface to the DDL cable. 51,
54

SMP: Symmetric multiprocessor

65

SPD: Silicon Pixel Detector

38

SPS: Super Proton Synchrotron

3, 12

SSD: Silicon Strip Detector

38

STAR: Solenoidal Tracker At RHIC

An experiment at RHIC ix, 4, 26

TDR: Technical Design Report

61

TOF: Time of Flight

38, 40

TORQUE: Terascale Open-Source Resource and QUEue Manager

A resource manager developed as an extension over OpenPBS which offers several improvements including fault tolerance scalability and feature extensions
137

TPC: Time Projection Chamber

The main tracking detector in ALICE 39

TRD: Transition Radiation Detector

38, 40

USB: Universal Serial Bus

64

XML: EXtensible Markup Language

A general-purpose markup language for describing different kinds of data 137

ZDC: Zero degree calorimeter

x, 42

Bibliography

- [1] H. Kastrup, P. Zerwas, eds., QCD 20 Years Later (World Scientific, Singapore, 1993)
- [2] D. J. Gross, F. Wilczek, Phys. Rev. Lett. 30 (1973) 1343
- [3] J. C. Collins, M. J. Perry, Phys. Rev. Lett. 34 (1975) 1353
- [4] E. V. Shuryak, Phys. Reports 61 (1980) 71
- [5] F. Karsch, Nucl. Phys. A 698 (2002) 199
- [6] H. Appelshaeuser *et al.*, Eur. Phys. J. C 2 (1998) 661
- [7] J. Cleymans and K. Redlich, Phys. Rev. C 60 (1999) 054908
- [8] F. Becattini and Ulrich W. Heinz, Z. Phys. C 76 (1997) 269-286
- [9] P. Braun-Munzinger, J. Stachel, J. P. Wessels and N. Xu, Phys. Lett. B 365 (1996) 1-6
- [10] Granddon. D Yen, Mark I. Gorenstein, Hörst Stöcker, Shin Nan Yangk and Walter Greiner, J. Phys. G 24 (1998) 1777 - 1784
- [11] H. Satz, Nucl. Phys. A715, 3 (2003).
- [12] CERN press release 2000, <http://pressold.web.cern.ch/PressOld/Releases00/PR01.00EQuarkGluonMatter.html>
- [13] K. Rajagopal and F. Wilczek, Nucl. Phys. B 399 (1993) 395-495
- [14] B. Mohanty and J. Serreau, Phys. Rept. 414 (2005) 263-358
- [15] E. Gladysz Dziadus, INP Report 1879, Cracow (2001)
- [16] T. C. Brooks *et al* (The MiniMax Collaboration) Phys. Rev. D 55 (1997) 5667
- [17] M. M. Aggarwal *et al.* (WA98 Collaboration) Phys. Lett. B 420 (1998) 169-179.

- [18] Miklos Gyulassy, New York Times June 2003
- [19] K. Adcox *et al* (PHENIX collaboration) Nucl. Phys. A **757** (2005) 184-283
- [20] A. Chodos, R. L. Jaffe, K. Johnson, C. B. Thorn and V. F. Weisskopf Phys. Rev. D **9** 3471 (1974)
- [21] S. Soff, S. A. Bass, D. Hardtke and S. Y. Panitkin, Arxiv : nucl-th/0202019
- [22] (NA 50 Collaboration) Physics Letters B 477 (2000) 28; CERN-EP-2000-013
- [23] R. Vogt, Phys. Rep. 310, 197 (1999); H. Satz, Rep. Prog. Phys. **63**, 1511 (2000)
- [24] K. G. Wilson, Phys. Rev. D, 2445 (1974)
- [25] M. Creutz, Phys. Rev. D 21, 2308 (1980)
- [26] L. D. McLerran and B. Svetitsky, Phys. Lett. B 98, 199 (1981)
- [27] J. Engels, F. Karsch, I. Montvay and H. Satz, Phys. Lett. B 101, 89 (1981)
- [28] For a review, see Karsch's lecture notes
- [29] F. Karsch hep-lat/0106019
- [30] R. J. Glauber, Phys. Rev. **100** 242 (1955)
- [31] J. D. Bjorken, Phys. Rev. D **27** (1983), 140-151
- [32] B. Back *et al* (PHOBOS Collaboration), Nucl. Phys. A **757**, (2005) 28
- [33] E. Anderson *et al* (WA57/NA57 Collaboration), CERN-EP/98-64 (1998)
- [34] E. Anderson *et al* (WA57/NA57 Collaboration), Phys. Lett B **449** (1999) 401
- [35] P. Braun-Munzinger, J. W. A. Cleymans, H. Oeschler and K. Redlich, Nucl. Phys. A **697** (2002), 902
- [36] J. Cleymans *et al.*, Euro. Phys. J. A **29** (2006), 119-121
- [37] M. Gaździcki (NA49 Collaboration), J. Phys. G **30** (2004), S701
- [38] M. Gaździcki and M. I. Gorenstein, Acta Phys. Polonica B **30** (1999) 2705
- [39] V. Friese *et al.* (NA49 Collaboration), J. Phys. G **30** (2004), S119-S128
- [40] F. Karsch, K. Redlich, and A. Tawfik, Phys. Lett. B **571** (2003) (67)
- [41] R. Stock, J. Phys. G **30** (2004) S633

- [42] S. Wheaton and J. Cleymans, "THERMUS - A Thermal Model Package for ROOT", ArXiv:hep-ph/0407174
- [43] S. Bass *et al.* (UrQMD Collaboration) Prog. Part. Nucl. Phys **41** (1998) 255-369
- [44] C. Alt *et al.* (NA49 Collaboration), Phys. Rev. C **68** (2003) 034903
- [45] C. Adler *et al.* (STAR Collaboration), Phys. Rev. C **66** (2002) 034904
- [46] A. Tang *et al.* (STAR Collaboration), AIP Conf. Proc. **698** (2004) 701
- [47] C. Adler *et al.* (STAR Collaboration), Nucl. Phys. A **757** (2005) 102
- [48] P. F. Kolb and U. Heinz, in Quark Gluon Plasma 3, eds. R.C. Hwa and X. N. Wang (World Scientific, Singapore, 2003); nucl-th/0305084
- [49] C. Adler *et al.* (STAR Collaboration), Phys. Rev. Lett. **87** (2001) 182301.
- [50] S. S. Adler *et al.*, ArXiv:nucl-ex/0503003
- [51] R. J. Fries, B. Müller, D. K. Srivastava, Phys. Rev. Lett. **90** (2003) 132301
- [52] S. S. Adler *et al.* (PHENIX Collaboration), J. Phys. G **30** (2004) S1003-S1006
- [53] J. W. Cronin *et al.* Phys. Rev. D **11** (1975) 3105
- [54] A. L. S. Angelis *et al.*, Phys. Lett. B **185** , 213 (1987)
- [55] E. Wang and X. N. Wang, Phys. Rev. C **64**, 034901 (2001)
- [56] D. Kharzeev, Y. V. Kovchegov, and K. Tuchin Phys.Rev. D**68** (2003) 094013
- [57] B. Back *et al.* (PHOBOS Collaboration) Phys. Rev. Lett. **91**, 072302 (2003)
- [58] S. S. Adler *et al.* (PHENIX Collaboration) Phys. Rev. Lett. **91**, 072303 (2003)
- [59] J. Adams *et al.* (STAR Collaboration) Phys. Rev. Lett. **91**, 072304 (2003)
- [60] I. Arsene *et al.* (BRAHMS Collaboration) Phys. Rev. Lett. **91**, 072305 (2003)
- [61] S. S. Adler *et al.* (PHENIX Collaboration), ArXiv:nucl-ex/0409015
- [62] S.S. Adler *et al.* (PHENIX Collaboration), ArXiv:nucl-ex/0409015
- [63] Jiangyong Jia (for the PHENIX Collaboration), arXiv:nucl-ex/05/10019 (2006)
- [64] J. Bielcikova, S. Esumi, K. Filimonov, S. Voloshin and J. P. Wurm, Phys. Rev. C **69** (2001) 021901
- [65] I. Vitev, arXiv:hep-ph/0506181

- [66] N. Armesto, C. A. Salgado and U. A. Wiedemann, arXiv:hep-ph/0411341
- [67] X. Zhu, H. Petersen and M. Bleicher J. Phys. **G 32** (2006) S365-S371
- [68] J. Adams *et al.* (STAR Collaboration) Phys. Rev. Lett. **92** (2004) 052302
- [69] P. Huovinen, P. F. Kolb, U. Heinz, P. V. Russkanen and S. A. Voloshin, Phys. Lett. **B 503** (2001) 58
- [70] J. Casalderrey-Solana, E. V. Shuryak and D. Teaney, arXiv:hep-ph/0411315
- [71] K. Eskola, Nucl. Phys. **A 698** (2002) 78-87
- [72] N. Armesto, C. Pajares, Int. J. Mod. Phys. **A 15** (2000) 2019-2052
- [73] A. Hebecker, ArXiv:hep-ph/0111092
- [74] K. J. Eskola, K. Kajantie, P. V. Ruuskanen and K. Tuominen, Nucl. Phys. **B 570** (2000) 379-389
- [75] B. Back *et al* (PHOBOS Collaboration) Phys. Rev. Lett. **86** (2001), 1970
- [76] H. Appelshauer *et. al*, Phys. Rev. Lett. **82** (1999), 2471
- [77] BRAHMS collaboration, Nucl. Phys. **A715** (2003) 171c
- [78] K. Redlich *et al* J.Phys. **G 30** (2004) S1271-S1274
- [79] T. Matsui and H. Satz, Phys. Lett. **B 178** (1986)
- [80] S. Eidelman, *et al.*, Physics Letters **B 592**, 1 (2004)
- [81] F. Karsch and H. Satz Z. Phys. **C 37** (1987) 617,1988
- [82] R. Vogt, Physics Reports 310 (1999)
- [83] M. C. Abreu *et al* (NA38 Collaboration) Phys. Lett. **B 444** (1998) 516
- [84] M. C. Abreu *et al* (NA 38 Collaboration), Phys. Lett. **B 449** (1999) 128
- [85] M. C. Abreu *et al* (NA38 Collaboration) Phys. Lett. **B 466** (1999) 408
- [86] *et al* (NA 50 Collaboration), Nucl. Phys. **A 698** (2002) 543
- [87] P.L. McGaughey *et al.* AIP Conf.Proc. **243** (1992) 65-78
- [88] M.J. Leitch *et al.* " Nuclear Effects on Heavy Quark Production: Results From Fermilab Experiments E-772 and E-789." LA-UR-91-3962 (1991)
- [89] S. J. Brodsky and A. H. Mueller, Phys. Lett. **B 206** (1988) 685
- [90] D. Kharzeev, C. Lourenço, M. Nardi and H. Satz Z. Phys. **C 74** (1997) 307-318

- [91] N. Armesto, A. Capella and E. G. Ferreira, *Phys. Rev. C* **59** (1999) 395
- [92] N. Armesto, A. Capella, E. G. Ferreira, A. Kaidalov and D. Sousa, *Nucl. Phys. A* **698** (2002) 583-586
- [93] D. M. Alde *et al.* *Phys. Rev. Lett* **66** (1991) 2285-2288
- [94] M. Nardi, *Nucl. Phys. A* **774** (2006) 353-366
- [95] K. J. Eskola, V. J. Kolhinen and P. V. Ruuskanen, *Nucl. Phys. B* **535** (1998) 351-371
- [96] K. J. Eskola, V. J. Kolhinen and P. V. Ruuskanen *Nucl. Phys. B* **535** (1998) 351-371
- [97] K. J. Eskola, V. J. Kolhinen and C. A. Salgado, *Eur. Phys. J. C* **9** (1999) 61-68
- [98] J. Cleymans and S. M. Wheaton, *J. Phys. G* **31** (2005) S1069-S1074
- [99] B. Andersson, G. Gustafson, G. Ingelman, T. Sjöstrand, *Phys. Rept.* **97** (1983) 31; B. Anderson in *"The Lund Model"*, Cambridge University Press (1998)
- [100] F. Becattini, J. Cleymans, A. Keranen, K. Redlich and E. Suhonen, *Phys. Rev. C* **64** (2001) 024901
- [101] W. Florkowski, W. Broniowski, M. Michalec, *Acta Phys. Pol. B* **33** (2002) 761
- [102] P. Braun-Munzinger, D. Magestro, K. Redlich, J. Stachel *Phys. Lett. B* **518** (2001) 41
- [103] J. Manninen, F. Becattini, A. Keranen, M. Gazdzicki, R. Stock *nucl-th/0405025*
- [104] E.L. Berger, D.L Jones, *PRD* **23** 1521 (1981)
R. Baier, R. Ruckl, *Z. Phys. C* **19** (1983) 251
- [105] F. Abe *et al* *Phys. Rev. Lett.* **75** (1995)4358
V. Papadimitriou *et al.* (CDF Collaboration), *Int. J. Mod. Phys A* **12** (1997) 3867
S. Abachi *et al* *Phys. Rev. Lett.* **B 370** (1996) 239
- [106] E. Braaten, T. C. Yuan, *Phys. Rev. D* **50** (1994) 3176 *ibid* **D 52** (1995) 6227
G. T. Bodwin *et al* *Phys. Rev. D* **51** (1995) 1125
- [107] J. F. Amundseon *et al* *Phys. Rev. Lett.* **B 390** (1997) 323
O. J. P. Eboli, E. M. Gregores and F. Halzen, *Phys. Rev. D* (1999) 117501

- [108] J. Collins, D.E. Soper and G. Sterman, in: *Perturbative quantum chromodynamics*, ed. A.H. Mueller (World Scientific, Singapore, 1989).
- [109] P. Nason, S. Dawson and R. K. Ellis, *Nucl. Phys. B* **303** (1998) 607
- [110] *Hard Probes in Heavy Ion Collisions at the LHC*, chapter 3, CERN-2004-009 (2004)
- [111] S. Digal, P. Petreczky and H. Satz, *Phys. Rev. D* **64** (2001) 094015
- [112] M. C. Abreu *et al* (NA50 Collaboration), *Phys Lett. B* **477** (2000) 28
- [113] S. S. Adler *et al.* (PHENIX Collaboration), *Phys. Rev. Lett.* **92** (2004) 051802
- [114] S. S. Adler *et al.* (PHENIX Collaboration), *Phys. Rev. C* **69** (2004) 014901
- [115] S. S. Adler *et al.* (PHENIX Collaboration) *Phys. Rev. Lett* **96** (2006) 012304
- [116] S. S. Adler *et al.* (PHENIX Collaboration) arXiv:hep-ex/0611020
- [117] L. Grandchamp, R. Rapp and G. E. Brown, *Phys. Rev. Lett.* **92** 212301 (2005)
- [118] A. Capella and E. G. Ferreira, *Eur. Phys. J. C* **42** (2005) 419
- [119] H. Ohnishi *J. Phys. G* **31** (2004) S1209-S1212
- [120] W.-M. Yao *et al.*, *Journal of Physics G* **33** (2006) 1
- [121] See the L3 experiment web page : <http://l3.web.cern.ch/l3/>
- [122] ALICE ITS Technical Design Report CERN/LHCC 99-12 ALICE TDR 4 (1999)
- [123] ALICE TPC Technical Design Report CERN/LHCC 2000-001 ALICE TDR 7 (2000)
- [124] ALICE TDR Technical Design Report CERN/LHCC 2001-021 ALICE TDR 9 (2001)
- [125] ALICE TOF Technical Design Report CERN/LHCC 2002-016 Addendum to ALICE TDR 8 (2002)
- [126] ALICE HMPID Technical Design Report CERN/LHCC 98-19 ALICE TDR 1 (1998)
- [127] ALICE PHOS Technical Design Report CERN/LHCC 99-4 ALICE TDR 2 (1999)

- [128] ALICE Dimuon Spectrometer Technical Design Report CERN/LHCC 99-22 ALICE TDR 5 (1999)
ALICE Dimuon Spectrometer Technical Design Report Addendum CERN/LHCC 2000-046 Addendum 1 to ALICE TDR 5
- [129] F. Manso, *et al*, ALICE-INT-96-13
- [130] P. Lautridou *et al*, ALICE Internal Note ALICE-INT-1998-27
- [131] R. Divià, P. Jovanovic and P. Vande Vyvre, ALICE-INT-2002-10 revised edition (2004).
- [132] E. Mathieson, Nucl. Instr. and Meth. in Phys. Res., **A270**(1988) 602
- [133] A. Zinchenko, ALICE-INT-2003-006 “A new approach to cluster finding and hit reconstruction in muon chambers of ALICE”
- [134] L. B. Lucy, Astronomical Journal **79** (1974) 745
- [135] E. A. Kolganova, E. L. Kosarev and G. A. Ososkov, Nucl. Instr. and Meth. **A 443** (2000) 464
- [136] “Raw data format for the ALICE dimuon spectrometer”, ALICE dimuon spectrometer collaboration (Christian Finck corresponding author), in preparation.
- [137] B. Espagnon *et al*, ALICE-INT-2004-026 (2004)
- [138] “Etude et réalisation d’un circuit intégré spécifique de mise en forme de signaux délivrés par un détecteur pour une expérience de Physique des particules.”, Ph.D thesis, Laurent Royer (Conservatoire National des Arts et Métiers, 2001).
- [139] C. Finck (corresponding author) ALICE-INT-2005-012 (2005)
- [140] G. Blanchard, P. Crochet and P. Dupieux, ALICE-EN-2003-10 (2003)
- [141] ALICE Collaboration (F. Carminati, Y. Foka, G. Paic, J-P. Revol, K. Safarik, Y. Schutz and U. A. Wiedemann (editors), J. Phys. G: Nucl. Part. Phys. **30** (2004) 1517-1763
- [142] CERN Yellow Report on hard probes : PDFs, shadowing and pA collisions, A. Accardi *et al* (K. J. Eskola editor), ArXiv:hep-ph/0308248
- [143] ALICE Technical Design Report (CERN-LHCC-2003-062), Chapter 2
- [144] P. Vande Vyvre and György Rubin, “ALICE DDL Interface Control Document”, ALICE Internal Note ALICE-INT-1996-43 V9.1
- [145] ALICE DAQ, DATE v4 User’s Guide, ALICE Internal Note ALICE-2002-036

- [146] R. Brun, F. Rademakers, Nucl. Instrum. Meth A **389**, 81-86 (1997)
ROOT home page : <http://root.cern.ch>
- [147] T. Sjostrand *et al*, Journal of High Energy Physics (2006) 026
- [148] GEANT Detector Description and Simulation Tools, CERN Program Library, Long Writeup W5013
- [149] G. E. Moore, Electronics **38**, (1965) 114-116
- [150] "The ALICE Trigger, Data Acquisition, High Level Trigger, Control System Technical Design Report", ALICE-DOC-2004-001 v.2
- [151] "A Modular and Fault-Tolerant Data Transport Framework", T. Steinbeck, Ph.D. thesis, University of Heidelberg (2004)
Arxiv cs.DC/0404014
- [152] S. Bethke *et al* "Report of the Steering Group of the LHCC Computing Review", CERN/RRB-D 2001-3 (2001)
- [153] "Geometry, mapping and segmentation conventions and software for the ALICE dimuon spectrometer", ALICE MUON collaboration, G. Martinez corresponding author (in preparation)
- [154] Sukalyan Chattopaday, private communication
- [155] F. Manso *et al*. ALICE-INT-2002-04. (2002)
- [156] ALICE High Level Trigger Conceptual Design Report
- [157] Private communication, S. Chattopadhyay and S. Chatterjee
- [158] S. Chattopadhyay, presentation at the dimuon week, Berg-En-Dal, South Africa (June 2004)
- [159] Artur Szostak, Private Communication (May 2004)
- [160] A. Szostak, M.Sc. thesis, University of Cape Town (in preparation)
- [161] T. Steinbeck, V. Lindenstruth and M. W. Schultz, IEEE Trans.Nucl.Sci **49**, (2002) 455-459
The ALICE High Level Trigger homepage :
<http://wiki.kip.uni-heidelberg.de/ti/HLT>
- [162] CERN-LHCC-2005-018 : ALICE Computing TDR (ALICE Collaboration)
- [163] S. Chattopaday, International Workshop on Dimuon Physics in Ion-Ion Collisions at the LHC, Berg-En-Dal (3-8 June 2005)

- [164] "A General-Purpose Computing Facility for Scientific Research using OSCAR", B. Becker and M. J. N. Horner, Proceedings of the OSCAR Symposium 2004 <http://www.csm.ornl.gov/oscar04/proceedings.shtml>
- [165] G. Devaux, M.Sc. thesis, University of Cape Town (in preparation)
- [166] B. Nowicki (Network Working Group, Sun Microsystems) "NFS: Network File System Protocol Specification" - RFC 1094 (1989) <http://www.ietf.org/rfc/rfc1094.txt>
- [167] R. Guernane, P. Crochet, A. Morsch and E. Vercellin, ALICE-INT-2005-018 (2005)
- [168] R. Vogt, arXiv hep-ph/0203151
- [169] ALICE Physics Performance Report, Volume II : ALICE Collaboration, B. Alessandro, F. Antinori, J. A. Belikov, C. Blume, A. Dainese, P. Foka, P. Giubellino, B. Hippolyte, C. Kuhn, G. Martinez, M. Monteno, A. Morsch, T. K. Nayak, J. Nystrand, M. López-Noriega, G. Paić, J. Pluta, L. Ramello, J-P. Revol, K. Safarik, J. Schukraft, Y. Schutz, E. Scomparin, R. Snellings, O. Villalobos Baillie and E. Vercellin editors J. Phys. **G 32** (2006) 1295-2040
- [170] P. Crochet, P. Braun-Munzinger, Nucl. Instr. and Meth. **A 484** (2002) p 564 - 572.
- [171] J. F. Gunion and R. Vogt Nucl. Phys. **B 492** 301-337 (1997)
- [172] E. Dumonteil, P. Crochet, ALICE-INT-2005-002
- [173] Eric Dumonteil, Ph.D. Thesis, CEA Saclay 2004
- [174] J. P. Blaizot and J. Y. Ollitrault, Phys. Lett **B 199**, 499-503 (1987)
- [175] D. Kharzeev and H. Satz, Phys. Lett. **B 366**, 316-322 (1996)
- [176] CDF Collaboration, Fermilab Preprint Fermilab-Conf-94/136-E May 1994
- [177] B. Petersson, Nucl. Phys. **A 525** (1991) 237c
- [178] ALICE Computing TDR : <http://aliweb.cern.ch/>
- [179] A. Andronic, P. Braun-Munzinger, K. Redlich, J. Stachel Phys. Lett. **B 571** (2003) 36-44
- [180] System Installation Suite : <http://sisuite.org/>
- [181] Ganglia home page : <http://ganglia.sourceforge.net>
- [182] <http://www.lam-mpi.org/>
- [183] Boden, N.J. *et al*, IEEE Micro **15** (1) (1995) 29 - 36

- [184] The IPMI specification : <http://www.intel.com/design/servers/ipmi>
- [185] P. Mausbach and H.-O. May, Fluid Phase Equilibria **213**, (2003)
- [186] A. D. Martin, R. G. Roberts, W. J. Stirling and R. S. Thorne, European Physics Journal **C4** (1998) 463-496

University of Cape Town

University of Cape Town