Identified Particle Transverse Momentum Distributions from Au+Au Collisions at 62.4 GeV per Nucleon Pair

by

Conor Henderson

MSci, Physics and Applied Mathematics, 1999
Queen's University, Belfast, Northern Ireland.

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2005

Author ....
Department of Physics
7 July, 2005

Certified by ....
Gunther M. Roland
Associate Professor of Physics
Thesis Supervisor

Accepted by ....
Thomas J. Greytak
Associate Department Head for Education

Identified Particle Transverse Momentum Distributions from Au+Au Collisions at 62.4 GeV per Nucleon Pair

by

Conor Henderson

Submitted to the Department of Physics
on 7 July, 2005, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Transverse momentum (p_T) distributions for pions, kaons, protons and antiprotons have been measured near mid-rapidity for Au+Au collisions at $\sqrt{s_{NN}} = 62.4$ GeV using the PHOBOS detector at the Relativistic Heavy-Ion Collider (RHIC) in Brookhaven National Laboratory.

Particle identification is performed using the PHOBOS Time-of-Flight plastic scintillator walls and specific energy loss in the multi-layer silicon Spectrometer, which is also used for track reconstruction and momentum-determination. The spectra are corrected for all detector-dependent effects, including feed-down from weak decays.

At $p_T \sim 3$ GeV/c, protons are measured to be the dominant species of charged hadrons and scale much faster with respect to collision centrality than mesons.

This behaviour at 62.4 GeV is found to be remarkably similar to that observed in Au+Au collisions at 200 GeV, an interesting observation which should serve as an important constraint on the various mechanisms which have been proposed to describe particle production over this p_T range.

Baryon stopping, the transport of baryon number from intial beam rapidity, is explored through the net proton ($p - \bar{p}$) yields at mid-rapidity. These results fill a large gap between the SPS and higher RHIC energies and as such form an important set of data for comparing to models of baryon transport mechanisms.

Thesis Supervisor: Gunther M. Roland
Title: Associate Professor of Physics

Contents

Chapter 1

Introduction

The strong interaction is understood to be a force between elementary particles called quarks, mediated by bosons called gluons and described by the theory of Quantum Chromodynamics (QCD). QCD has the unusual property that the interaction strength increases with distance from the source, a phenomenon called 'asymptotic freedom.' Strongly-interacting matter may exist in a variety of phases, including a state of deconfined quarks and gluons called the Quark-Gluon Plasma (QGP). Heavy-ion collisions at high energies are an attempt to create this QGP and study its properties. Identified particle transverse momentum distributions will be important in elucidating the mechanisms that govern particle production in the complex system formed in relativistic heavy-ion collisions.

1.1 The Quark Model of Hadrons

Cosmic ray studies and the development of particle accelerators in the 1950s led to the discovery of a vast proliferation of baryon and meson resonances. Regularities in the data were seen, and Gell-Mann and Ne'eman independently developed a system of organisation called the 'Eightfold Way' [1] (borrowing the name from the Buddhist path to enlightenment). At that time, the particle at the apex of the baryon decuplet triangle, the Ω^-, was not known experimentally: based on his organisation scheme, Gell-Mann boldly predicted its existence and its properties. When the Ω^- was later seen by experiments at the Brookhaven Alternating Gradient Synchrotron [2], it was a striking confirmation of the theory.

Later, Gell-Mann [3] and Zweig [4] independently proposed that the patterns could be explained if hadrons were actually composed of fractionally-charged, spin-$\frac{1}{2}$ particles called *quarks*.[1] Baryons would be built from three quarks qqq and mesons from a quark-antiquark pair $q\bar{q}$. The three quark types or 'flavours' known at the time were called 'up'(u), 'down'(d) and 'strange'(s); the quark model was later enlarged with the discoveries of the charm quark [5, 6], the bottom quark [7] and eventually the top [8, 9]. The properties of the six Standard Model quarks are displayed in Table 1.1.

Quark	Name	Charge	Mass
u	Up	$+\frac{2}{3}$	1.5 – 4.5 MeV
d	Down	$-\frac{1}{3}$	5 – 8.5 MeV
s	Strange	$-\frac{1}{3}$	80 – 155 MeV
c	Charm	$+\frac{2}{3}$	1.0 – 1.4 GeV
b	Bottom	$-\frac{1}{3}$	4.0 – 4.5 GeV
t	Top	$+\frac{2}{3}$	174.3 ± 5.1 GeV

Table 1.1: Table of quark properties, with masses taken from the Particle Data Book [10].

Further evidence that nucleons had constituents came from deep-inelastic scattering of electrons from protons [11]. While probing the structure of the proton, the experimental results at large momentum transfer were found to exhibit a relationship called 'Bjorken scaling' [12], with the interpretation that the electrons were scattering off point-like, charged, spin-$\frac{1}{2}$ particles inside the proton - quarks.

---

[1] The name 'quark' is taken from a phrase in James Joyce's famously difficult novel *Finnegan's Wake*: "Three quarks for Muster Mark."

1.2 Quantum Chromodynamics

1.2.1 Color

A problem with the quark model was that some states appeared to violate the Exclusion Principle. Consider the Δ^{++} which is interpreted as being uuu. In the absence of any spin-orbit coupling, the wavefunction for the state can be written as a product of separate spatial, spin and flavour functions: $\Psi_{\Delta^{++}} = \psi_L(r)\chi_{spin}\Phi_{flavour}$. The Δ^{++} is known experimentally to have spin $\frac{3}{2}$, so all 3 spin-$\frac{1}{2}$ quarks are in the same spin-state $(\uparrow\uparrow\uparrow)$ and hence χ_{spin} is symmetric under the interchange of any two quarks. If it is uuu, clearly $\Phi_{flavour}$ is also symmetric, and since it is also known to have positive parity, the spatial wavefunction must be symmetric as well. The total wavefunction is therefore symmetric with respect to the interchange of any two quarks. But quarks are fermions, so according to the Exclusion Principle, a collection of them should be *antisymmetric* under exchange.

A solution was to propose the existence of a new, hidden quantum number. The total wavefunction would then actually be: $\Psi_{\Delta^{++}} = \psi_L(r)\chi_{spin}\Phi_{flavour}\Theta_{new}$ and if Θ_{new} were antisymmetric, the situation would be saved. In order to construct an antisymmetric state with three particles, the new quantum number had to be three-valued. Nature suggested a good analogy: white light is composed of equal parts of red, green and blue. Thus the new quantum number was called *color* [13]. It was hypothesised that all natural particles had to be 'colorless' - baryons are RGB, while mesons could be $R\bar{R}$, for example.

Experimental evidence for the existence of color came from measurements of the process $e^+e^- \rightarrow q\bar{q} \rightarrow hadrons$. At collision energies substantially above the mass of $q\bar{q}$ pair and away from meson resonances, the cross-section for this process simplifies considerably. It can be compared to the process $e^+e^- \rightarrow \mu^+\mu^-$, and phase space factors will cancel in the ratio, so that

$$\frac{\sigma(e^+e^- \rightarrow q\bar{q} \rightarrow hadrons)}{\sigma(e^+e^- \rightarrow \mu^+\mu^-)}$$

is essentially proportional to the number of degrees of freedom for hadron production. In this sense, creation of an $R\bar{R}$ pair is distinct from a $G\bar{G}$ pair, thus the existence of color as a true feature of quarks contributes a factor of 3 to the hadron production cross-section. This factor of 3 is in good agreement with the experimental data.

1.2.2 The QCD Lagrangian

The strong interaction between quarks is described by the theory of Quantum Chromodynamics (QCD). Color is the 'charge' of this interaction and it is mediated by eight massless vector bosons called *gluons.* For historical reasons, quarks and gluons are often collectively referred to as *partons.* The QCD Lagrangian is:

$$\mathcal{L}_{QCD} = \sum_{i}^{n_f} \sum_{j} \bar{q}_{ij}(i\gamma^{\mu} D_{\mu} - m_i) q_{ij} - \frac{1}{4} G^{a}_{\mu\nu} G^{\mu\nu}_{a} \tag{1.1}$$

where q_{ij} is the quark field for a quark of flavour $i = u, d, s, c, b, t$, mass m_i and color $j = R, G, B$. γ^{μ} are the 4×4 Dirac matrices. In terms of the eight gluon fields G^{a}_{μ}, the covariant derivative D_{μ} is:

$$D_{\mu} = \partial_{\mu} + igG^{a}_{\mu}\frac{\lambda^{a}}{2} \tag{1.2}$$

g is the coupling constant of the theory. λ^{a} are the eight 3×3 Gell-Mann matrices which satisfy the commutation relations $[\lambda_a, \lambda_b] = i f_{abc} \lambda_c$. The numbers f_{abc} are called the 'structure constants' of the group SU(3) - because they are non-zero, the group is said to be *non-Abelian.*

The QCD Lagrangian is invariant under local SU(3) gauge transformations:[2]

$$q(x) \rightarrow e^{i\alpha_a(x)\lambda_a} q(x) \tag{1.3}$$

To preserve gauge invariance given the above commutation relations for the λ_a operators, the gluon fields must be massless and transform as:

$$G^a_\mu \rightarrow G^a_\mu - \frac{1}{g}\partial_\mu \alpha_a - f_{abc}\alpha_b G^c_\mu \tag{1.4}$$

As a result, the gluon field strength tensor $G^a_{\mu\nu}$ is:

$$G^a_{\mu\nu} = \partial_\mu G^a_\nu - \partial_\nu G^a_\mu - g f_{abc} G^b_\mu G^c_\nu \tag{1.5}$$

The presence of this last term in the theory allows for the possibility of not just quark-gluon couplings, but also direct gluon-gluon coupling, in the form of three-gluon and four-gluon vertices.[3] These direct couplings occur because the gluons themselves carry color. The interactions of QCD are illustrated in Figure 1-1.

1.2.3 Asymptotic Freedom and Quark Confinement

It is well-known from Quantum Electrodynamics that the effective charge of an electron decreases as the test charge is moved away. This is due to vacuum polarisation: the presence of virtual e^+e^- pairs in the vacuum have a screening effect on the bare electron charge, similar to Debye screening in a plasma.

For a bare quark, there is a similar color-screening effect due to virtual $q\bar{q}$ pairs

---

[2] The concept of local gauge invariance was first used in Quantum Electrodynamics (QED), where it was discovered that requiring the Dirac Lagrangian for the electron to be locally gauge invariant introduced a massless vector field corresponding to the photon. Local gauge invariance is now a key principle of modern particle physics.

The idea of a locally gauge invariant theory can be heuristically understood by considering the quark wavefunction as being represented by a vector in color space. Just as in quantum mechanics, only the length of this vector is important, not its orientation (its phase). Thus the theory should have the freedom to arbitrarily rotate the phase of the wavefunction at any point in space: the gluon fields then exist in order to be able to perform this local phase transformation.

[3] It is worth noting that these gluon-gluon couplings would not exist if the group were Abelian, i.e. if the generators of the group commuted. This is the case in Quantum Electrodynamics (which is a U(1) group), and the boson mediator of this interaction, the photon, carries no electric charge.

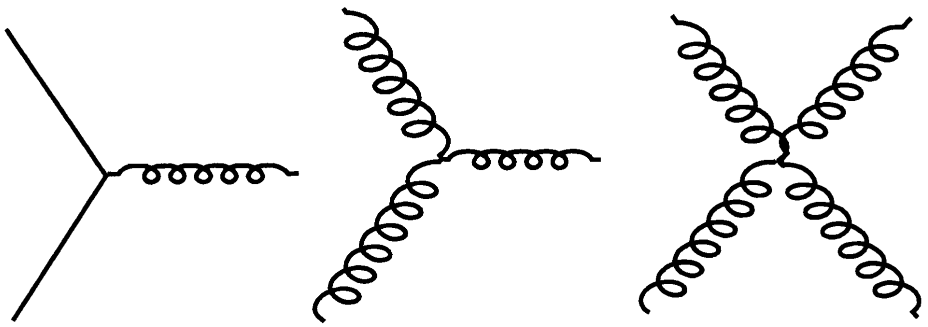

Figure 1-1: Illustration of the possible quark-gluon and gluon-gluon couplings in Quantum Chromodynamics.

in the vacuum. However, in QCD there are also gluons, which are colored, and their presence in the vacuum has an *anti-screening* effect. It is found that the gluon anti-screening effect dominates[4] and the coupling constant of the strong interaction, α_s, actually increases with distance from the source. This is known as *asymptotic freedom* [14, 15].

The variation of this so-called 'running coupling constant' with the momentum transfer Q (inversely proportional to separation distance) is shown in Figure 1-2. In the 'one-loop' approximation, the analytic form is:

$$\alpha_s(Q^2) = \frac{12\pi}{(33 - 2n_f)\log(\frac{Q^2}{\Lambda^2})} \tag{1.6}$$

where $n_f = 6$ is the number of quark flavours and $\Lambda \approx 200$ MeV is a dimensional parameter which sets the scale for QCD.

Asymptotic freedom is thought to be the reason for the phenomenon of quark confinement, the fact that an individual quark has never been seen (which had been a problem for the quark model). Even as the separation between quarks increases, the force between them never goes to zero and it is not possible to completely isolate a

---

[4] Very, *very* roughly speaking, gluon anti-screening dominates over $q\bar{q}$ screening because there are more gluons than quarks.

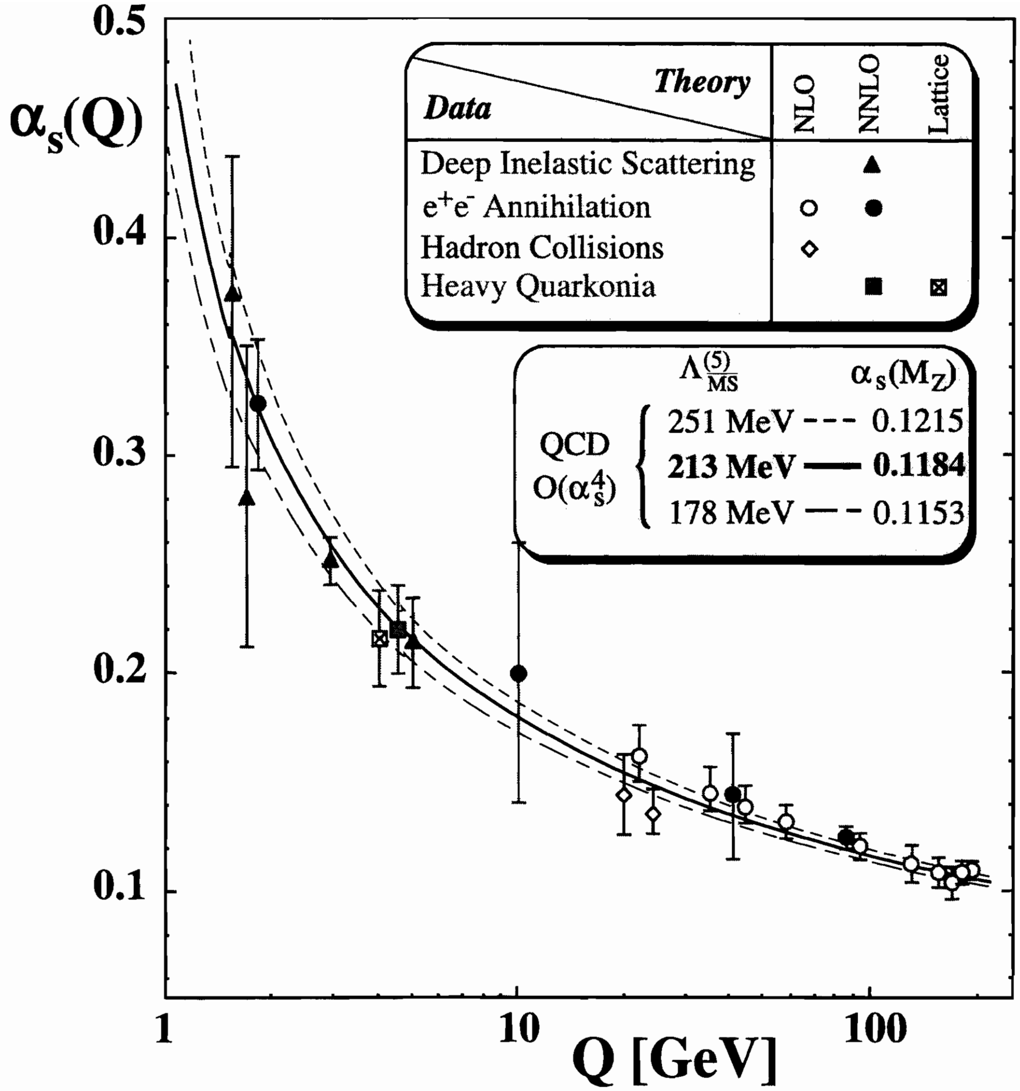

Figure 1-2: Summary of $\alpha_s(Q)$, taken from [16]. This is not by itself a direct experimental observable and has to be extracted by comparing experimental results to QCD calculations. The type of experiment is indicated for each data-point, along with information on the corresponding theoretical calculation used. The uncertainties shown are combined experimental and theoretical uncertainties.

single quark. Quantitatively, in the QCD string model, the interaction between quarks at large distances is modelled as a linear potential $V = \kappa r$, with the value of $\kappa \approx$ 1 GeV/fm obtained from experiment. This model of a relativistic string connecting massless quarks reproduces the observed relations between angular momentum and mass for hadronic states.

1.2.4 Evidence for Gluons

Deep-inelastic electron-proton scattering experiments only probe the quarks inside the proton (gluons, being uncharged, do not couple to the virtual photon). However, if the experimental results are integrated, it is found that the quark contribution to the momentum of the proton is only about 50% - the remaining ≈50% is believed to come from gluons.

But the strongest evidence for the existence of gluons is considered to be observations of three-jet events in collider experiments. A collision event produces a $q\bar{q}$ pair, one of which occasionally radiates a high-energy gluon that produces a third jet of hadrons. The frequency of three-jet relative to two-jet events is determined by the probability of gluon emission, which in turn is related to the strong coupling constant α_s. The experimental measurements confirm the predictions of QCD. The angular distributions in these three-jet events also support the hypothesis that gluons are vector (spin-1) bosons.

1.2.5 Lattice QCD

Calculations in quantum field theories are typically performed using a perturbative expansion in powers of the coupling constant. However, for many of the processes of interest in QCD the momentum scales are small, resulting in a strong coupling constant $\alpha_s > 1$; the perturbation series therefore fails to converge and accurate predictions are not possible. The technique of lattice QCD [17, 18] has been developed as a tool for doing QCD calculations in this non-perturbative regime.

The theory is formulated on a discrete lattice of space-time points. Physically meaningful results will be recovered in the limit as the lattice extent tends to infinity and the spacing goes to zero. It is an important fact that this lattice formulation preserves the gauge invariance of the original theory. The space-time is converted to a Euclidean signature by using imaginary time ($t \longrightarrow \imath t$); this allows the evaluation of path integrals using the Monte Carlo technique of 'importance sampling.' The path integrals can also be related to the partition function of the system, and from this thermodynamic quantities such as energy density and entropy can be extracted.

1.2.6 Chiral Symmetry in QCD

The helicity of a particle is defined as the projection of its spin-vector along its direction of motion: right-handed particles have a spin-vector parallel to their momentum vector, while for left-handed particles the two are anti-parallel. In the massless limit, the left-handed and right-handed components of the quark wavefunction are completely independent of each other and the QCD Lagrangian is said to exhibit chiral symmetry.[5]

However, the QCD vacuum does not possess chiral symmetry. Consider a $q\bar{q}$ pair created from the vacuum, with zero total linear and angular momentum. As shown in Fig. 1-3, this pair has a net chiral charge. If the quarks are light or massless, then, since interactions between them are strong, it is energetically favourable to create such pairs, so the QCD vacuum will consist of a condensate of $q\bar{q}$ pairs which breaks chiral symmetry.

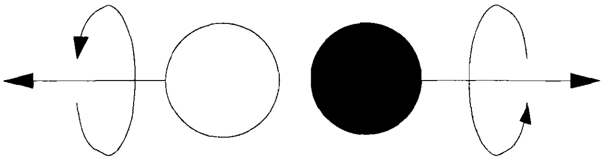

Figure 1-3: Illustration of how a quark-antiquark pair, created from the vacuum with zero total linear and angular momentum, can have a net chiral charge.

When the Lagrangian of a theory possesses a symmetry but the vacuum state does not, the symmetry is said to be spontaneously broken. According to Goldstone's

---

[5] 'Chiral' comes from the Greek word *chier*, meaning 'hand'.

theorem [19], spontaneously broken continuous symmetries produce massless scalar particles called Goldstone bosons.

Thus, for the massless limit of QCD, one would expect the existence of zero-mass Goldstone bosons arising from the spontaneous breaking of chiral symmetry. In real QCD, however, quarks are not massless, so chiral symmetry is not a true symmetry of the theory. But the real quark masses ($m_u, m_d \approx 5 - 10$ MeV) are actually small compared to the fundamental mass-scale of QCD ($\Lambda \approx 200$ MeV), so chiral symmetry is still an approximate symmetry of the strong interaction and this approximate symmetry is still spontaneously broken by the vacuum. Thus it is possible to identify the lightest strongly-interacting bosons, the pions, as pseudo-Goldstone bosons.

1.3 The QCD Phase Diagram

The nature of the strong interaction leads to a rich array of possible phases of QCD matter [20], as illustrated by the phase diagram, Fig. 1-4. The four main phases are the nucleus, hadron gas, color superconductor and quark-gluon plasma. The two variables used to describe the phases are the temperature T and the baryon chemical potential μ_B, defined as the change in energy of the system when one baryon is added.

It should be noted that this phase diagram is for three-flavour QCD: u,d and s. The masses of the charm, bottom and top quarks (see Table 1.1) are too large for them to play a role here. The masses of up, down and strange quarks are taken to be $m_u \approx 5$ MeV, $m_d \approx 10$ MeV and $m_s \approx 100$ MeV.

1.3.1 The Nucleus and Hadron Gas phases

The most familiar phase of strongly-interacting matter is the nucleus. This is a low temperature phase. It is an empirical fact that the nuclear density remains approximately constant, so when a nucleon (mass $m_N \approx 940$ MeV) is added to a nucleus, the change in energy of the system is:

$$\Delta E \approx m_N - \text{Binding Energy} \approx 940\text{MeV}$$

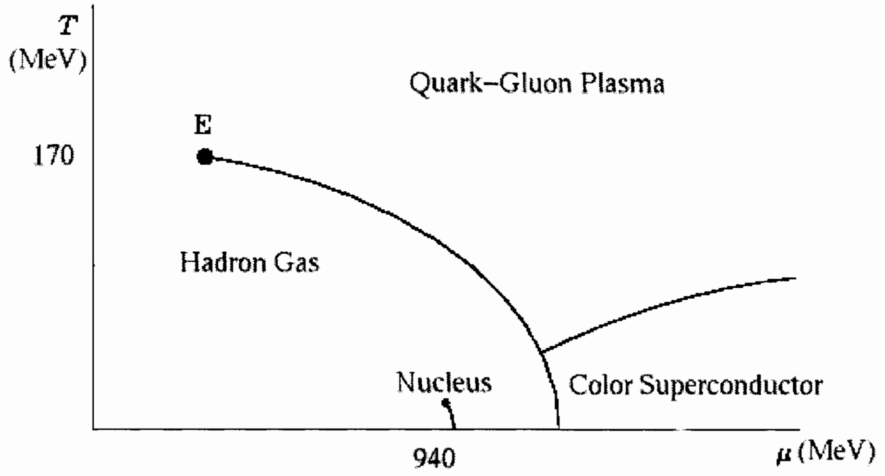

Figure 1-4: Schematic representation of the phase diagram of Quantum Chromodynamics. The vertical axis is the temperature of the system, the horizontal axis is μ_B, the baryon chemical potential.

Hence the nucleus is located on the phase diagram at $T \approx 0$, $\mu_B \approx 940$ MeV.

At temperatures higher than the binding energy of a nucleon (1-10 MeV), the nucleus dissolves into a gas of hadrons. The pion-exchange forces between colorless nucleons in a nucleus are analogous to the induced-dipole Van der Waals forces between electrically-neutral molecules, so a liquid-gas phase transition is expected for nuclei. However, this question has not yet been resolved experimentally [21, 22].

1.3.2 Color Superconductivity

At low temperature and high baryon density, the inter-particle spacing is small, and hence, by asymptotic freedom, the interactions are weak. Quarks are no longer bound in hadrons but instead form Fermi seas. As in solid-state physics, the energy of the ground state can be lowered by the formation of Cooper pairs, if there exists any attractive interaction between the fermions which would allow them to form bosonic 'quasi-particle' states.[6] In QCD, there are attractive interactions between pairs of

---

[6] Note that only quarks near the top of the Fermi surface are affected by the formation of Cooper pairs. Quarks in the bulk of the Fermi sea are 'Pauli-blocked': all accessible states are already filled, so by the Exclusion Principle, there can be no interactions between these quarks.

quarks, as can be seen for instance by single-gluon exchange diagrams, which results in the formation of a condensate of diquark Cooper pairs. These diquarks[7] cannot be color-singlets and this gives rise to *color superconductivity* [25].

At temperatures greater than the strength of the diquark interaction, the diquark condensates melt and the system enters the Quark-Gluon Plasma phase. This is estimated to occur at around $T \approx 50$ MeV. It is interesting to note that this temperature is much higher, even in relative terms, than in solid-state superconductors. This is because the Cooper pairing mechanism in color superconductors comes from the primary (color) interaction, while in solid-state superconductors it is a secondary effect, arising from electron spin-pairing mediated by the atomic lattice.

At large enough baryon density that $\mu_B \gg m_s$, the strange quark mass can be neglected and $SU(3)$ flavour symmetry is approximately restored. The most favourable diquark pairing is then one which is invariant only under simultaneous transformations in both $SU(3)$ color and flavour. In a sense, color and flavour become 'locked' together and this phase is described as Color-Flavour Locked (CFL) matter [26].

At intermediate μ_B, the non-zero strange quark mass must be taken into account and this complication renders the phases in this region much harder to analyse: possibilities include a crystalline color superconductor [27, 28] or a 'gapless' CFL phase [29].

Although chances of creating a color superconductor in the laboratory remain slim (how does one compress matter beyond nuclear density?), it may be possible to observe this phase in astrophysical systems. Gravitational collapse of massive stars can overcome the strong interaction to create a super-dense state. While no such 'quark star' has yet been conclusively discovered, the search is ongoing and many phenomenological studies are being done to identify potential observable effects that would indicate the presence of a color-superconducting state in the star's core - see [30] for a review.

---

[7] This concept of correlated diquarks is now thought to exist not only in color superconductors but in hadrons in general at low temperatures, and has been recently used in analysis of the newly-seen exotic pentaquark states [23] and in other aspects of hadron phenomenology [24].

1.3.3 The Quark-Gluon Plasma

At high temperatures, the interactions that bind quarks into hadrons can be overcome, and the ground state of the system will be a soup of deconfined quarks and gluons called the *quark-gluon plasma* (QGP).

This phase transition is expected from first principles, and is supported by lattice QCD studies. An example of this is shown in Figure 1-5. The quantity ϵ/T^4 counts the number of degrees of freedom in a thermodynamic system: above the critical temperature T_c, we see a sharp rise followed by a plateau, indicating the system becomes deconfined. The value of this critical temperature is believed to be $T_c \approx$ 170 MeV (see e.g. [31]).

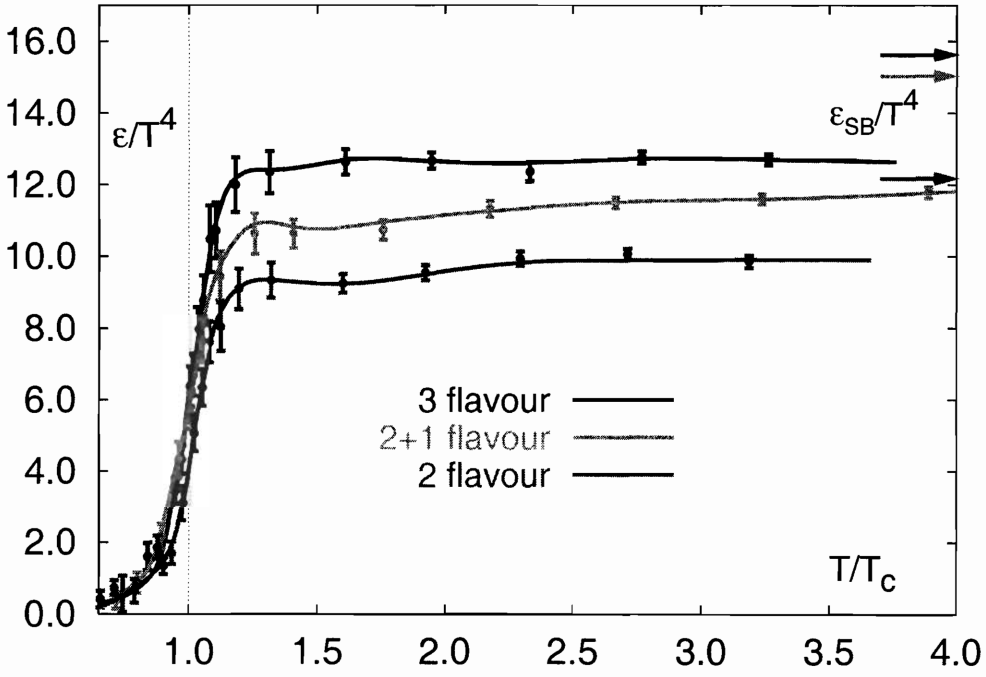

Figure 1-5: Illustration of transition from hadronic to QGP state, from lattice QCD calculations. The quantity ε/T^4 counts the number of degrees of freedom in a system; a sharp rise is seen as the system temperature crosses the critical temperatute T_c. ε_{SB}/T^4 represents the Setfan-Boltzmann limit of a non-interacting system. Configurations with different numbers/masses of quark flavours are shown. This figure is taken from [32].

Another important phase transition associated with high-temperature QCD is the restoration of chiral symmetry [33]. This process is illustrated in Figure 1-6, in analogy with the ferromagnetic transition in a metal. As discussed in Section 1.2.6, the QCD vacuum spontaneously breaks chiral symmetry. The expectation value $\langle q\bar{q}\rangle \neq 0$ and for all quark-antiquark pairs is oriented in the same direction in flavour space (represented by the arrows in the figure), analogous to spin-ordering in metals at zero temperature. As the temperature increases so does the disorder, until a critical temperature is reached where there is no longer any preferred direction and thus the expectation value $\langle q\bar{q}\rangle = 0$, meaning chiral symmetry is restored (the metal analogy is the disappearance of ferromagnetism for temperatures above the Curie temperature).

No fundamental reason is known why the deconfinement and chiral symmetry restoration transitions should occur at the same temperature, but lattice results appear to indicate that they do: this issue is still under investigation [34, 35].

The nature of the phase transition has been the focus of much theoretical analysis (see [36] for a review of the arguments, for both the deconfinement and chiral symmetry restoration transitions). The consensus view, supported by lattice QCD results, is that the transition is first-order, changing to a smooth cross-over at the critical end-point E.

This critical-point should have clear phenomenological signatures, such as enhanced fluctuations of thermodynamic variables. No such signature has yet been conclusively observed, but there are suggestive preliminary experimental measurements [37] at a μ_B in rough agreement with the location of the critical-point estimated from lattice QCD studies [31].

1.4 Relativistic Heavy-Ion Collisions

The most effective method for probing the high temperature regions of the QCD phase diagram and potentially creating a quark-gluon plasma is to collide large nuclei at relativistic energies.

Fixed target experiments at the Alternating Gradient Synchrotron (AGS) at Brookhaven National Laboratory and CERN's Super Proton Synchrotron (SPS) have

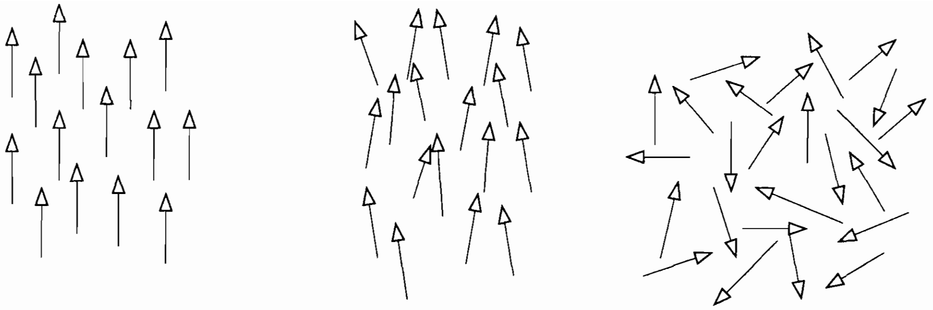

Figure 1-6: Illustration of the restoration of chiral symmetry to the QCD vacuum at a critical temperature T_c. The temperatures of the three panels, from left to right, are: $T = 0$; $T \neq 0$ and $T = T_c$.

studied collision systems from light ions up to Au+Au and Pb+Pb at a range of beam energies from 2 to 158 GeV per nucleon. The collider era for heavy-ion physics began with the Relativistic Heavy-Ion Collider (RHIC) at Brookhaven in 2000. RHIC has provided Au+Au, d+Au and Cu+Cu collisions at centre-of-mass energies from 19.6 to 200 GeV per nucleon pair. In 2008, the Large Hadron Collider (LHC) at CERN is expected to raise the energy regime even further by colliding Pb+Pb at $\sqrt{s_{NN}} = 5.5$ TeV.

Overviews of the experimental results and theoretical work in the field of heavy-ion physics can be found in the proceedings of the 'Quark Matter' series of international conferences on ultra-relativistic nucleus-nucleus collisions [38, 39, 40]. The four RHIC experimental collaborations have also each recently released a 'white paper' summarizing their observations [41, 42, 43, 44].

Even if a QGP is created in a heavy-ion collision, it can only be a short-lived phase and the system will have hadronized once more before it can be directly detected; this makes inferring its creation difficult.

A wide variety of experimental observables have been studied, such as charged particle multiplicity, azimuthal asymmetry, transverse momentum distributions and jet correlations. Many of the results were unexpected and no unambiguous QGP signatures have been seen, but there are strong indications that the system possesses an energy density above the critical value suggested by lattice QCD and that there are interactions between the produced particles, leading to the establishment of thermodynamic equilibrium. There are also signs that the system may exhibit color deconfinement.

1.4.1 Kinematic Variables

In high-energy physics, it is customary to parametrise the longitudinal axis in terms of *rapidity* (y). For a Lorentz transformation along the z-axis with boost parameter β_z, rapidity is defined by:

$$\beta_z = \tanh y \tag{1.7}$$

and it is a convenient variable because it is additive under Lorentz boosts. In this sense, rapidity in relativistic mechanics is analogous to velocity in non-relativistic mechanics, since velocity is additive under Galilean transformations.

For a particle with energy E and longitudinal momentum p_z, the rapidity can be written as:

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right) \tag{1.8}$$

In the high-energy limit where the particle mass is small compared to its momentum, rapidity can be approximated by *pseudorapidity* (η), which is defined in terms of the polar angle θ as:

$$\eta = -\ln[\tan(\theta/2)] \tag{1.9}$$

Note that for a particle travelling precisely along the z-axis, the pseudorapidity is infinite while the rapidity is always finite.

For a particle of rest mass m_0, the momentum transverse to the collision axis, p_T, can be used to define a *transverse mass* $m_T = \sqrt{m_0^2 + p_T^2}$. These give rise to the following useful relations:

$$E = m_T \cosh y \tag{1.10}$$

$$p = p_T \cosh \eta \tag{1.11}$$

where p is the total momentum of the particle ($E^2 = m_0^2 + p^2$).

1.4.2 Impact Parameter and Collision Centrality

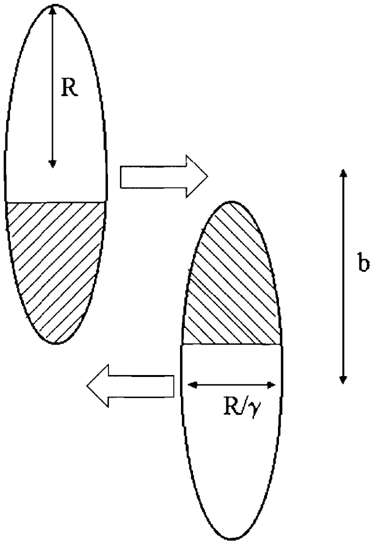

Figure 1-7: Illustration of the impact parameter b in a heavy-ion collision. The cross-hatched regions denote the *participant* nucleons; the remainder are *spectators*. The nuclei are Lorentz-contracted by a factor of γ.

As illustrated in Figure 1-7, collisions of heavy-ions possess an impact parameter which defines the *centrality* of the collision. At these relativistic collision energies, only nucleons in the overlap region will be involved in the interaction: these are called *participants* and the rest are *spectators.*

The number of participating nucleons N_{part} is an important way of characterizing a heavy-ion collision: many observables vary with the collision centrality and knowledge of this gives insights into the underlying dynamics. It is also useful to know the number of binary nucleon-nucleon collisions N_{coll}.

N_{part} and N_{coll} are calculated from a Glauber model of the collision [45, 46]. In this picture, a nucleus-nucleus collision is treated as a superposition of individual nucleon-nucleon collisions, using the free nucleon-nucleon inelastic cross-section obtained from *pp* scattering data. The mass density distribution of the nucleus is assumed to follow the Woods-Saxon form of a smoothed square well. Within this framework, Monte Carlo simulations, such as that implemented in the 'event generator' HIJING (Heavy

Ion Jet Interaction Generator) [47] are a natural way to calculate N_{part} and N_{coll} for a heavy-ion collision with a given impact parameter.

Scaling with N_{part} is expected to dominate for 'soft' (low-p_T) particle production, according to the superbly-named *wounded nucleon* model [48], where a nucleon which has interacted once is assumed to be unable to interact again.

An alternative picture is to regard a nucleus-nucleus collision as an independent superposition of individual nucleon-nucleon collisions. If this were the case, then heavy-ion observables would be expected to scale with N_{coll}. Direct photon production, for example, has been found to exhibit N_{coll}-scaling [49].

1.4.3 Particle Multiplicity in Heavy-Ion Collisions

The simplest measurement that can be made for a heavy-ion collision is to count the number of particles which are emitted.

The PHOBOS collaboration has measured the distribution of charged particles along the longitudinal axis in Au+Au collisions at centre-of-mass energies of 19.6, 130 and 200 GeV per nucleon [50]. These results are displayed in Figure 1-8. The total number of charged particles emitted for the most central collisions at each energy are summarised in Table 1.2.

$\sqrt{s_{NN}}$	N_{ch}
19.6 GeV	1680 ± 100
130 GeV	4170 ± 210
200 GeV	5060 ± 250

Table 1.2: Total number of charged particles produced by central Au+Au collisions at different energies, as determined by the PHOBOS experiment.

Estimating the Energy Density

The energy density of the system formed in a heavy-ion collision can be estimated from the measured particle multiplicity. A common approach follows that of Bjorken [51], which focuses on the mid-rapidity region and presumes that particle production as a function of rapidity exhibits a Lorentz-invariant 'central plateau' structure.

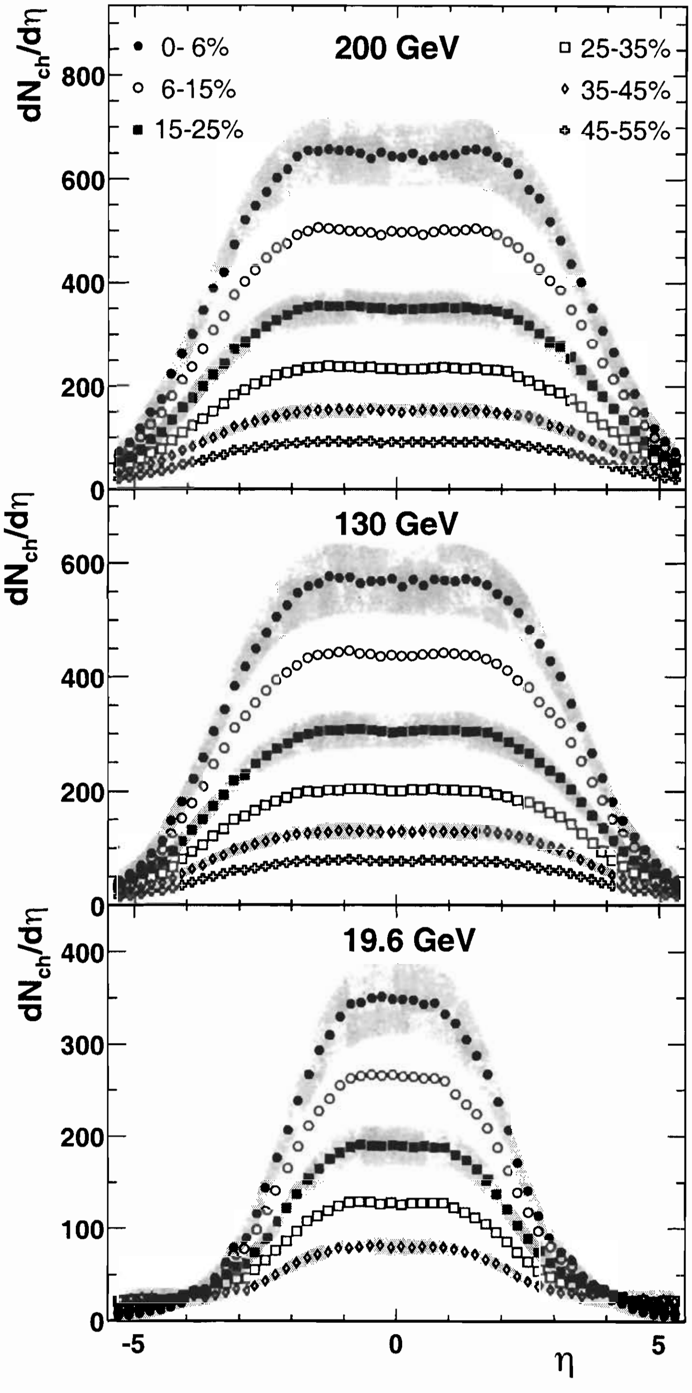

Figure 1-8: Charged particle pseudorapidity distributions in Au+Au collisions at centre-of-mass energies of 19.6, 130 and 200 GeV, measured by the PHOBOS collaboration. At each energy, results are shown for collisions with different mean impact parameters.

The colliding system is hypothesized to reach thermal equilibrium in a time of about 1 fm/c and thereafter follow a hydrodynamic longitudinal expansion. The transverse dimension of the system is taken to be the size of the original nuclei. With a mean energy per particle of $\langle E \rangle \sim 0.5$ GeV, the PHOBOS multiplicity results for central Au+Au collisions at $\sqrt{S_{NN}} = 200$ GeV indicate an energy density of roughly 5 GeV/fm^3. Although not a rigorous calculation, this estimate is significantly greater than the threshold energy density for creating a quark-gluon plasma expected from lattice QCD studies, which is $\varepsilon \sim 1$ GeV/fm^3.

1.4.4 Azimuthal Anisotropy and Interactions

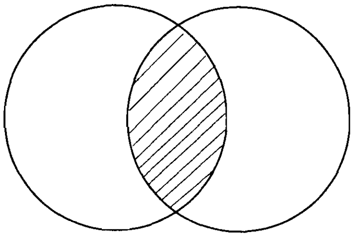

Figure 1-9: Transverse slice-view of a collision with non-zero impact parameter, showing the initial azimuthal asymmetry of the system.

Figure 1-9 shows the view along the longitudinal axis of a heavy-ion collision with non-zero impact parameter. The initial state of the colliding system clearly displays a high degree of azimuthal anisotropy. However, if each point in this system simply acts as an independent source of particle production, with no interactions between the particles, then the observed azimuthal particle distribution would be perfectly symmetric - the collective effects of interactions are required to translate the initial state asymmetry into final state asymmetry.

Therefore, experimental measurements such as [52] which demonstrate a significant azimuthal anisotropy in the final state are clear indicators that the system formed in relativistic heavy-ion collisions is a highly-interacting one.

1.4.5 Thermal Models of Heavy-Ion Collisions

If the system is highly-interacting, it is feasible that it could reach thermal equilibrium, at which point it can be characterised by macroscopic thermodynamic quantities such as temperature, pressure and chemical potentials. Particle populations will then follow a statistical distribution, with the number density n_i of species i being given by:

$$n_i = \frac{g_i}{(2\pi)^3} \int \frac{d^3p}{e^{(E_i - \sum_j \mu_j X_j^i)/T} \pm 1} \tag{1.12}$$

where g_i is the spin-degeneracy, $g_i = 2J_i + 1$ for a particle with spin J_i

$E_i = \sqrt{p^2 + m_i^2}$ is the energy of a particle with momentum p and mass m_i

X_j^i is the extensive variable for species i appropriate for the chemical potential μ_j

T is the temperature of the ensemble

and ± 1 is the quantum mechanical contribution for ensembles of identical particles: $+1$ for fermions and -1 for bosons.

Only hadrons constructed from u, d or s quarks are typically considered in thermal models. Therefore, there are three chemical potentials: μ_u, μ_d and μ_s; alternatively, the same information can be expressed in different degrees of freedom, using the three conserved quantities of charge Q, strangeness S and baryon number B. Thus $\sum_j \mu_j X_j^i = \mu_Q Q_i + \mu_S S_i + \mu_B B_i$.

Thermal models (see [53] for a recent review) typically have μ_B and T as free parameters – the rest are constrained by conservation laws. The fitted models have been found to give a good description of a multitude of measured particle yields.

Since they allow a determination of (μ_B,T), thermal models can be used to estimate the location on the QCD phase diagram of the system formed in a heavy-ion collision. For example, experimental measurements of antiparticle-to-particle ratios by the PHOBOS collaboration ([54]) have been used in conjunction with a thermal model ([55]) to estimate that central Au+Au collisions at $\sqrt{s_{NN}}$ =200 GeV produce a system with a temperature of $T \sim 170$ MeV and a baryochemical potential of

$\mu_B \sim 25$ MeV. Thus collisions at RHIC are seen to be approaching the baryon-free regime ($\mu_B = 0$) where lattice QCD calculations are expected to be reliable and may be able to be directly compared to experimental results.

Thermal models for heavy-ion collisions at the SPS also have a *strangeness suppression factor* γ_s, which is needed to accurately describe the yields of strange particles. There is a factor of γ_s for each strange quark in the hadron. The common explanation for this suppression factor is that strange quarks, having a higher mass than u or d, are not fully equilibrated. At RHIC energies, however, this factor has been found to be unnecessary (or equivalently, $\gamma_s = 1$); this is attributed to the higher collision energy being sufficient to produce a complete equilibrium of all three quark flavours.

It should be noted, however, that a statistical description of particle yields has also been found to work well for elementary e^+e^- collisions ([56]) where thermal equilibrium is not expected to be established. An interpretation of this unexpected success is that the process of hadronisation might statistically populate the available phase space. Thus, the agreement of the observed particle yields with the predictions of a thermal model for heavy-ion collisions is a necessary but not sufficient condition for demonstrating thermal equilibrium in the complex system formed by the collision.

1.4.6 Baryon Stopping

The energy available for particle production, and hence the subsequent evolution of system formed in a heavy-ion collision, is determined by the energy loss of initial nuclei, referred to as *baryon stopping.*

A phenomenological approach based on data from proton-nucleus collisions with 100 GeV projectiles on a variety of targets ([57]) has been used to estimate that in central collisions with a heavy-ion, the average rapidity loss of a projectile nucleon is on the order of 2 units. Studies of net-baryon rapidity distributions for central Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV have measured a similar mean rapidity loss [58], implying that around 75% of the collision energy is available for particle production.

The primary candidate mechanism which could explain such a large degree of

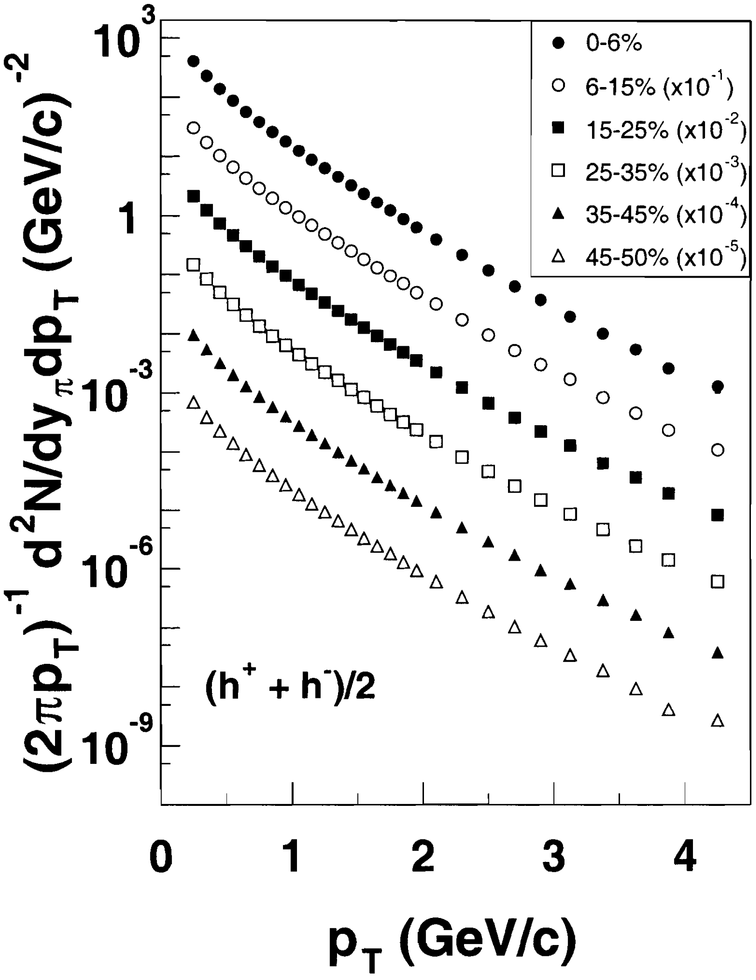

Figure 1-10: Transverse momentum distributions of non-identified charged hadrons from Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV.

baryon transport from beam rapidity proposes that baryon number is carried not by valence quarks but instead by *gluon junctions* [59]. In this model, a baryon is pictured as three valence quarks linked by gluon strings which join at a central 'Y' junction, and it is this junction which traces baryon number. In a high energy heavy-ion or nucleon-nucleon collision, this gluon configuration is believed to be more easily transported from beam rapidity than valence quarks, which tend to pass through and populate the fragmentation regions.

1.4.7 Transverse Momentum Distributions

A strong indication that the system formed in Au+Au collisions at RHIC energies is deconfined comes from the measurement of the transverse momentum p_T distributions of charged hadrons [60, 61, 62, 63, 64, 65]. The spectra measured by PHOBOS at $\sqrt{s_{NN}} = 200$ GeV are shown in Figure 1-10. At high p_T, the particle yields were found to be strongly suppressed relative to the expectation based on measured p+p spectra scaled by the number of binary nucleon-nucleon collisions, N_{coll}.

The standard explanation for this is that the heavy-ion collision produces a dense partonic medium, where high-p_T partons undergo significant energy loss through gluon radiation and emerge as hadrons with a much lower transverse momentum. Thus the high-p_T hadron yields are suppressed. This is known as 'jet-quenching.' Calculations can be done in a perturbative QCD framework where the treatment of individual parton-parton scattering events (with large enough momentum transfers that $\alpha_s < 1$) is combined with initial parton distribution functions and 'fragmentation functions,' obtained from e^+e^- collision data, that describe the production of final-state hadrons from the scattered partons. This approach is reviewed in, for example, [66]. The free parameter in such theories is the gluon density of the medium, and fitting to the experimental data gave values for this density that strongly indicate a color-deconfined system.

An alternative explanation for the high-p_T suppression attributed it to an initial state effect, the Color Glass Condensate[8] (CGC) [67]. In this semi-classical approach to QCD at high density, low-momentum gluons are not independent but instead fuse together, creating a coherent system defined by a dimensional parameter, the *saturation scale*. This picture was successful in explaining many features of the bulk particle production in Au+Au collisions [68].

To distinguish between these two scenarios, RHIC collided d+Au ions, with the motivation that an initial state effect such as the CGC should still be present in d+Au, but there should be no dense partonic final state formed, and thus no parton

---

[8] Some theorists have argued that the Color Glass Condensate should not properly be called a 'condensate', because no symmetries are broken – I'm not going to get into that.

energy loss. The experimental results [69, 70, 71, 72] were clearly in favour of the jet-quenching scenario, as no high-p_T suppression was observed in d+Au collisions.[9]

1.4.8 The Proton/Pion Puzzle

Measurements from the PHENIX collaboration of identified particle transverse momentum distributions at 200 GeV Au+Au collisions [74] produced a surprise: at intermediate p_T ($2 < p_T < 4$ GeV/c) in central collisions, the proton yield was found to be comparable to the pion yield. This is illustrated in Figure 1-11.

This proton/pion ratio is much larger than in proton-proton collisions, and the result has been referred to as the 'baryon anomaly.'

By considering the centrality dependence of proton and pion p_T spectra separately (Figure 1-12), it is seen that the large proton/pion ratio at intermediate p_T arises because the pion yields are strongly suppressed as described in the previous section, while protons are consistent with N_{coll}-scaling in this p_T range.

Preliminary results from the PHOBOS collaboration [75] on identified particle p_T spectra in d+Au collisions at the same energy showed no signs of a similar enhancement in proton yields, indicating that the large proton/pion ratio observed in Au+Au collisions at 200 GeV is most likely not an initial state effect due to having a nuclear target, but rather a final-state effect that arises from the dense partonic system formed in these high-energy heavy-ion collisions.

Two physical scenarios have been proposed to explain this baryon enhancement in Au+Au collisions.

The first [76] combines the gluon junction mechanism for baryon stopping (described in Section 1.4.6) with jet-quenching (see Section 1.4.7). The gluon junction mechanism allows for signficant baryon and antibaryon yields at mid-rapidity; the suppression of pion production due to parton radiative energy loss in the dense deconfined medium then enhances the relative ratio of protons and antiprotons to charged hadrons. This model predicts that the baryon enhancement should be limited to

---

[9] It has been proposed, however, that the CGC is still an influence on particle production at forward rapidities in d+Au collisions [73].

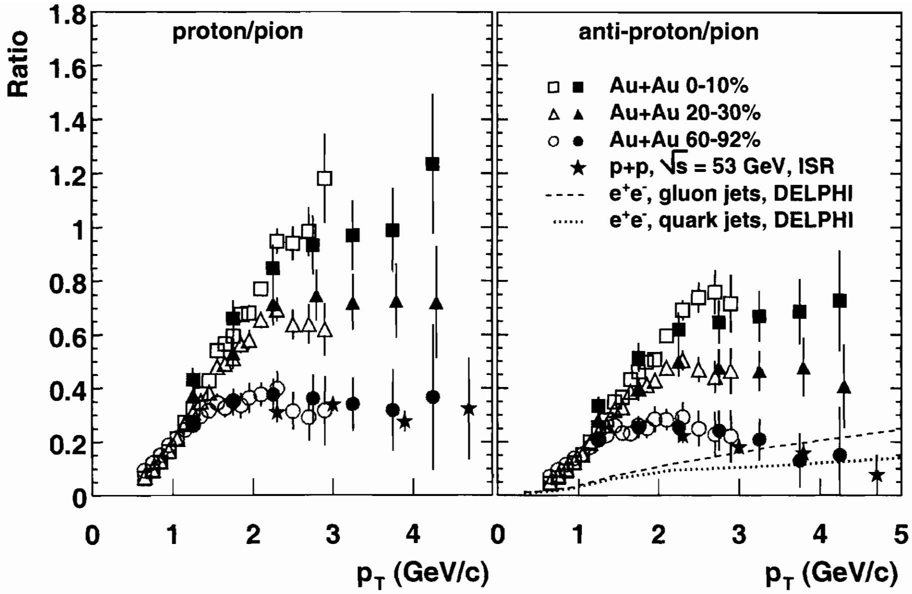

Figure 1-11: Proton/pion ratio as a function of p_T for a variety of collision systems. This figure is taken from [74].

$p_T < 6$ GeV/c, after which the relative yields of protons to charged hadrons should be similar to those in p+p collisions.

The second scenario involves parton coalescence or recombination as the dominant source of hadron production at intermediate p_T [77, 78, 79]. In the standard perturbative QCD picture of a high-energy collision, hadrons are produced from partons after interaction by a process called 'fragmentation,' where the parton radiates gluons until enough $q\bar{q}$ pairs are formed that hadronisation can occur. This is the only possible hadronisation process available to partons in a vacuum. In a dense partonic system, however, there may already be enough quarks around that hadronisation can simply occur by a process of quark coalescence. It has been shown that for a thermalised, exponential parton p_T spectrum, recombination is more effective than fragmentation at producing intermediate p_T hadrons. Recombination also enhances the baryon/meson ratio at intermediate p_T. Consider a simple example of a baryon and a meson at the same $p_T = 3.0$ GeV/c: in a recombination picture, the baryon would have come from three quarks with $\langle p_T \rangle = 1.0$ GeV/c, which, for a steeply-falling p_T-spectrum, are

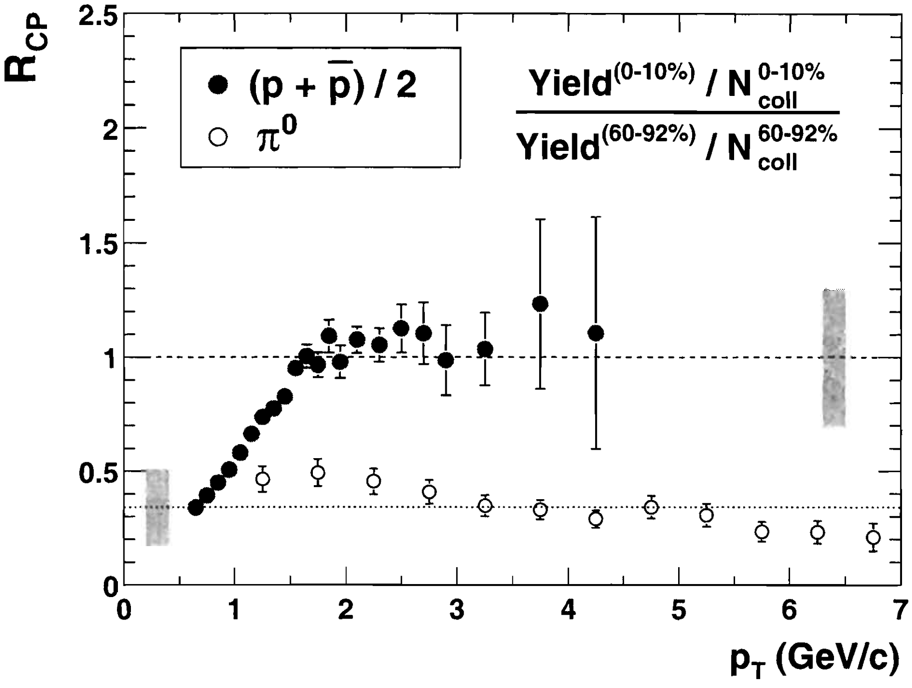

Figure 1-12: Ratio of central to peripheral p_T spectra for protons and pions from 200 GeV Au+Au collisions. This figure is taken from [74].

more numerous than the $\langle p_T \rangle = 1.5$ GeV/c partons that would comprise the meson.

1.4.9 Goal of this Thesis

This thesis presents the first results on identified particle transverse momentum distributions from Au+Au collisions at a new collision energy of $\sqrt{s_{NN}} = 62.4$ GeV. We will use these spectra to investigate how the relative behaviour of protons and pions at intermediate p_T evolves with collision energy. This will provide important data for helping to understand the processes which govern particle production in the complex system formed in relativistic heavy-ion collisions.

Chapter 2

The PHOBOS Experiment at RHIC

The Relativistic Heavy-Ion Collider (RHIC) at Brookhaven National Laboratory, NY, is a machine for colliding nuclei at the highest centre-of-mass energies ever achieved in a laboratory, in the hope of creating and studying the quark-gluon plasma.

PHOBOS is one of four experiments at RHIC. The experiment is named after the larger moon of the planet Mars[1]. The choice of name reflects the historical origin of the experiment as the off-shoot of an earlier proposed experiment named MARS (Modular Array RHIC Spectrometer).

The main mission of PHOBOS is to provide a global characterisation of the heavy-ion collisions, with about 1% of the particles undergoing more detailed analysis in the magnetic spectrometer and Time-of-Flight system.

2.1 The Relativistic Heavy-Ion Collider

Construction of the RHIC facility [80] was completed in 1999; first collisions were obtained on 12 June 2000, and the design goal of collisions at a beam energy of 100 GeV per nucleon was achieved on 18 July 2001.

---

[1] Phobos, the moon of Mars, takes its name from a blend of Greek and Roman mythology: Phobos was the Greek god of fear and was a son of the War-God Ares, whom the Romans named Mars.

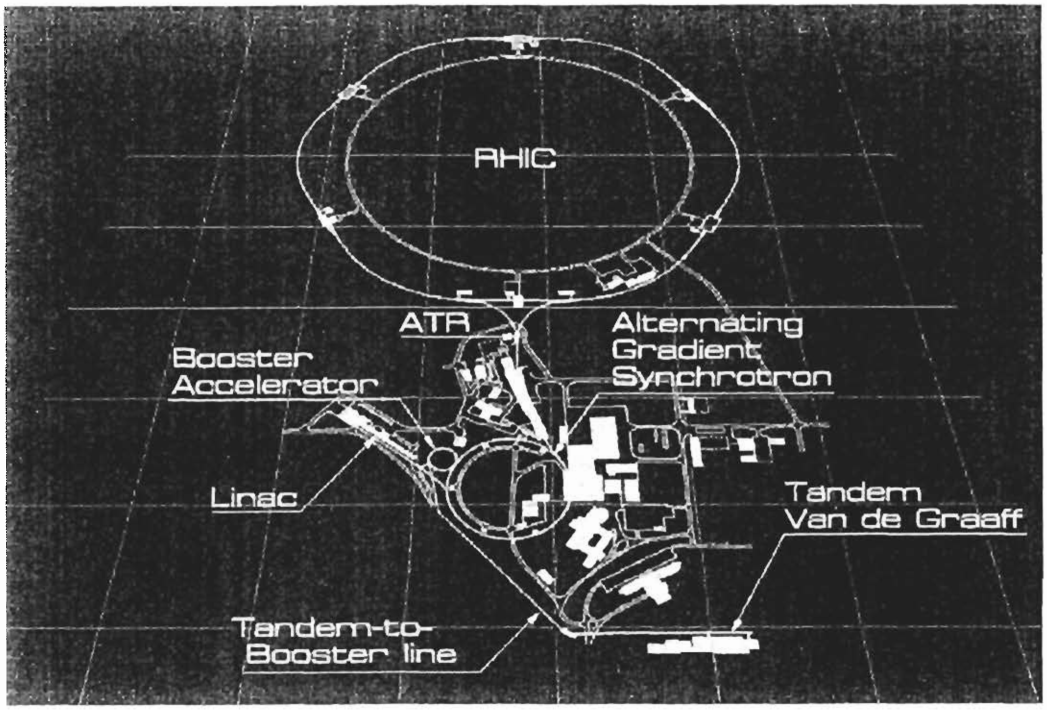

Figure 2-1: The Relativistic Heavy-Ion Collider (RHIC) acceleration complex.

Figure 2-1 shows the components of the accelerator complex. RHIC heavy-ion collisions begin with the production of gold ions by a pulsed sputter source at the Tandem Van der Graaff machine. Negative ions formed by electron capture are accelerated to the high voltage terminal of the Van der Graaff, where a metal foil is used to partially strip them of some electrons; the now positively-charged ions are then accelerated back again in the second stage of the Tandem, reaching an energy of 1 MeV per nucleon. The ions undergo further foil-stripping on exit from the Tandem and the charge-state +32 is picked by magnetic selection and sent to the Booster synchrotron which accelerates it to 95 MeV per nucleon.

From the Booster, ions are stripped to +77 and injected into the Alternating Gradient Synchrotron (AGS). They are accelerated to 10.8 GeV per nucleon at this stage, and upon exit are fully stripped to +79. The AGS-to-RHIC transfer line (ATR) carries these ions to the RHIC ring for final acceleration.

RHIC consists of two quasi-circular concentric accelerator rings, with a circumference of 3.8 km. It is built from six arc sections and six straight sections for beam insertion or interaction. The bending magnets in the arc sections are superconducting dipole magnets with a field strength of 3.458 T and a bending radius of 243 m. The operating temperature is ≈4.5 K and they run at a current of 5050 A. Quadrupole magnets are used to focus the beams.

Typically, around sixty ion bunches are present in the rings simultaneously. Some gaps are left in the bunch pattern to allow for controlled beam dumps. A 200 MHz RF is used to 'rebucket' the beams after acceleration - this produces a longitudinal beam size of $\sigma_L \approx 25$ cm. The bucket size is much smaller than the separation between bunches, so there are many empty buckets and a *crossing clock* is used to indicate the times when full buckets should be interacting at the collision sites.

The collision event rate at a collider is given by:

$$R = \frac{f n_1 n_2 \sigma_{int}}{4\pi\sigma_x\sigma_y} \tag{2.1}$$

where f is the frequency of bunch collisions, n_1 and n_2 are the number of ions in the colliding bunches, σ_{int} is the interaction cross-section, and σ_x and σ_y measure the transverse beam profile.

At RHIC, bunches collide every few hundred nanoseconds, each bunch has the order of 10^9 ions and the beam size is in the millimetre range. For nuclei, $\sigma_{int} \approx$ 1-10 barns (10^{-28} m^2). The result is a collision rate on the order of tens of kHz.

RHIC also has the capability to collide polarized proton beams for spin physics studies [81]. Protons are generated at the 200 MeV Linac (Linear Accelerator). The use of spin rotators and complicated spin-flipping magnet structures called 'Siberian Snakes' allows RHIC to minimise the effect of depolarizing resonances and operate at by far the highest energies achieved for polarized proton beams.

2.2 Overview of the PHOBOS Detector

The PHOBOS experiment [82] is sited at the ten o'clock position on the RHIC ring. A diagram of the complete PHOBOS detector in 2003 is shown in Figure 2-2.

PHOBOS is comprised of four main sets of sub-detectors: the multiplicity array,

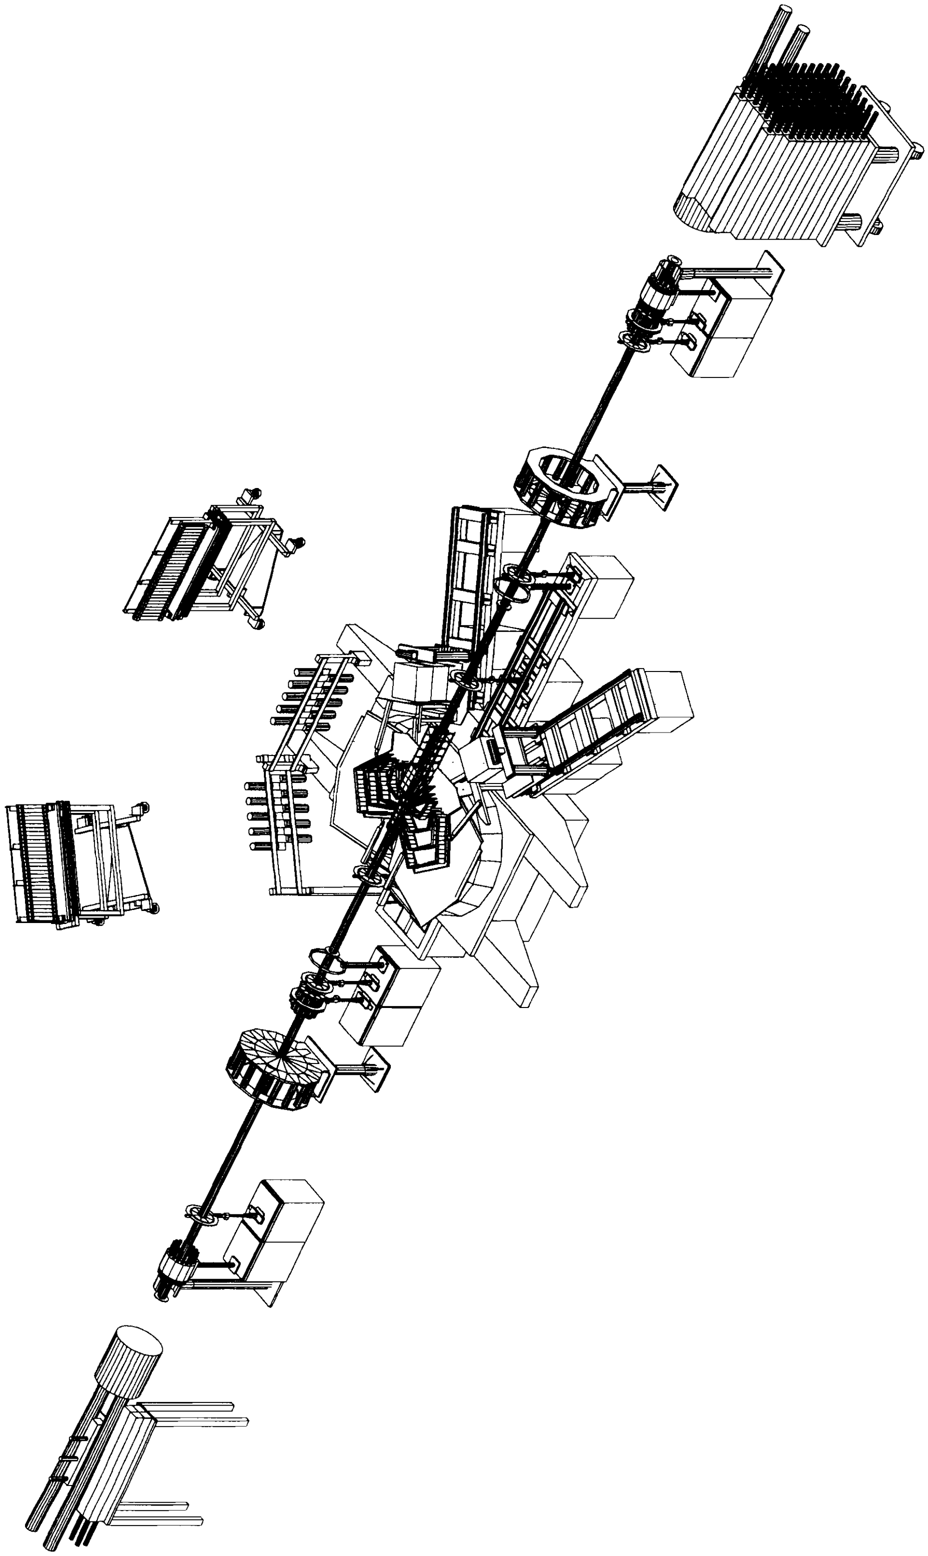

Figure 2-2: The complete PHOBOS Detector in 2003.

consisting of Octagon and Ring silicon detectors; the vertex detector; the silicon Spectrometer and particle-identification system, including the Time-of-Flight walls; and the trigger system, which consists of the Paddle Counters, Zero-Degree Calorimeters, and Spectrometer Trigger. There is also a forward proton calorimeter (PCAL), and for the 2004 run an additional calorimeter, the Spectrometer Calorimeter (not shown in Figure 2-2) was also installed.

2.2.1 PHOBOS Beam-pipe and Coordinate System

The beam-pipe in the PHOBOS interaction region is constructed from Beryllium; the low-Z material helps keep multiple scattering of primary particles and the production of secondary particles to a minimum. The Be beam-pipe is 12 m long, in 3 sections of 4 m, with a diameter of 76 mm and a wall thickness of approximately 1 mm.

The z-axis of the PHOBOS coordinate system is chosen to be the beam-line, with the positive direction defined as the direction of clockwise circulation. It is also desirable to have the positive y-axis correspond to 'up' in physical space. But with these two choices made, in order to have a right-handed coordinate system, the x-axis must defined in an unusual way: when one looks down on a plan of the PHOBOS detector, the positive x-axis is on the left-hand side. Hence the TOF walls are at negative x, for example.

It is very important to accurately know the positions of all detectors in the PHOBOS experiment. Individual silicon modules were surveyed on a survey-station with a precision of 10 μm and the position of the modules relative to their mounting frames is known to 100 μm. The optical survey of the detectors in their final positions within the PHOBOS interaction region has a precision of 100 μm.

2.3 Silicon Sensor Technology in PHOBOS

The PHOBOS multiplicity array, vertex detector and spectrometer are all constructed from silicon pixel detectors. These provide good energy resolution and high-precision position determination for measurements involving charged particles.

2.3.1 Principles of Silicon Detectors

Semiconductors such as silicon or germanium are materials which have a small (of order 1 eV) energy gap between the valence electron band and the conduction band. This gap is called the forbidden band. With a small gap, electrons are easily excited into the conduction band, leaving behind a 'hole' in the valence band. Under the influence of an external electric field, both electrons and holes can move (although electrons have more mobility), generating an electric current. Pure semiconductors are 'doped' by adding impurities. If the impurity has more valence electrons than Si or Ge (which have four), then it acts as an electron donor and increases the number of electrons relative to holes. The semiconductor is then said to be *n-type*, because the dominant charge carrier is negative. On the other hand, if the impurity has less valence electrons, this produces more holes and the semiconductor is called *p-type.*

Silicon charged-particle detectors make use of the special properties that arise when a p-type and an n-type semiconductor form a pn-junction[2]. Electrons from the n-type material will initially diffuse into the p-type and recombine with the holes there (and vice verse for holes from the p-type). The result is a 'depletion zone' where there are no mobile charge carriers. The size of this depletion zone can be greatly increased by reverse-biasing the junction (applying a positive voltage to the n-type and a negative voltage to the p-type).

This depletion zone is the sensitive volume of a silicon detector: when charged-particles pass through it, they produce electron-hole pairs which then get swept out to the electrodes at either side. Since it takes less energy to produce electron-hole pairs in a semiconductor than it does to ionize a gas, silicon detectors produce a larger number of charged signal-carriers than a gas-ionization detector and thus have better energy resolution.

There should in principle be no current across a pn-junction, even when biased, but in practice, a variety of factors such as contamination, etc... contribute to a small, fluctuating current called the 'leakage current'. It is desirable to minimise this

---

[2] A pn-junction cannot be formed by just pressing the two materials together; special fabrication techniques are needed to produce the correct contact.

current, since it contributes to signal noise and can indicate an unhealthy sensor.

A drawback of silicon detectors for high-energy physics is that excessive radiation can cause structural defects in the lattice, which create discrete energy levels in the forbidden band that can trap and reduce the number of charge carriers. Radiation damage also typically produces a rise in the leakage current. Radiation levels in the experiment and the health of the Si detectors need to be carefully and constantly monitored.

2.3.2 Design of PHOBOS Silicon Sensors

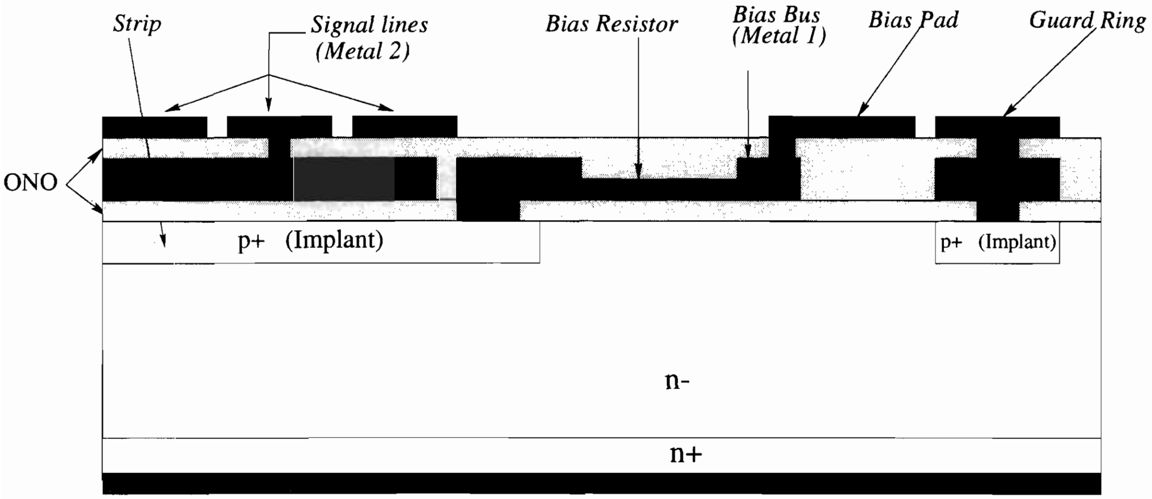

Figure 2-3: Schematic diagram of design of silicon pixel detector used in PHOBOS.

A cross-section of one pad of a PHOBOS silicon sensor [83, 84] is shown in Figure 2-3. The substrate is a 300 μm-thick n-type wafer, with a heavily-doped p+ implant to create the pn-junction. A guard ring surrounds the active area. Bias is applied through a 5 MΩ polysilicon resistor; approximately 70 V is required to fully deplete the sensor.

These sensors use an AC-coupling design, where the charge collected at the p+ implant is coupled to the output via a capacitor. This capacitor is formed from the rectangular p+ implants and the first metal layer, separated by an oxide-nitride-oxide

(ONO) insulating layer that is 0.2 μm thick. The signal from each pixel in the sensor is carried away by the second metal trace-lines, which are all routed along the face of the sensor to a common destination: the bonding pad where the read-out chip is attached. A further ONO layer of 1.2 μm thickness isolates the trace-lines from the first metal layer (except at the single connection point), to minimise signal noise.

The silicon sensors went through extensive tests before being incorporated into the PHOBOS experiment. To find the depletion voltage, the capacitance was measured versus the applied bias, and only sensors with a depletion voltage $\leq$70 V were accepted. The leakage current was also measured as a function of the applied bias: Spectrometer and vertex sensors were required to have $< 5\mu$A at the depletion voltage, while multiplicity array sensors had a less stringent criterion of $< 10\mu$A. The presence of short circuits due to pinhole connections between the two metal layers was tested for, and the capacitance of each channel was also checked and compared to a template, because anomalous capacitance readings also indicate a broken channel. Only sensors with less than 5% non-functioning channels were considered acceptable.

Modules are built comprising one or more sensors. The sensors are physically supported by circuit-boards called hybrids, which also carry the read-out chips and associated circuitry for supplying power, etc... Photographs of the PHOBOS silicon detector modules are shown in Figure 2-4. The modules are then mounted on their frames and installed in the PHOBOS experiment.

2.3.3 Silicon Sensor Read-out System

The silicon read-out chips are commercially available and designed specifically for this purpose. Each chip has 64 or 128 channels and a large dynamic range for integrating the charge collected from the silicon pads. Sampling of the silicon signals is timed carefully relative to the event trigger in order to capture the peak of the signal and hence maximise the signal-to-noise ratio.

Front-End Controllers (FECs) manage the read-out chips. A fixed time after receiving the L1 trigger signal from the Data Acquisition (see Section 2.10), the silicon signals are sampled and the values held until the L2 trigger is received, at

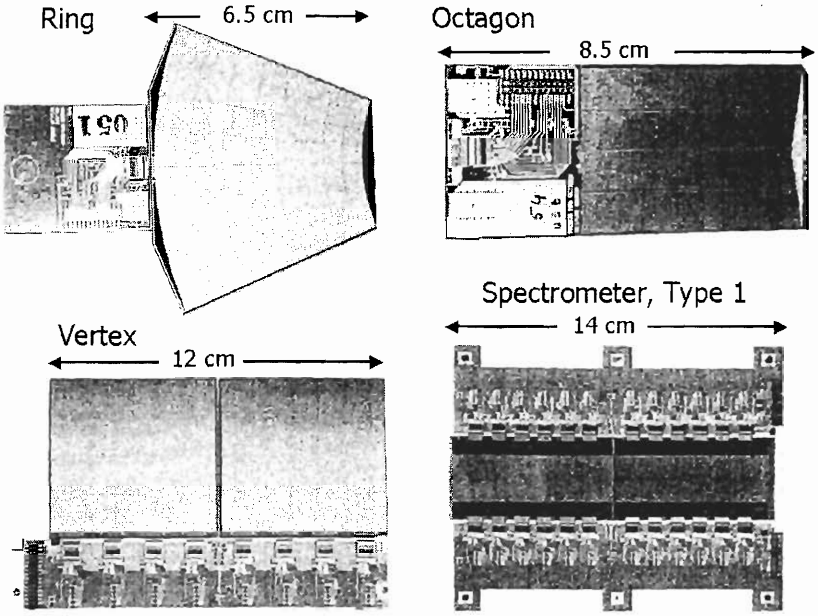

Figure 2-4: Photographs of four different types of silicon sensor modules used in the PHOBOS detector.

which point read-out of the data begins. If L2 is not received within 10 microseconds of L1, the FECs are cleared to be ready for subsequent events. Each FEC has four ports, with each port responsible for reading out a single detector module. The FECs also record information such as chip-bias and hybrid temperature, and this digitised information is appended to the event.

Data from each FEC is sent via G-link interface to a Data Multiplexing Unit (DMU) in the Data Concentrator. Data from all the DMUs are collated into serial streams and transmitted over optical fibres to the Data Acquisition (DAQ) in the PHOBOS Counting House. The Data Concentrator handles the interface to the DAQ: it distributes the L1 and L2 signals to the FECs.

The silicon read-out system is summarised in Figure 2-5. More details can be found in [85].

A special feature of the silicon front-end electronics is that it must be able to cope with 'latchups', where ionizing radiation deposits a large amount of energy in a read-out chip, causing the transistors to switch on and conduct an abnormally large amount of current which disables the chip. This can happen often in the environment

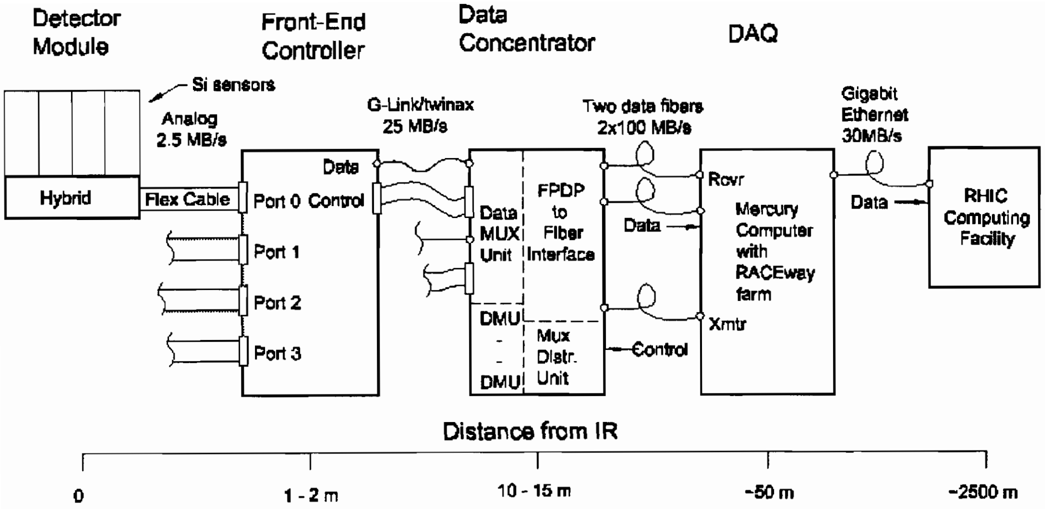

Figure 2-5: Illustration of read-out chain for PHOBOS Silicon sensors. The DAQ stages are described in Section 2.10.

of high-luminosity, high-energy heavy-ion collisions. The latched-up chip needs to be power-cycled before returning to normal operation. A latchup protection circuit was designed to detect and alert for these occurrences - a global power reset can then be issued by the operations-monitoring personnel in the Counting House.

2.4 PHOBOS Multiplicity Array

The PHOBOS Multiplicity Array is a set of single-layer silicon detectors used to measure the number and angular distributions of charged particles. It consists of a central barrel detector called the Octagon, which covers a pseudorapidity range $|\eta| < 3.2$, and two sets of three Ring counters mounted around the beam-pipe which extend the coverage up to $|\eta| < 5.4$.

2.4.1 Octagon

The octagon is 1.10 m long with a face-to-face diameter of 9 cm. It is supported by an aluminum frame (Figure 2-6). Certain sections of the Octagon have not been instrumented, in order not to shadow the higher precision Spectrometer and Vertex detectors. Each Octagon sensor is 84 mm long and 36 mm wide, with the active area divided into 4 rows of 30 pads; each pad measures 2.71 × 8.71 mm for a total active area of 34.9×81.3 mm^2.

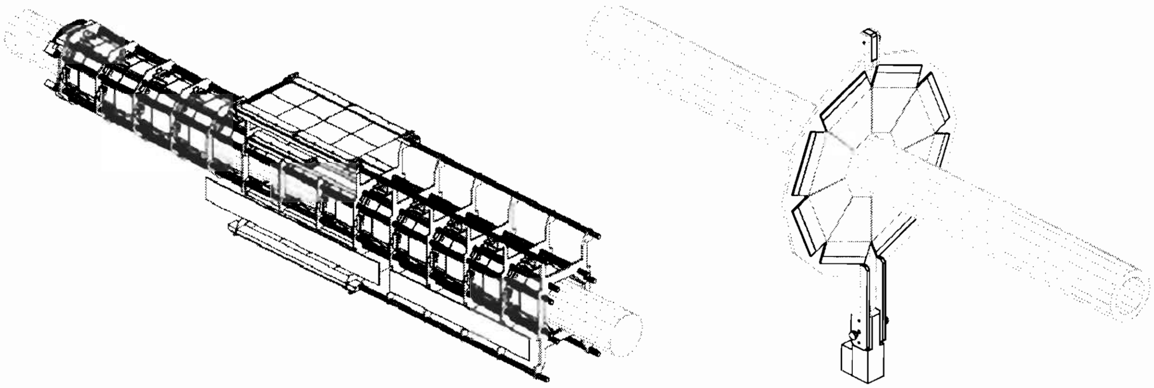

Figure 2-6: Left: The Octagon detector. Right: A Ring detector.

2.4.2 Rings

Mounted on carbon-fibre frames, the Ring detectors (Figure 2-6) are located at $z = \pm 1.13$, ± 2.35 and ± 5.05 m. The Ring inner radius is 5 cm, the outer radius 11 cm. Each Ring has 8 trapezoidal sensors, consisting of 64 pads arranged into 8 rows of 8 radial columns. The pad size is chosen so that each pad has approximately the same pseudorapidity coverage of $\Delta\eta \approx 0.1$ and ranges from 3.8 × 5.1 mm up to 10.2 × 10.2 mm.

2.5 Vertex Detector

The PHOBOS vertex detector consists of two pairs of silicon layers, located above and below the beam-pipe in the y direction. The design goal for this detector was to determine the location of the event vertex with a resolution of 0.2 mm for events which occurred within $|z| \leq 10$ cm.

The 'Inner Vertex' layers are 25 cm and 4.8 cm wide, and are located at $y = \pm 5.6$ cm. They consist of two modules of two sensors. The silicon pad dimensions are 0.473 mm × 12 mm, in a 4 × 128 arrangement.

The 'Outer Vertex' layers are also 25 cm long but are 9.6 cm wide and are located at $y = \pm 11.8$ cm. This time, each layer is comprised of four modules, each with two sensors, and the pads are 0.473 mm × 24 mm, arranged into two coulmns of 128. The silicon pad sizes were chosen to give good pointing resolution along z and to have low occupancy even in the highest multiplicity events.

For collisions at the nominal vertex $z = 0$, the pseudorapidity coverage of the vertex detector is $|\eta| < 1.54$ for the Inner Vertex and $|\eta| < 0.92$ for the Outer Vertex, while the azimuthal angle subtended by the detector is $42.7°$.

Details of the vertexing procedures and resolution achieved can be found in Section 3.5.

2.6 Particle Identification in PHOBOS

A multi-layer silicon spectrometer in a 2 T magnetic field is used to reconstruct the tracks of charged particles and determine their momenta. For low momentum particles, the spectrometer can also perform species identification. The Time-of-Flight walls, in conjunction with the Time-Zero detectors, provide additional particle identification capabilities at higher momentum.

2.6.1 PHOBOS Magnet

The PHOBOS magnet is a double dipole design, with vertical fields of opposite polarity on each side of the beam-pipe. The mechanical construction of this conventional room-temperature magnet is shown in Figure 2-7. The gap between the poles is 158 mm, and the magnet is fully energised at a current of 3600 A. It has a total bending power of ≈ 1.5 Tm. As illustrated in Figure 2-8, the magnet is designed such that the first six Spectrometer layers sit in almost zero field, while the remaining layers are in a region of approximately constant field. The maximum value is $B_y = 2.18$ T and the B_x and B_z components are less than 0.05 T in this region. The field strength has been measured at many points with a Hall probe and combined with software simulations of the design to produce a complete field map, from the beam-pipe through

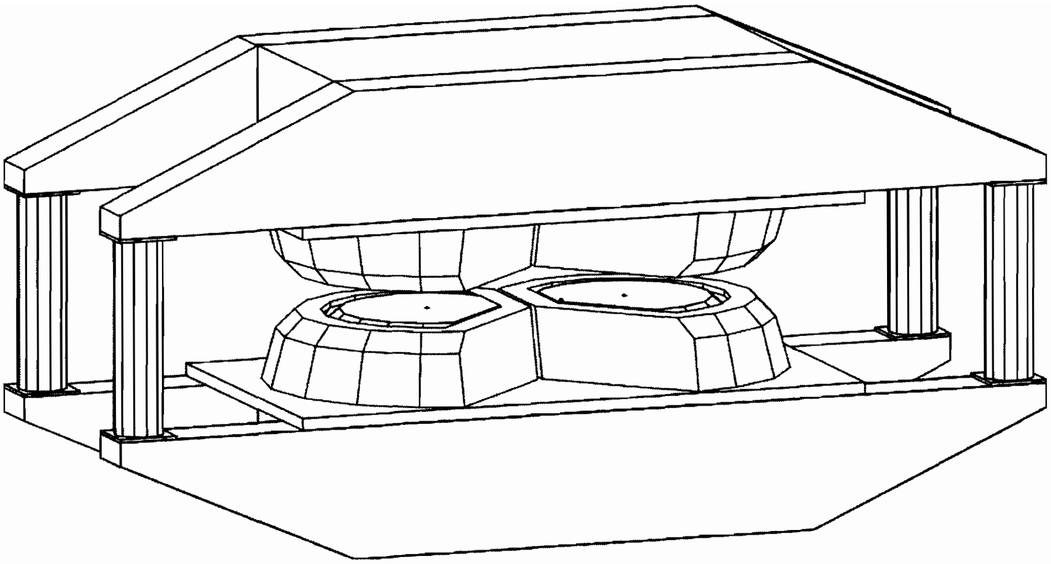

Figure 2-7: The PHOBOS magnet. The beam-pipe (not shown) goes through the centre of the structure.

the Spectrometer region all the way out to the Time-of-Flight walls.

2.6.2 Spectrometer

Sensor Type	Layers	Active Area (mm^2)	Number of Pads	Pad Size(mm^2)
1	0-3	70.0×22.0	70×22	1.0×1.0
2	4-7	42.7×30.0	100×5	0.427×6.0
3	8-15 inner	42.7×60.0	64×8	0.667×7.5
4	8-11	42.7×60.0	64×4	0.667×15.0
5	12-15	42.7×60.0	64×4	0.667×19.0

Table 2.1: Table of Spectrometer sensors.

The PHOBOS Spectrometer (Figure 2-9) is constructed from two arms, each containing 137 silicon sensors arranged into sixteen layers. The sensors are mounted on water-cooled aluminum frames, connected to a non-conducting carbon-epoxy carrier plate which minimises vibrations due to eddy currents when the magnet is being ramped. The Spectrometer enclosure is light- and air-tight, and the relative humidity is maintained below 10% by a dry nitrogen system.

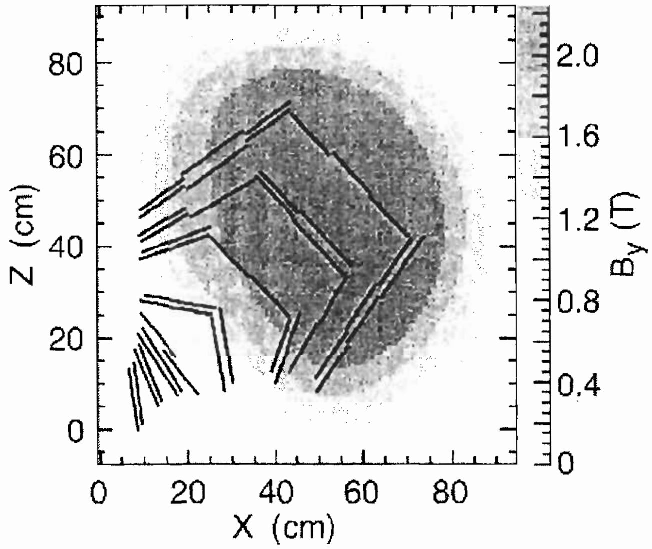

Figure 2-8: Magnetic field map showing the strength of the B_y component as a function of position in the xz-plane at $y = 0$. The position of the PHOBOS spectrometer layers are overlaid.

The dimensions of the Spectrometer in the $x - z$ plane are shown in Fig 2-10. Each arm covers an azimuthal angle of $\Delta\phi \approx 0.1$ radians.

The Spectrometer is designed to have a high degree of pixelisation for precise position determination in the $x - z$ plane. Five types of silicon sensors are used in the construction; their details are given in Table 2.1.

The trajectories of charged particles in the Spectrometer can be reconstructed based on the hit positions in each layer traversed. The algorithms used to do this are discussed in Chapter 4. The momentum resolution achieved is illustrated in Figure 4-7.

As well as simply determining the hit position, the energy deposited by charged particles in the silicon pixels can be used for species identification of low-momentum particles ($p < 1$ GeV/c). This procedure is also described in Chapter 4.

2.6.3 Time-of-Flight Walls

The purpose of the Time-of-Flight (TOF) detectors is to extend PHOBOS's particle identification capabilities to high transverse momentum. There are two TOF walls:

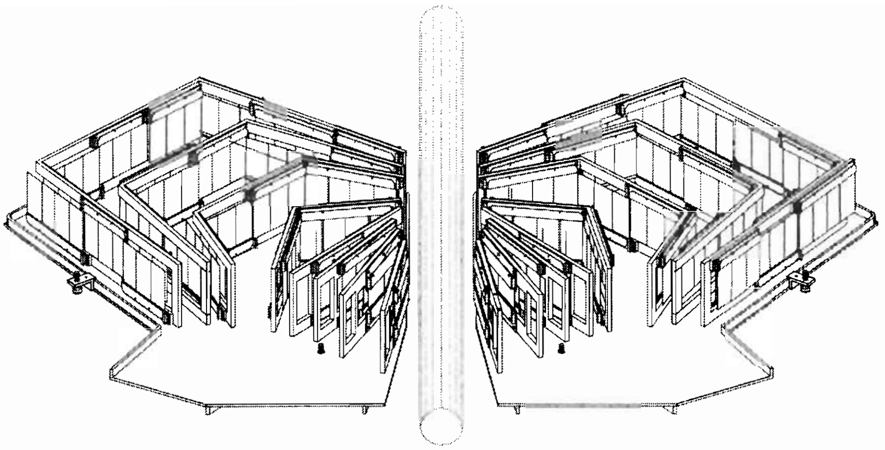

Figure 2-9: The PHOBOS multi-layer silicon Spectrometer

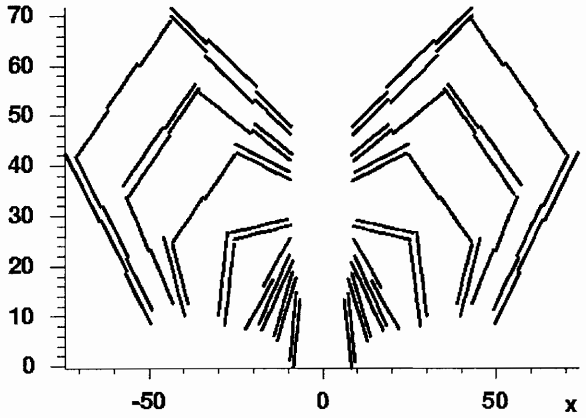

Figure 2-10: Dimensions of the Spectrometer in the $x - z$ plane.

wall 'B' is at 45^o to the beam-axis, 5.4 m from the origin; and wall 'C' is parallel to the beam-axis, at a distance of 3.9 m. Each wall is 1.1 m long and together they span a pseudorapidity range of $0 \leq \eta \leq 1.24$.

There are 120 Bicron BC404 plastic scintillators per wall. This material provides good timing resolution (by virtue of its 1.8 ns decay constant) and it has a moderate attentuation length. Also, its wavelength of maximum emission is 408 nm, which is very close to the peak response of fast photomultiplier tubes (420 nm). The scintillators are 20 cm in length with a cross-section of 8×8 mm.

The scintillators are connected to Hamamatsu R5900 photomultiplier tubes (PMTs).

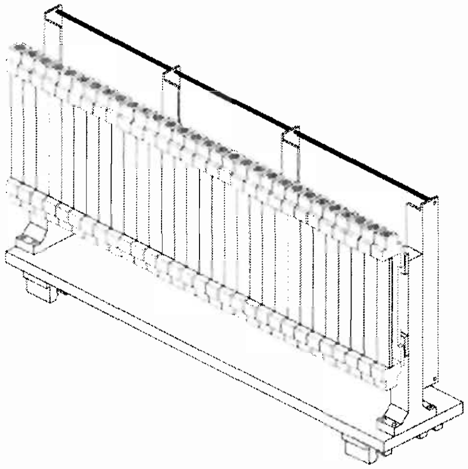

Figure 2-11: Time-of-Flight wall.

These PMTs have a segmented 2×2 anode and four scintillators are connected to one PMT, sharing the same photocathode. Channels are not isolated so there is the possibility for cross-talk, because electrons generated by photons from one scintillator can drift inside the tube and be collected by the anode for a different scintillator. Bench tests found this cross-talk to be less than 10% for adjacent channels and negligible for diagonal channels.

This PMT design also requires an unusual light-guide and mounting fixture for the TOF sensors. Cross-talk in the acrylic light-guides is minimised by leaving a 1 mm spacing between them and wrapping them in aluminum foil.

The signal from each PMT is split in two: one output is digitised by FASTBUS 13-bit ADC (LeCroy 1881M) to give a pulse-height reading, while the other is discriminated by CAMAC leading-edge discriminator (LeCroy 3412) then digitised by FASTBUS TDC (LeCroy 1875A) to give timing information with 25 ps/channel sensitivity. Each scintillator is read-out from the top and from the bottom. This has the advantage that the vertical position of the hit can be obtained using either the time differences (with 10 mm resolution) or the ratio of the pulse-heights (37 mm resolution). Furthermore, the consistency of pulse-height ratios and time differences from top and bottom readouts can be used to eliminate instances of 2 hits in the same scintillator, which would mess up the timing information.

The procedure for calibrating the Time-of-Flight detectors is described in Section 3.3. Based on bench tests, the TOF sensors are found to have an intrinsic timing resolution of 80 ps.

2.6.4 Time-Zero Detectors

The Time-Zero (T0) detectors are Cerenkov detectors, used to provide a precise measurement of the start-time for the TOF walls. They are constructed from Bicron BC800 Cerenkov radiator, shaped as cylinders with radius 2.5 cm and length 3 cm and connected to Hamamatsu R2083 fast PMTs, also 50 mm in diameter. Ten sensors are arranged around a circle of radius 10 cm from the beam-pipe. The T0s are located at $z = \pm 5.2$ m, but during d+Au collisions, the T0 on the positive (outgoing deuteron) side was moved to a position of $z = 2.5$ m.

The T0s are calibrated in a similar way to the TOF (see Section 3.3); their intrinsic time resolution has been measured to be 110 ps.

T0 Time Equaliser

The T0 signals can provide an online measurement of the event vertex position, to be used as a vertex trigger. It is therefore critical that all T0 counters provide a uniform timing response, so that the vertex trigger is not distorted depending on which counters fire.

But inevitably, small differences in PMT timings, cable lengths, etc... can cause variations that would have a significant effect on an online trigger. It is impractical to make high-precision corrections for such variations using delay cables or by fine-tuning the sensor positions. Instead, a custom-designed Time Equaliser (TEQ) is used to align the output signals of all counters in real-time.

There are two TEQs, one for each set of T0s. The boards slot into a CAMAC crate in the PHOBOS interaction region. They take the individual T0 counter signals as inputs and produce an output signal that is used by the vertex trigger. This output is generated by the first input received from any T0 channel - but the time between firing and emitting the signal is adjustable separately for each channel, so the output can be programmed to be emitted uniformly, no matter which channel fired it. The necessary delays for each channel are determined from calibration data and are uploaded to the boards via the CAMAC interface. The TEQ can operate at rates up to the MHz level, comfortably above the PHOBOS event rate. The T0 vertex resolution achieved after fine-tuning the TEQ settings is roughly 3-4 cm.

2.7 Calorimeters

A forward proton calorimeter (PCAL) is located adjacent to the Zero Degree Calorimeters (described in Section 2.8.2). The PCAL is composed of modules of lead-scintillator sandwich, arranged into 12 rows and 8 columns; signals are recorded using VME ADCs. The PCAL is designed to study spectator and recoil protons from the heavy-ion collision.

Using the same type of modules as the PCAL, a new calorimeter called the SpecCAL was installed for the 2004 run. The goal of the SpecCAL is to provide calorimetry for high-p_T particle tracks reconsructed in the Spectrometer.

2.8 PHOBOS Trigger Detectors

The main trigger detectors for the PHOBOS experiment are the Paddle counters, the Zero Degree Calorimeters and the Spectrometer Trigger. There are also Cerenkov detectors [86], two sets of 16 Lucite radiators located around the beam-pipe at ±5.5 m, but these are no longer used, having been supplanted as a vertex trigger by the T0s.

2.8.1 Paddle Trigger Counters

The primary event trigger for the PHOBOS detector in Au+Au running is provided by the 'Paddle Counters' [87]. These are two sets of sixteen scintillator counters located at ± 3.21 m along the beam-axis from the nominal interaction point. They cover the pseudo-rapidity region $3 < |\eta| < 4.5$, with the active area comprising 99% of the subtended solid angle. A diagram of one of these detectors is shown in Figure 2-12.

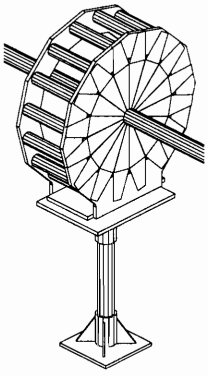

Figure 2-12: A Paddle detector and a section of the PHOBOS beam-pipe.

Each module consists of a tapered BC-400 plastic scintillator joined to a BC-800 light-guide, which couples to a hybrid photomultiplier tube assembly (Hamamatsu H1151-2). The individual scintillators are 18.6 cm in length, 0.95 cm thick, with a 9.5 cm long outer edge and 1.9 cm inner edge. The photomultiplier tube is surrounded by a 2 mm mu-metal magnetic shield. When mounted on their frame, the inner tip of the scintillator is 2 cm from the PHOBOS beam-pipe.

The pulse-height signals from the photomultiplier tube are proportional to the energy deposited in the scintillator; the value for a Minimum Ionizing Particle (MIP) is $\Delta E = 1.7$ MeV. The pulse-heights are measured by an analogue-to-digital convertor (ADC) and were calibrated using bench-tests with a radioactive beta-source.

The time resolution of the paddle detectors is $\sigma_T \approx 1$ ns, which is sufficient to distinguish between collision and background events at the PHOBOS interaction point.

2.8.2 Zero-Degree Calorimeters

Collisions of heavy-ions cause the evaporation of neutrons. At RHIC collision energies, these fragmentation neutrons diverge from the beam axis by less than 2 mrad and are the dominant source of particles in this 'zero degree' region. The RHIC Zero-

Degree Calorimeters [88] (ZDCs) are designed to detect and measure the energy of these spectator neutrons. The detectors are identical in all four RHIC experiments, to provide a common event characterisation. They are also useful as a luminosity monitor for RHIC beam operations.

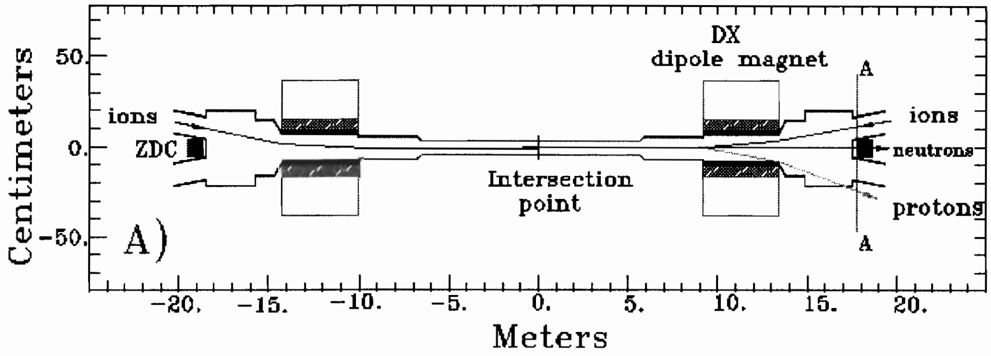

Figure 2-13: Location of the Zero Degree Calorimeters within a RHIC interaction region.

As shown in Figure 2-13, the ZDCs are positioned at ±18 m at each of the four RHIC heavy-ion experimental interaction regions, in the 'elbow' behind the 'DX' dipole magnets, which direct the ion beams into their individual circulating rings. The DX magnets sweep charged particles out of the ZDC acceptance.

The ZDCs are hadron shower calorimeters. Their location constrains the detectors to be no more than 10 cm wide, and they were therefore designed to maximise the transverse shower containment within this radius. Tungsten alloy absorber is used, and the shower is sampled by Cerenkov radiation from charged shower products in commercial Poly-methyl-methacrylate (PMMA) optical fibre.

Each Zero Degree Calorimeter is composed of three modules. The length of each module is equivalent to two nuclear absorption lengths or fifty radiation lengths ($2\lambda_I$, $50X_0$). The mechanical construction of a ZDC module is shown in Figure 2-14. Tungsten plates sandwich single-layer ribbons of optical fibres (diameter 0.5 mm), which are oriented at 45° relative to the incident beam direction because this approximately coincides with the Cerenkov angle of $\beta = 1$ particles in PMMA.

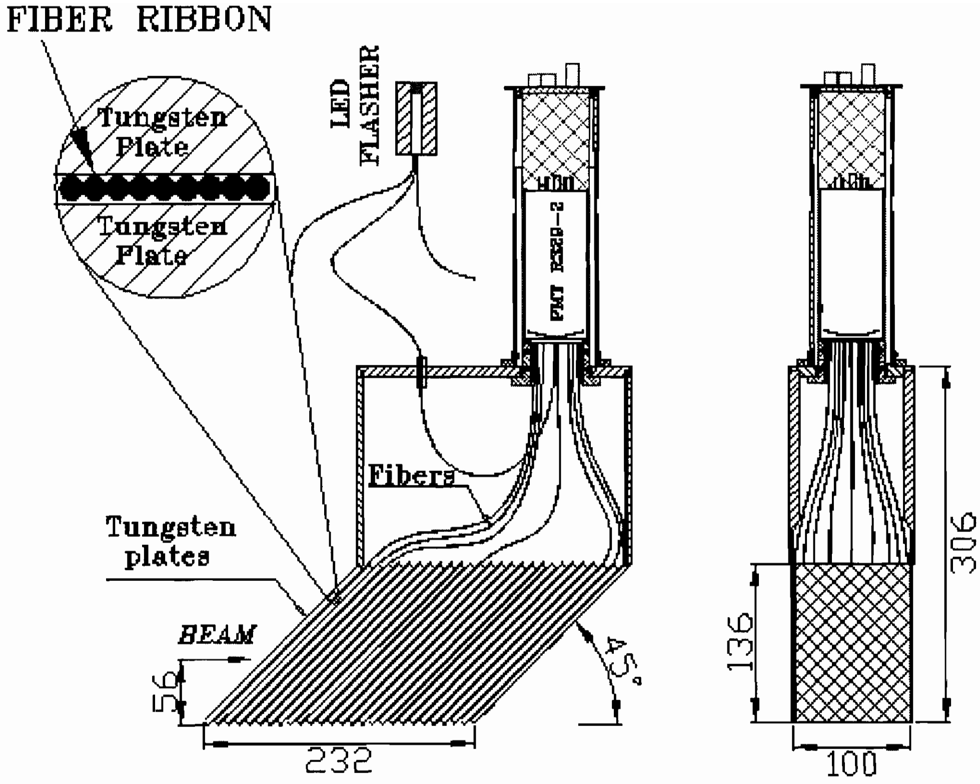

Figure 2-14: Mechanical design of Zero Degree Calorimeter modules.

2.8.3 Spectrometer Trigger

The purpose of the Spectrometer Trigger ('SpecTrig') is to enhance the yield of high transverse momentum particles for particle identification in the Time-of-Flight. It is designed for use in the low-multiplicity environments of d+Au and p+p collisions. The operating principle of the SpecTrig is that hit positions in the segmented scintillator detectors, in combination with the known event vertex, can be used to identify and trigger on the straight-line trajectories that correspond to high-momentum particles traversing the Spectrometer and hitting the Time-Of-Flight detector.

The SpecTrig consists of two walls, labelled 'B' and 'C'. Each wall comprises 10 scintillators and is 73 cm long. The location of the SpecTrig walls are chosen so that they cover the solid angle subtended by the Time-of-Flight detector. Wall C is at 90^o, 1.6 m from the origin, and wall B is at 45^o and a distance of 2 m.

The SpecTrig sensors are made from BC-408 plastic scintillator and are 11cm long, 7.24 cm wide and 0.5cm thick. They connect to a trapezoidal light-guide which tapers down to a width of 4.4cm over a distance of 15cm; optical cement glues this to a PMT. The PMT is encased in a double-layer magnetic shield. The scintillators are wrapped in foil and the whole arrangement wrapped in black tape to eliminate light leaks that would contribute to noise. The sensors are mounted on an aluminum frame which bolts to the PHOBOS magnet structure - the design therefore had to account for compression of the structure when the magnet is energised.

Stability of the detector is checked using a monitoring unit mounted for each wall. Light from a blue LED is transmitted directly into each PMT via fibre-optic cables. This allows constant monitoring of the tubes, checking for malfunctions, drifts in gain, etc.. The LED is driven by an electronic pulser circuit. To ensure no 'false alarms' caused by an LED failure, a photo-diode is also installed to monitor the LED light output itself.

The pulse height of the SpecTrig signals are recorded by VME ADCs and the timing measured by VME TDCs, after the signal passes through a discriminator. The discriminated signal is also used as part of the trigger, indicating whether each sensor has recorded a 'hit'.

For trigger purposes, the 120 sensors in each TOF wall (which, as described in Section 2.6.3, have PMTs attached to both top and bottom) are grouped into 14 sets of 8 sensors each. These are connected to a discriminator and the SUM output is used, so that a 'hit' is recorded for this group if any one (or more) of the individual channels are above threshold.

Low momentum particles will bend in the 2 T magnetic field region, while high-p_T ($p_T \geq 2$ GeV/c) particle trajectories remain effectively straight. For both pairs of SpecTrig and TOF walls, the hit combinations which point back to a vertex position in a designated range ($|v_z| \leq 20$ cm for wall B and $0 \leq v_z \leq 20$ cm for wall C) are calculated from the geometry. If the event vertex is known independently to be in this fiducial range, then these hit combinations should correspond to high-p_T particles. It is the selection of these desired SpecTrig+TOF hit combinations that forms the basis for the high-p_T Spectrometer trigger. Thus it is clear why the SpecTrig is only effective in low multiplicity collision environments - the detectors need to have low occupancy, or the system will be swamped by false triggers arising from just combinatorics.

The online trigger decision is made using an XLM-72 programmable electronic logic module, manufactured by JTEC Instruments of Rochester, NY. The XLM-72 has a custom VME control interface and an 80MHz internal clock. It has 72 ECL ports that can be configured by the user as either input or outputs: for this design, 68 ports are inputs, for the 2×10 SpecTrig and 2×14 TOF signals, each carried on a separate 17-channel ribbon cable; and 4 are outputs, for the trigger decision plus diagnostic outputs. The main feature of the XLM-72 is a Xilinx XCS40XL Field Programmable Gate Array (FPGA). This can be configured (or, indeed, reconfigured, since connections in an FPGA are not permanent and can always be remade) to any design containing up to 40,000 logic gates.

Since the SpecTrig and TOF signals arrive simultaneously at the XLM, the operation is quite straight-forward and simply consists of a series of AND gates for each of the desired SpecTrig+TOF hit combinations - if at least one of these is satisfied, the XLM outputs a trigger signal, known as the SILVER decision. Diagnostic outputs include indicators of whether the SILVER decision was due to the B walls or the C walls or both. This design is programmed by the user in VHDL, a hardware description language for very high-speed integrated circuits, and loaded onto the FPGA using the JTAG connection protocol.

The event vertex is independently measured online by the Time-Zero Cerenkov detectors (see Section 2.6.4) and a vertex trigger identifies the events that fall in the fiducial vertex region. The complete high-p_T Spectrometer trigger, then, requires a coincidence between the SILVER and Vertex triggers.

In the d+Au run of 2003, the Spectrometer Trigger produced an enhancement by a factor of twenty in the number of high-p_T Time-of-Flight tracks per recorded event.

2.9 Event Trigger

The main PHOBOS event triggers are a Minimum Bias collision trigger, a vertex trigger and the high-p_T Spectrometer trigger. There are also special non-collision events recorded for diagnostic and calibration purposes.

The Spectrometer trigger has already been fully described in Section 2.8.3. The Minimum Bias collision trigger uses the Paddle counters, which have a relatively large acceptance and adequate timing resolution. A coincidence is required between signals from both sets of Paddles on either side. A further constraint is imposed that at least 2 out of the 16 Paddle segments on each side are hit: this restriction makes for better online selection of true collision events, but necessarily reduces a little the efficiency for triggering on very peripheral collisions.

The vertex trigger is designed to enhance the sample of useful events for physics analysis by triggering on collisions which occur in the fiducial vertex range $-20 \leq v_z \leq 20$ cm. The T0N and T0P Time Equalizer outputs (see Section 2.6.4) are connected to a Time-to-Analogue Convertor (TAC): T0N acts as the start and T0P as the stop signal, with an artificial delay added to ensure that Start always arrives before Stop. If the time difference $\Delta T = T0P - T0N$ falls within an adjustable window, the TAC emits a fast pulse which is the signal for the vertex trigger. The window size therefore determines the accepted vertex range, and is typically made a little wider than 40 cm, to make sure events at the edges of the fiducial vertex range are not missed; the offset can be shifted to make this range asymmetric about the origin, if desired.

The PHOBOS trigger logic has three levels:

- L0 - starts all ADC and TDC gates;
- L1 - controls sampling of Si signals;
- L2 - controls read-out of Si channels by DAQ;

The timing of these trigger signals relative to the event occurrence is very important, in order to properly capture all the detector responses. L0 triggers are required to be in coincidence with the RHIC bunch-crossing clock (see Section 2.1). L2 provides an option for high-level, complex trigger decisions - it can arrive up to 10 μs after L1; if it does not arrive within this time, the Si channels are not read-out and are cleared in preparation for the next event. In practice, however, all current PHOBOS trigger decisions are completed in time for L1, and there is no trigger selection performed at the L2 level.

Custom-designed electronic logic boards are used to handle the L0 and L1 trigger processes. Since the different types of trigger occur at quite different rates, these boards have the ability to prescale their inputs to obtain the desired trigger mix. For example, in regular Au+Au running, the bulk of the events will be vertex triggers, with a few percent Minimum Bias triggers, for balance.

It is possible for the same event to be fired by more than one trigger (a simple example is an event which satisfies both the vertex and collision triggers). The L0 and L1 boards record a bit pattern indicating all the triggers which fired for each event and this information is stored in the event data.

2.10 Data Acquisition

Once an event has been been accepted by the trigger, the signals in all PHOBOS subdetectors are read-out and recorded by the data acquisiton (DAQ) system [89]. The DAQ has to be able to combine data from 135,000 Silicon channels, 2000 FASTBUS channels (for the scintillator detectors - Paddles, TOF, T0s, ZDCs) and 300 VME channels (for PCAL, SpecCAL and SpecTrig detectors).

Processing of the Si data received from the FECs (see Section 2.3.3) is performed by 24 PowerPCs, organised into a master-workers scheme. These run the Mercury operating system, which is designed for fast real-time parallel-processing, and are linked by a RACEway network switch that allows a high rate of data transfer and the ability to make many concurrent connections. The master distributes the incoming data among the workers, with each worker processing the output of 2 FECs. Workers check the received data for consistency and compress it using Huffman coding.

Huffman coding (see [90] for a fuller description) is a lossless compression algorithm for fixed-length input, producing variable-length output. It takes advantage of a situation where some inputs are more common than others, and can therefore be encoded in fewer bits than the more rare cases. This is ideal for the PHOBOS silicon data, since the Spectrometer and Vertex detectors were designed to have low occupancy: most channels are therefore 'empty' and contain only pedestal data, which can be coded in less bits than true signals. The compression achieved is almost a factor of four.

Workers put their checked and compressed data into the output queue. An 8-bit global counter is used to ensure correct synchronisation of all data from the same event. An Event Builder combines the output Si data with the significantly smaller (in terms of number of channels) digitised data from the FASTBUS and VME subsystems; the same synchronisation marker is again used. The whole operation is under the control of an Event Manager, which responds to the event trigger signals and distributes the synchronisation marker among all subsystems.

In older versions of the PHOBOS DAQ, complete events were written to a disk array controlled by SUN workstations - a major upgrade was to replace this with a dedicated data recorder, to allow for faster event through-put. An event distributor samples events from the DAQ data-stream and makes them available for online monitoring in the Counting House. The final stage of the DAQ process is transfer to the data tape storage system, via Gigabit ethernet. The unit of transfer is called a 'sequence' and is 1 GB in size, consisting of typically about 10,000 events.

The PHOBOS DAQ is capable of a sustained event rate of around 400-500 Hz. The main bottleneck in the system is the processing of the Si data, which takes 1-2 ms.

Dead-time, Applied Trigger Rate and Recorded Rate

The rate at which events are recorded is related to the rate of applied triggers (a) and the 'dead-time' associated with processing each event (d). If the system is running at rate r, then the fraction of time during which it is 'live' and able to accept triggers is $1 - rd$, so assuming the applied triggers are uniform, the rate of accepted triggers fulfills the self-consistency relation:

$$r = a(1 - rd) \tag{2.2}$$

If we define the maximum theoretical rate m as the inverse of the dead-time, $m = 1/d$, then the above equation can be re-arranged into the convenient form:

$$\frac{1}{r} = \frac{1}{a} + \frac{1}{m} \tag{2.3}$$

2.11 PHOBOS Computing

The raw data recorded by the PHOBOS experiment is stored on tape drives in the HPSS storage system. A web-based file catalog provides convenient and efficient access to this data. An Oracle Database coordinates almost all aspects of the experimental operations.

The RHIC Computing Facility [91] (RCF), consisting of Linux Farms containing hundreds of computers, is used for data processing, detector simulation and physics analysis.

PHOBOS software is based on the ROOT [92] object-oriented C++ framework. ROOT aims to provide a common framework for all aspects of high-energy physics computing, from detector simulation to data analysis, and it implements basic functionality such as I/O, as well as a convenient interpreted environment for programming. The PHOBOS Analysis Toolkit (PhAT) is built on top of ROOT and implements PHOBOS-specific classes and programs. Simulations of the detector use the popular GEANT package [93].

Chapter 3

Event Characterisation

A basic selection procedure to remove non-collision events is common to all PHOBOS physics analyses. The centrality of the collision is then determined, and data from the plastic and silicon sub-detectors must be fully calibrated. A variety of procedures are used to obtain the position of the event vertex.

3.1 Event Selection

Selection of collision events is primarily done by means of timing information from the Paddle detectors. A plot of the Paddle timing is shown in Figure 3-1. The central region is the collision events, where collision products have hit paddle segments on either side of the interaction region. The satellite peaks at either end are due to 'beam-gas' events, where a beam ion scatters off a gas molecule in the beam-pipe, and then the deflected beam ion passes through both sets of paddles. These two peaks are centred on ±21 ns, which is the time taken for a particle to travel from one set of paddles to the other (a distance of 6.4m) at the speed of light. A cut on the Paddle time difference being within ±4 ns is used to reject such beam-gas events and other outliers.

As well as this Paddle timing cut, collision events are required to have good ZDC timing also. Events which were triggered by the DAQ heartbeat or by the Time Calibrator (see Section 3.3) are manifestly not collision events, so these must be

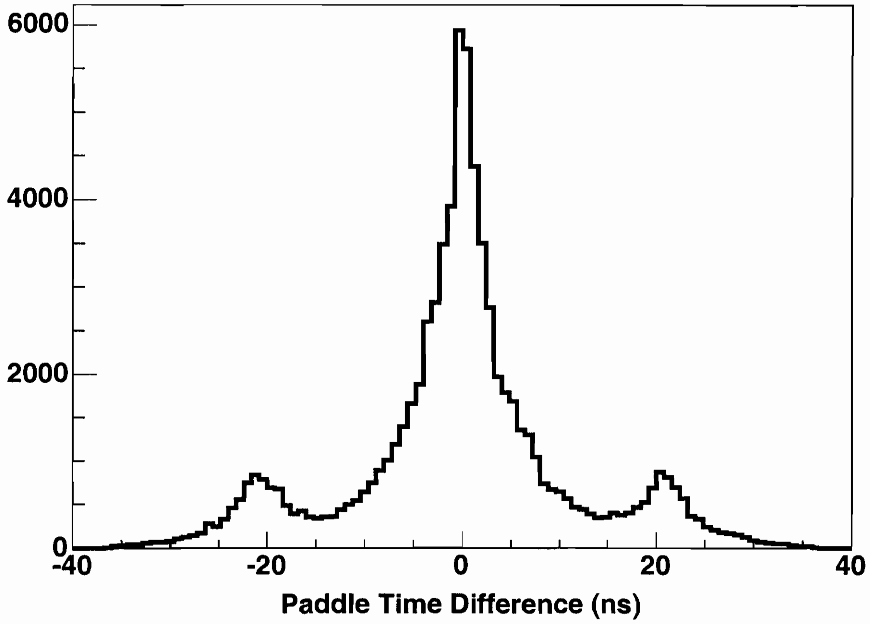

Figure 3-1: Paddle timing from a PR04 200GeV MinBias run. Events known not to be collisions, eg DAQ heartbeat events, are not included.

discarded. They are identified by the L0 and L1 trigger bit patterns which are stored in each event. As a further test of the integrity of the event, a check is done to ensure that data from all the major read-out subsytems - Silicon, FASTBUS and VME - is present.

3.2 Centrality Determination

As explained in the Introduction, it is important to know the impact parameter of a heavy-ion collision, but it is impossible to focus and steer colliding ion beams accurately enough to be able to control this. Instead, when a collision occurs, its centrality has to be determined after the fact by means of some measurable quantity. Ideally, this measurable quantity would have a one-to-one correspondence with a centrality measure such as the impact parameter or the number of participating nucleons (N_{part}); in practice, it is sufficient for them to have a strictly monotonic relationship and a tight correlation.

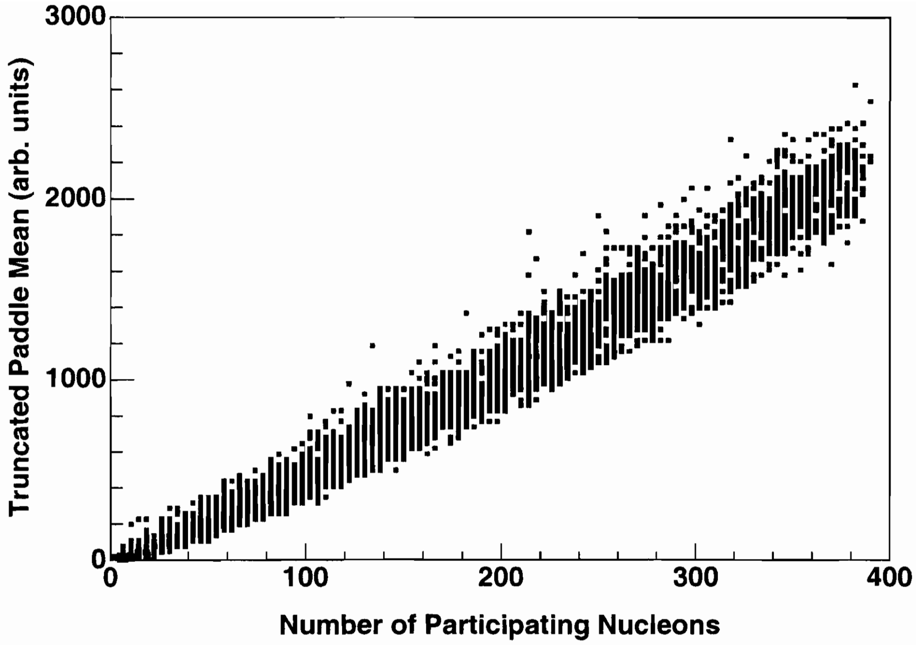

Figure 3-2: Relationship of Truncated Paddle Mean signal to number of participating nucleons in the heavy-ion collision, from a Monte Carlo simulation of the PHOBOS detector. This relationship forms the basis of the centrality determination process in PHOBOS.

Monte Carlo simulations of the PHOBOS detector have shown that the Truncated Paddle Mean signal satisfies these requirements, as can be seen from Figure 3-2 where it is plotted against N_{part}.

The truncated mean is a standard way of dealing with the fact that fluctuations in energy loss by charged particles are quite asymmetric, with a large tail extending to higher energy (this is described in more detail in Section 4.3.2). In each event, the paddle segments with a signal above the pedestal value are sorted in order of size, and the top 25% are discarded; the remainder are averaged. This truncated mean reduces the effect of energy loss fluctuations and improves the correlation of the Paddle signal with N_{part}.

As a cross-check on the use of the Paddle signal as a centrality measure, it is compared to the ZDC signal in Figure 3-3. As described in Section 2.8.2, the ZDCs measure spectators in a heavy-ion collision, not participants, so an anti-correlation is expected between ZDC and Paddle signals. This anti-correlation is indeed seen over most of the range. There is also a 'turn-over' observed at very peripheral events, where the ZDC signal is reduced because few neutrons are entering its limited acceptance.

Centrality classes are defined by partitioning the truncated Paddle mean distribution into bins; Monte Carlo simulations then estimate the mean impact parameter b and N_{part} for each bin. The centrality bins defined for 62.4 GeV Au+Au collisions are shown in Table 3.1.

An important input to this is the trigger efficiency, i.e. the percentage of the total inelastic collision cross-section sampled by the PHOBOS online trigger. This varies depending on the exact choice of trigger configurations, but is typically in the range 85-95%. The number is determined by detailed comparisons of trigger data to Monte Carlo simulations.

For d+Au collisions, or collisions at lower energies, other centrality measures based on signals in the Octagon and Ring silicon multiplicity detectors have been used.

Bin	% σ_{int}	TruncPdlMean	$\langle N_{part} \rangle$	N_{part} rms	$\langle b \rangle$ (fm)	b rms
0	0-20%	-5.2761 - 0.2401	4.236	2.79	14.596	1.49
1	20-25%	0.2401 - 11.432	9.148	4.097	13.553	0.952
2	25-30%	11.432 - 24.8081	13.699	4.84	13.023	0.767
3	30-35%	24.8081 - 42.2818	19.496	5.777	12.507	0.677
4	35-40%	42.2818 - 65.2173	27.004	6.757	12.015	0.629
5	40-45%	65.2173 - 92.7655	36.585	8.046	11.505	0.615
6	45-50%	92.7655 - 126.117	48.117	9.332	10.996	0.586
7	50-55%	126.117 - 164.889	62.078	10.729	10.442	0.572
8	55-60%	164.889 - 210.629	78.53	12.411	9.871	0.582
9	60-65%	210.629 - 260.863	97.851	14.274	9.276	0.586
10	65-70%	260.863 - 315.945	119.844	16.232	8.646	0.594
11	70-75%	315.945 - 378.342	145.228	18.318	7.966	0.601
12	75-80%	378.342 - 447.432	174.065	20.728	7.224	0.641
13	80-85%	447.432 - 524.313	208.807	23.728	6.357	0.686
14	85-90%	524.313 - 605.288	247.618	26.732	5.39	0.76
15	90-94%	605.288 - 674.527	287.95	28.163	4.325	0.854
16	94-97%	674.527 - 732.255	323.418	27.943	3.228	1.021
17	97-100%	732.255 - 99999	349.438	22.729	2.26	0.992

Table 3.1: Centrality classes used for analysis of 62.4 GeV Au+Au collisions.

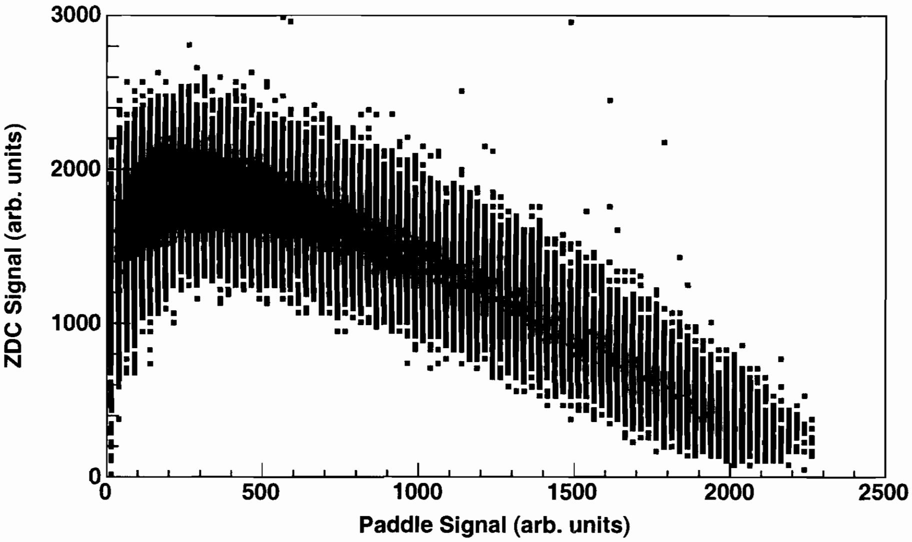

Figure 3-3: Correlation of ZDC and Paddle counter signals, for events with a valid vertex in the range $-20 < v_z < 20$ cm.

3.3 Time-of-Flight and Time-Zero Detector Calibrations

Analogue signals from the plastic scintillator Time-of-Flight and Cerenkov T0 sensors are digitised by FASTBUS ADCs and TDCs. ADCs integrate the input they receive over a fixed period of time, called the 'gate'. Random detector noise will therefore also be integrated and this leads to a non-zero ADC reading even in the absence of any true detector signal - this is called a 'pedestal' and should be subtracted. Pedestal values are obtained by regularly recording the ADC readings when there are no colliding beams in the RHIC rings. ADC signals after pedestal subtraction can be converted to actual energy deposited, based on bench tests with controlled radiation sources and Monte Carlo simulations of the TOF plastic scintillators. The MIP peak for these scintillators is found to be 1.45 MeV.

TDCs digitise the time between START and STOP signals. In PHOBOS, all TDCs receive START from the L0 event trigger (see Section 2.9). For each TOF and T0 sensor, the output of the PhotoMultiplier Tube is connected to a discriminator which produces a pulse if the input signal is greater than the defined 'hit' threshold; this discriminator output then acts as the STOP signal for the TDC for that channel.

Differences in cable lengths or PMT throughput times, for example, cause systematic differences in the timing responses of the TOF and T0 channels. A Time Calibrator device is used to account for this. The Time Calibrator produces an initial pulse which is routed to the trigger logic to provide L0 and start the TDCs. A second pulse generated at a fixed time relative to the first is directed to the TEST input of the aforementioned discriminators. A signal to the TEST input causes all discriminator channels to send an output pulse, which then stops each TDC channel. By doing this many times, the relative timing delays in each channel can be calibrated. The period between Time Calibrator initial and final pulses is also varied, so that the conversion from TDC units to real time can also be checked.

As well as intrinsic timing differences between channels, significant 'drifts' in timing have been observed over the course of the data-taking period. These have been attributed primarily to temperature-induced variations in the conductivity of the ribbon-cable carrying the signals, and also need to be accounted for. To this end, special-purpose Time Calibrator runs are recorded daily, and Time Calibrator events are also incorporated into the trigger mix for regular data-taking, comprising less than 1% of the total data.

A further correction which needs to be applied to the TOF and T0 recorded times is due to the 'slewing effect,' which arises because the PMT output signal is discriminated in order to generate the TDC STOP signal. All PMT pulses tend to peak at the same characteristic time; this means that a larger pulse will trigger the discriminator threshold sooner (as illustrated in Figure 3-4), producing a systematic timing shift. For a fixed discriminator threshold, the size of this slewing effect can be determined and an appropriate correction is applied based on the ADC reading for each channel.

A final point to note is that channel timing delay measurements using the Time Calibrator as described above only probe differences in timing paths from the discriminator onwards: differences due to scintillator or PMT response are not included

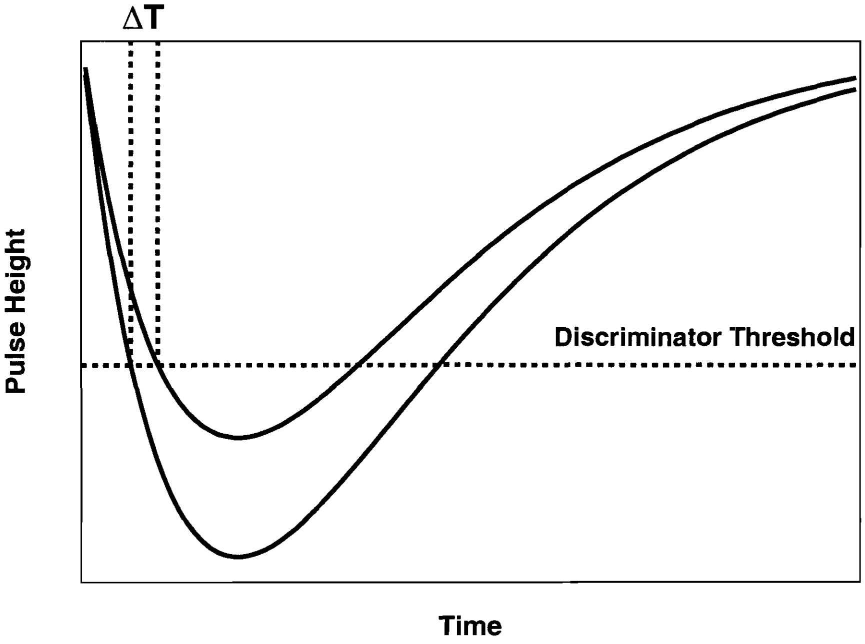

Figure 3-4: Illustration of slewing effect on TOF and T0 recorded times.

in this method, and so there is still room for some fine-tuning. This is performed after the TOF hits have been associated with the particle tracks in the Spectrometer; the physical fact that particle velocities should be very close to the speed of light can then be used to fine-tune the calibrations for each TOF channel.

3.4 Silicon Signal Processing

3.4.1 Pedestal, Noise and Gain Calibrations

The presence of leakage currents in silicon will generate a pedestal ADC reading even in the absence of any true signal. Pedestals are calculated from the regular silicon data in a two-pass system. First, the mean reading for each channel is calculated over a few hundred events - since the PHOBOS silicon sensors have typically low occupancy, this mean will be dominated by the pedestal. In the second pass, only readings which fall within a defined narrow range of the mean are averaged - this rejects signals and gives a better determination of the pedestal.

The voltage applied to the silicon read-out chip can fluctuate, which produces collective fluctuations in all channels controlled by that chip, called 'Common Mode Noise.' This is detected by special algorithms and subtracted, so that the corrected signal reading is defined by: signal = ADC - pedestal - Common Mode Noise.

The conversion from ADC units to actual energy deposited makes use of the internal calibration circuitry of the silicon read-out devices. During special calibration runs, a series of known signals are input to each chip and the ADC readings of each channel are recorded, producing a gain curve that is used to calibrate the real data. It is found that 1 ADC unit corresponds to 2.1 keV and that the most probable value of the energy deposition in these silicon sensors (called the Minimum Ionizing Particle or MIP peak) occurs at $\sim$80 keV.

After calibration, the silicon signals undergo zero-suppression, where pixels with very low energy signals (less than 10% of the MIP value) are simply negelected and set to zero. This saves considerable storage space, particularly since the silicon sensors have low occupancy.

3.4.2 Dead and Noisy Channels

Non-functioning channels in the silicon detectors need to be identified and accounted for. At the other extreme, channels which are very noisy will tend to produce spurious signals and should also be removed from consideration.

A map of dead and noisy channels is made based on the number of times each channel produces signals above a defined threshold, and on the energy associated with each such signal; this map is then used to mask these channels out before doing physics analysis of the data.

3.5 Vertexing

Heavy-ion bunches in the RHIC machine have a Gaussian longitudinal profile with $\sigma_L \approx$25 cm - thus there is a large range along the z-axis where collisions can occur. It is important to have a good measurement of the event vertex position in order to properly interpret the detector response.

3.5.1 Performance of the Vertex Detector

As is to be expected from its name, the Vertex detector (see Section 2.5) gives the most precise measurement of the event vertex, for collisions which occur within 20 cm of the nominal bunch crossing-point. Hits in the Inner and Outer Vertex layers are assembled into two-hit tracks, and then all tracks are extrapolated to the z-axis to estimate the most likely position of the event vertex. The geometry of the vertex detector locations and pad sizes leads to a vertex resolution of $\approx 450/\sqrt{N-1}\ \mu$m, where N is the number of reconstructed Vertex tracks in the event.

The vertex resolution is also limited by multiple scattering of particles in the beam-pipe and Inner Vertex layers. Studies have estimated this effect to be $\approx 85/\sqrt{N-1}\ \mu$m for 0.5 GeV/c pions.

A further effect arises from our lack of perfectly precise knowledge of the geometrical positions of the sensors. The extent of this 'misalignment' can be estimated by treating the top and bottom vertex detectors separately and finding an independent measure of the event vertex from each. Systematic shifts in the found vertex positions indicate relative misalignments of the sensors. This effect is found to dominate the vertex resolution for high-multiplicity events. Studies involving cosmic ray tracks which traverse all layers are used to fine-tune our knowledge of the sensor alignment after the initial survey measurements.

An important aspect of a good vertexing procedure is the ability not only to find a vertex but also to estimate just how reliable this vertex really is. All PHOBOS vertex-finders report a status flag to indicate the expected quality of the found vertex.

3.5.2 Vertices from other Sub-Detectors and the Composite Vertex

The PHOBOS Spectrometer, Octagon and T0 detectors are also capable of obtaining an estimate of the event vertex.

Straight-line tracks are reconstructed using the Spectrometer layers in the zero-field region; the event vertex is then determined as the most probable common point of origin of all reconstructed tracks. The Spectrometer provides good resolution along the z and x axes, making it a useful complement to the Vertex detector, which has good resolution only along y and z.

The average energy deposited by a charged particle in a silicon strip is proportional to the distance it travels; the size of the signal in an Si pixel can therefore be used to estimate the angle of incidence of the charged particle, and this forms the basis of the Octagon vertex-finding method. The vertex of a given event is estimated as the position most likely to have produced the observed distribution of Octagon hits. However, even for a fixed path length, fluctuations in energy loss are large and this limits the resolution of Octagon-based vertex-finders to around 1-2 cm. On the other hand, the Octagon has a large angular coverage ($|\eta| < 3.2$) and this gives it two advantages over vertices from the Spectrometer and Vertex detectors: its fiducial vertex range is much larger, and it is able to determine a vertex for more peripheral (and hence lower multiplicity) events.

The T0 Cerenkov detectors can measure the event vertex with a resolution of roughly 4 cm. They are used for an online vertex trigger (see Section 2.9) and can provide a broad sanity-check for the more precise silicon-based vertex procedures.

To obtain the best possible vertex position, from the wealth of information available, a composite vertex is created. For each event, all valid vertices from the different sub-detectors are checked for compatibility and the best x, y, z positions are determined independently, using an algorithm tuned for the properties of each vertex-finder, as found from Monte Carlo simulations.

The results of a Monte Carlo study of the resolution of vertices from the Vertex, Spectrometer and Octagon detectors, along with the composite 'Selected Vertex,' are shown in Figure 3-5.

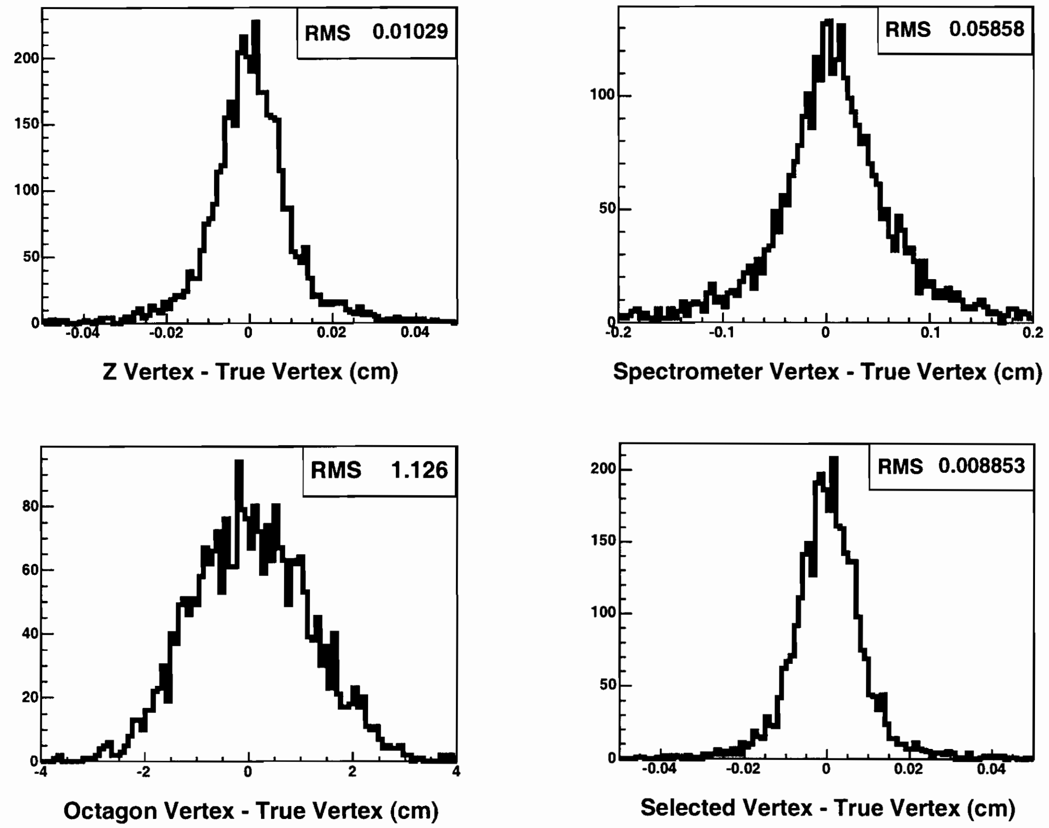

Figure 3-5: Resolution of vertices from different sub-detectors, obtained from Monte Carlo simulations. To avoid sounding silly, 'Z Vertex' is the name given to the vertex found using the Vertex detector; 'Selected Vertex' is the composite vertex created by combining information from all available vertices.

3.5.3 Beam Orbit

While the longitudinal z-position of the collision can vary over a large range, the location in the transverse xy-plane is closely controlled by the colliding-ring's beam optics and is much more stable: the 'beam spot' has a typical radius of 1 mm or less. The central beam orbit position for each run is determined from the vertex finders and stored; this information can be useful for analysis of the data.

Chapter 4

Track Reconstruction and Particle Identification

After calibration, signals in the Spectrometer pixels are combined using clustering algorithms to create hits. Hits from different layers are used to reconstruct the trajectories of charged particles through the Spectrometer; their momenta are determined from the curvature in the magnetic field. The energy deposited in the silicon can also be used to provide information on the species (pion, kaon or proton) of the reconstructed particle track. Track trajectories are extrapolated to the Time-of-Flight walls to search for matches with the hits recorded there - this provides particle identification up to higher momenta than just using the Spectrometer energy loss method.

4.1 Spectrometer Hits

The PHOBOS silicon Spectrometer has a high degree of pixelisation, to enable precise position determination. However, this means that charged particles which traverse the sensor at an oblique angle are likely to deposit energy in more than one pixel. Signals in neighbouring pads must therefore be combined by merging algorithms to recreate properly the original hit. The PHOBOS Spectrometer has a much narrower angular coverage and much coarser pixelisation in the vertical direction than in the horizontal direction, so hit-merging is only performed along the horizontal direction.

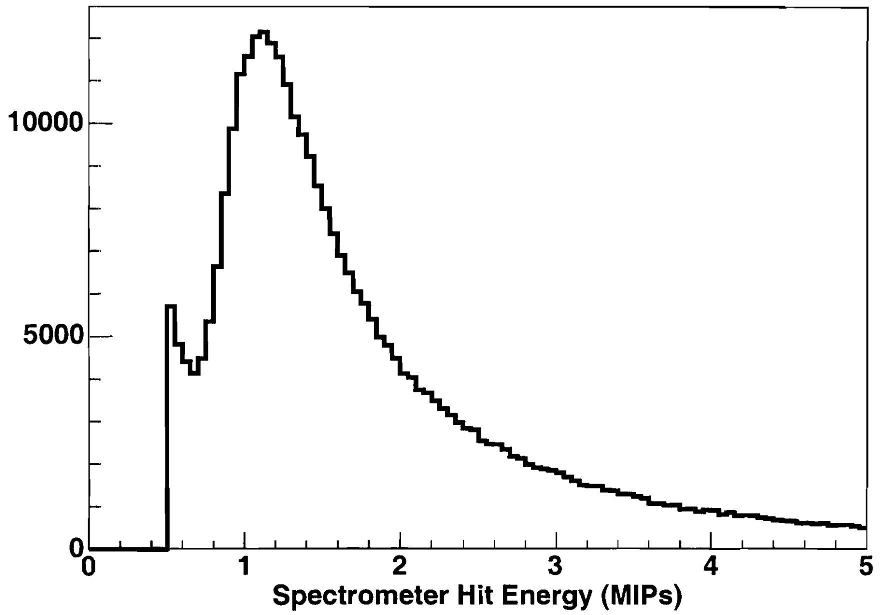

Figure 4-1: Energy distribution of Spectrometer hits, created from calibrated pixel readings by the merging algorithm. The units of energy are Minimum Ionizing Particles, ~80 keV for the PHOBOS silicon sensors. The minimum cut-off for hit-reconstruction at 0.5 MIP is clearly seen.

For each Spectrometer sensor, the hit-merging procedure scans all pads along a row. A pad becomes a candidate for merging if it contains an energy greater than the defined noise threshold of 0.15 MIP (the MIP is defined in Section 4.3.1; it has a value of ≈80 keV). The algorithm continues to merge adjacent pads in the row provided each one is above the noise threshold, or until it reaches a defined limit on the maximum number of merged pads (8). The summed energy of each set of merged pads is then checked to ensure that it is greater than the defined hit threshold of 0.5 MIP - if so, a hit is created and its position in the sensor is assigned based on the energy-weighted mean of the positions of the contributing pads. The hit stores the list of pads from which it was created, since this information is very useful, particularly for Monte Carlo simulation studies.

Figure 4-1 shows the energy of the Spectrometer hits, after merging. The cut-off at 0.5 MIP is clearly seen. The peak energy of the distribution is larger than 1 MIP because the hit energies are not corrected for angle-of-incidence.

4.2 Track Reconstruction

Figure 4-2 shows the PHOBOS Spectrometer tracks found in a typical Au+Au event.

The PHOBOS track reconstruction package[1] uses a straight-line road-following algorithm to find tracks in the low magnetic field region, and an algorithm based on Hough transformations to find tracks in the high-field region of the Spectrometer. The two track segments are then combined to form tracks which traverse the entire Spectrometer; a track-fit procedure then determines the momentum which best fits this full track trajectory.

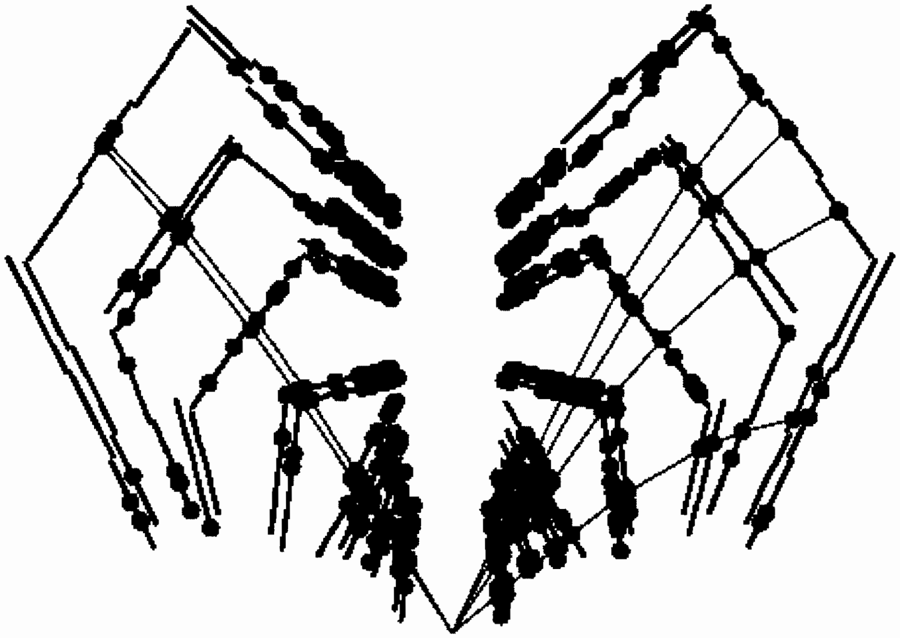

Figure 4-2: Example showing reconstructed tracks in an actual PHOBOS Au+Au collision event. Blue dots are Spectrometer hits; red dots are those hits which have been assigned to a reconstructed track.

---

[1] See also [94, 95] for other descriptions of the PHOBOS track reconstruction procedures

4.2.1 Multiple Scattering of Charged Particles

As charged particles pass through matter, they undergo elastic collisions with nuclei, resulting in deflections from their original path. Understanding this process is important for reconstructing tracks.

For a single collision, the cross-section for Coulomb scattering is given by the well-known Rutherford formula (see for example [96]). This cross-section has a $1/\sin^4(\theta/2)$ angular dependence, so that most collisions only produce very small-angle scattering of the incident charged particle.

When the material is sufficiently thick that the incident particle suffers many collisions, the process is called 'multiple scattering.' This case was treated by Moliere and a description can be found in [96]. The distribution of final deflection angle has a Gaussian core resulting from many small-angle scatterings, comprising 98% of the distribution, and large tails arising from single large scatterings. The tails therefore broadly follow the Rutherford formula.

For a given material, the mean scattering angle depends on the velocity and momentum of the incident charged particle according to:

$$\theta_0 \propto \frac{1}{\beta c p} \tag{4.1}$$

Thus low momentum tracks have larger mean deflections and are therefore more difficult to reconstruct accurately.

4.2.2 Finding Straight Tracks

Straight tracks are formed from hits in the first six Spectrometer layers. As can be seen from Figure 2-8, the magnetic field strength is very small in this region and charged-particle trajectories are to a good approximation straight lines.

The Au+Au tracking procedure assumes that the event vertex position is independently known to high precision.

Track seeds are generated from two-hit combinations of hits from the first and second layers. For each hit, the θ and ϕ angles relative to the event vertex are calculated, then $\Delta\theta$ and $\Delta\phi$ are computed for each possible two-hit combination (one from each layer); if both are less than a maximum cut-off value, then this combination is considered a suitable track-seed. Each layer combination has its own maximum allowed value for $\Delta\theta$ and $\Delta\phi$.

The track seeds are then propagated one-by-one to the next layer. All SpecHits in the next layer are compared to the last hit of the current track and the new hit is added to the track if it passes the $\Delta\theta$ and $\Delta\phi$ criteria for that particular layer-pair. If more than one hit passes the criteria for the current track, then the track gets copied and the hit is added to the copy - there is no selection of a 'best hit.'

Track candidates are re-evaluated after each layer, based on their number of hits and the quality of their straight-line fit. Tracks are allowed to miss a hit in one layer, but once two hits are missed, that track is removed from the candidate list. The remaining tracks are fitted to a straight line using a standard χ^2 fitting procedure [90] and the minimum fit-probability required to keep the candidate is 0.0005.

This process is repeated for each layer. When all six layers have been processed, final selection among the candidates is performed. The requirement is that final tracks cannot share more than one hit: if two tracks share two or more hits, the track with the lowest straight-line fit probability is discarded. The fit is performed in the xz-plane using all six layers and in the yz-plane using only the first four layers (because the vertical pixelisation increases considerably in layer 5 onwards, as can be seen from Table 2.1).

4.2.3 Finding Curved Tracks

For Spectrometer layers 8 and beyond, the effect of the magnetic field becomes important and a different method must be used to reconstruct charged particle tracks.

Neglecting the effects of multiple scattering, the trajectory of a charged particle in the PHOBOS xz-plane is determined by its vertex position, total momentum p and polar angle θ.

The event vertex position is found independently of the tracking algorithms (see Section 3.5 for a description of the PHOBOS vertexing procedures). With the vertex known, a combination of two Spectrometer hits in different layers can then be translated into the (p, θ) variables of the track which would have produced these hits.

This translation is performed by means of a pre-calculated look-up table. The hit combination is parametrised by two angles: the polar angle of the first hit relative to the event vertex, and the angle of the second hit relative to the first. Tables relating these two variables to (p, θ) are generated for 0.5 cm vertex bins in the range $-15 < v_z < 10$ cm from Monte Carlo simulations of pions passing through the Spectrometer. In practice, it is convenient to work with $1/p$ as a variable instead of p. Based on the pixelisation of the Spectrometer in a given layer, the uncertainty in this $(1/p, \theta)$ assignment can vary, so the errors for each combination are also stored in the table. There are separate tables for positive and negative particles, which of course bend in opposite directions in the magnetic field and therefore follow different trajectories.

The curved-tracking algorithm then iterates through the Spectrometer layers in a pre-defined sequence of layer-pairs. Due to the geometry of the Spectrometer, different layer-pairs are used for tracking in the central and outer wing regions. Tracking is not currently performed in the inner wing region, because the hit density there is so large that this tracking method would consume too large amounts of computation time.

For each layer-pair, all possible two-hit combinations are made and $(1/p, \theta)$ for that combination is determined from the look-up table (not all two-hit combinations are necessarily possible for a given event vertex). This two-hit combination is called a 'stick.' Once sticks have been produced for all layer-pairs, the algorithm attempts to join sticks from different layer-pairs into 'chains.' This is done by looking for groups of sticks that form a cluster in $(1/p, \theta)$-space, which would indicate a track that passed through all appropriate layers. The cuts on this cluster are deliberately chosen to be quite loose, so as to assemble as many track candidates as is compatible with an acceptable processing time. The procedure is to start with one stick from the first layer-pair, then take the second hit from this stick and search all sticks in the next layer-pair that begin with this hit. Δp and $\Delta\theta$ relative to the first stick are calculated, and if they satisfy the cuts then this stick is added to the chain. If more than one stick passes the cuts, the chain is duplicated and each stick is added to a new copy of the chain. Chains are propagated through all layers in this way.

The momentum and polar angle of the complete chain is taken as the error-weighted average of all component sticks. The chains are assigned a charge based on their bend-direction in the magnetic field.

The result of this curved-tracking stage is a container of chains comprising six Spectrometer hits, beginning in layer 8, and with approximate values of p and θ for each chain.

4.2.4 Matching straight and curved tracks

There are now straight tracks from layers 0-5 and curved tracks from layers 8 onwards - these must be joined to form full tracks that traverse the entire Spectrometer.

The two track pieces are checked for consistency in their θ angle (because θ is known for both straight and curved tracks) and their average energy per hit. The latter check would prevent, for example, a low momentum proton straight track being joined to a pion curved track. There is also a consistency test along the azimuthal angle, where the straight track is propagated outwards in the vertical direction and compared to the vertical position of the hits from the curved-part.

If a straight and curved piece satisfy the above requirements, then a full track is created from all component hits. This track is assigned a preliminary total momentum from the curved chain and the Cartesian momentum components are then calculated based on the θ and ϕ angles from the straight part. A truncated mean energy loss per hit is also calculated; this is described in Section 4.3.2.

Full tracks are also compared to each other and are not allowed to share more than two hits. If any pair of tracks does share more than two hits, then the track with the lower fit-probability (described in the next section) is discarded. These stringent requirements are designed to enhance the purity of the reconstructed track sample.

4.2.5 Determining the Track Momentum

Once the hit pattern of a Spectrometer track has been reconstructed, the goal is to accurately determine the track momentum - that is to say, what is the most probable track momentum that would produce this particular pattern of hits?

Maximum Likelihood Estimation

The best momentum is found using the traditional χ^2-fitting technique. Suppose there are N data-points (x_i, y_i) with a measurement errors σ_i, and the data are known to obey a functional relationship $y(x; a_1 \ldots a_M)$ which depends on the set of parameters a_i whose values are not *a priori* known. Then the values of the parameters a_i which maximise the likelihood that the function actually produced the observed data is found by minimising the quantity:

$$\chi^2 = \sum_{i=1}^{N} \left(\frac{y_i - y(x_i; a_1 \ldots a_M)}{\sigma_i} \right)^2 \tag{4.2}$$

Even though large-angle scatterings of the charged particle and the pixelisation of the silicon detector introduce errors in the hit position that are not Gaussian, χ^2-minimization remains a good method for finding the optimum momentum for a reconstructed track.

Charged-Particle Trajectory in Magnetic Field

Given an initial momentum vector and using a detailed map of the PHOBOS magnetic field, the trajectory of a charged particle is obtained in a series of steps, by propagating the momentum vector at a given point through a certain step-size. This propagation is performed by a Runge-Kutta algorithm (see [90]). A momentum-dependent variable step size is used in regions of large field gradient to give a more accurate trajectory. At this stage, the trajectory does not yet include the effects of multiple scattering or energy loss by the charged particle.

The deviation of a hit from the trajectory is called the hit *residual* and is calculated as the minimum distance of the hit from the trajectory, interpolating between trajectory points as necessary. In the χ^2 definition above, $y_i - y(x_i; a_1 \ldots a_M)$ is the residual.

Covariance Matrices

The correct error σ_i also has to be assigned to each point. This error has two sources: pixelisation of the Spectrometer sensors, and deflection of the charged particle due to multiple scattering.

Another effect of multiple scattering is that deflection in one layer tends to produce systematic offsets in the residuals of hits in later layers, which introduces correlations between the errors on different points. These correlations must be taken into account in order to properly determine the errors on the best-fit parameters that come out of the χ^2-fitting procedure.

The correlation between the errors on two points x_i and x_j is encapsulated in the *covariance matrix*:

$$cov(x_i, x_j) \equiv C_{ij} = \langle (x_i - \bar{x}_i)(x_j - \bar{x}_j) \rangle \tag{4.3}$$

The diagonal elements of the covariance matrix are seen to be $C_{ii} = \sigma_i$, the errors on each individual point.

Covariance matrices for the PHOBOS tracking package are pre-generated by simulating the passage of pions through the detector, including the effects of energy loss and multiple scattering. This means that realistic errors σ_i are used in the χ^2-fitting stage, and that the correlation between errors is correctly accounted for, so that a meaningful estimate of the 'goodness-of-fit' can be made. The covariance matrices are stored in look-up tables to be used during the track reconstruction.

Downhill Simplex Minimisation

Minimisation of the χ^2 to find the best trajectory is performed using the *downhill simplex* method [90]. This method of multi-dimensional minimisation does not require knowledge of the derivatives of the function being minimised; this is important for the PHOBOS tracking because the non-uniform magnetic field means that there is no analytic form for the particle trajectory, so derivatives cannot easily be computed.

In N dimensions, a *simplex* is defined as the geometric figure consisting of $n+1$ points: in $2D$ a simplex is a triangle, in $3D$, it is a tetrahedron, for example. The minimisation technique involves manipulating this figure in a well-defined way through the N-dimensional functional topography until it finds a (local) minimum. The most common steps are reflections of the highest point of the simplex through its opposite face. The simplex can also expand along a direction to take bigger steps, or it can contract as it attempts to 'squeeze' through the landscape in search of a minimum. The process ends when the incremental improvement after a step falls below a user-defined tolerance. It is recommended procedure to restart a multi-dimensional fit at the point where it converged, to ensure that the process was not fooled by a single anomalous step.

The momentum reconstruction package attempts to simultaneously fit the best track vertex position and the best initial momentum vector. Initial values for the parameters are obtained from the curved-track chain, then randomized by a small amount, to avoid potential systematic bias in the minimisation.

Momentum Correction for Kaons and Protons

The covariance matrices are generated from pion trajectories and at low momentum, kaons and protons undergo larger energy loss than pions as they traverse the silicon sensors (this is described in Section 4.3.1). The result is that the momentum reconstruction for low-p kaons and protons will be systematically shifted towards a lower assigned momentum than the particle's true momentum. This is illustrated in Figure 4-3. A fit to this graph is used to correct the reconstructed momentum for particles which are identified as kaons or protons.

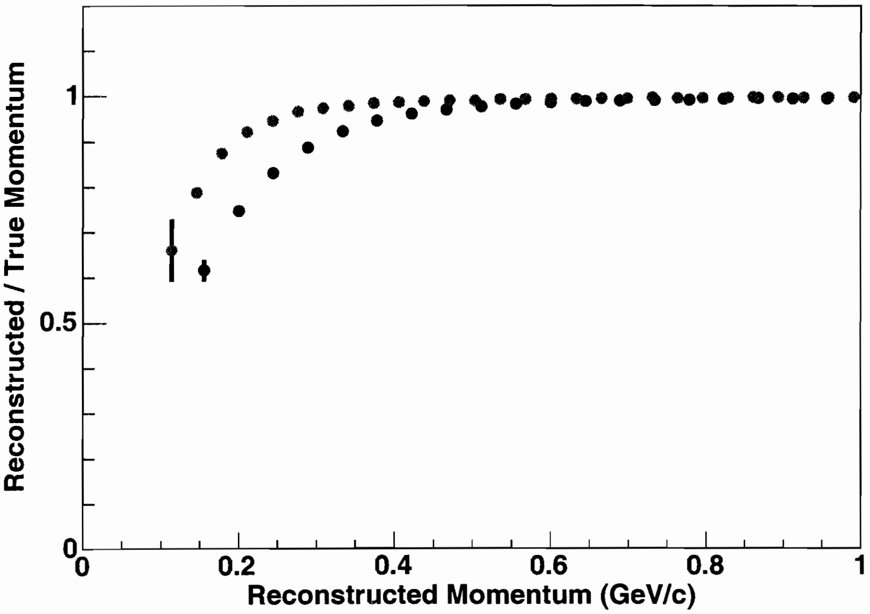

Figure 4-3: Reconstructed divided by true momentum for simulated kaons (green) and protons (blue), illustrating the systematic deviation in reconstructed momentum that arises from the inherent assumption in the track-fitting procedure that all particles are pions. This is plotted versus the reconstructed momentum value, and a fit to this graph is used to correct the reconstructed momentum values for identified kaons and protons.

4.2.6 Final Track Selection

There are two important selection cuts made on the reconstructed tracks: a cut on the 'goodness-of-fit' of the track using the χ^2 fit-probability; and a cut on the distance of closest approach of the track to the event vertex.

Track fit probability

The standard χ^2-fitting method not only determines the set of parameters that best fit the observed data but also provides estimates of the errors on those parameters and a statistical measure of the goodness-of-fit. For example, it is possible to find the best momentum that fits a given hit pattern, but if all hits wildly deviate from the predicted trajectory then this collection of hits is probably not actually a true track.

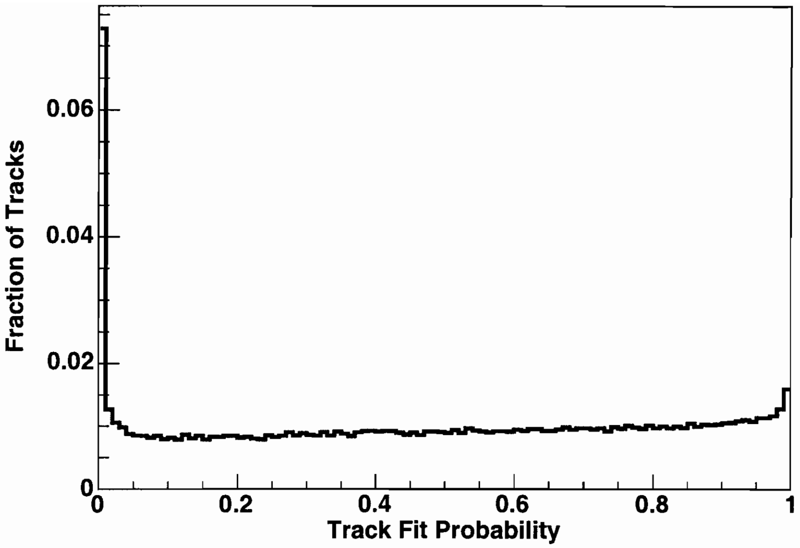

Figure 4-4: The distribution of the track fit-probability is shown for reconstructed tracks from Au+Au data. A cut on $prob > 0.04$ is applied to the data to remove the low-probability tracks, most of which are due to incorrect hit associations.

Once the best-fit parameters are known, the probability that a true track with these parameters would actually produce the observed hits can be calculated. This formally requires that the variances in the hits positions be Gaussian; since this is not the case, the 'fit-probability' cannot be interpreted as a true probability, but it is nevertheless a very useful indicator of goodness-of-fit.

The distribution of fit-probability should theoretically be a constant ranging between 0 and 1. The distribution which is observed in reconstructed tracks from PHOBOS data is shown in Figure 4-4. There are two contributions to the spike at low fit-probability. First, perfectly valid tracks which happen to undergo large-angle multiple scatterings are assigned a low fit-probability, because the assumption of Gaussian variances in the hit positions significantly underestimates the probability of this occurring. The remaining low-probability tracks are due to incorrect hit associations and should be discarded. In PHOBOS track-based analyses, the standard cut is to require that a track have a fit-probability $prob > 0.04$.

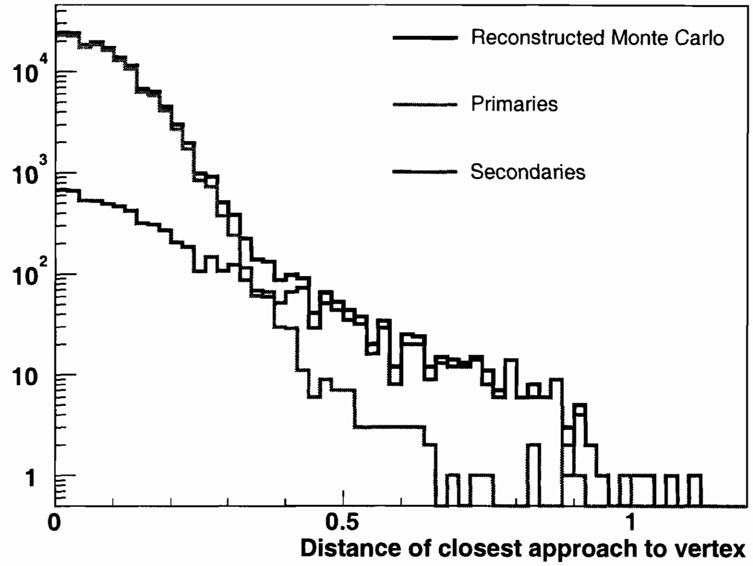

Figure 4-5: Distance of closest approach of track to event vertex (DCA), obtained from Monte Carlo simulations where reconstructed tracks can be identified as being from primary or secondary particles. The DCA is shown to have sensitivity for distinguishing between primaries and secondaries. A cut on DCA < 0.35 cm is applied to the data.

Distance of Closest Approach to Vertex (DCA)

Although the track reconstruction package starts from the position of the event vertex, the final found tracks after momentum-fitting are not constrained to originate there. The distance of closest approach of the track to the event vertex (DCA) can then be calculated, and as illustrated in Figure 4-5, Monte Carlo studies have demonstrated that this DCA quantity can be useful for distinguishing between primary particles and secondaries which did not originate from the event vertex. A cut on DCA< 0.35 cm is applied to reconstructed tracks to reduce the contamination from non-primary particles.

4.2.7 Tracking Efficiency and Momentum Resolution

The tracking efficiency and the momentum resolution achieved by the complete tracking procedure are shown in Figures 4-6 and 4-7 respectively. The efficiency is roughly

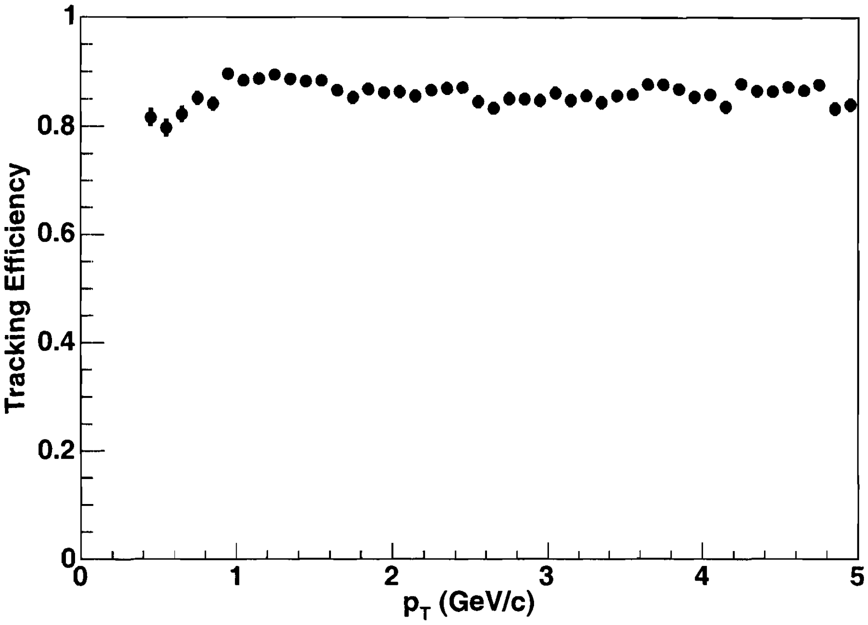

Figure 4-6: The tracking efficiency as a function of transverse momentum p_T. This is averaged over vertices in the range $-10 < v_z < 10$ cm and includes both central and outer-wing Spectrometer regions.

90% for $p_T > 1$ GeV/c, tailing off for lower p_T due to the effects of multiple scattering. The momentum resolution ($\delta p/p$) is about 1% for total momentum $p \sim 1$ GeV/c and rises linearly with p, but is still less than 5% for $p = 8$ GeV/c.

4.3 Particle Identification using Specific Energy Loss (dE/dx) in Silicon

Charged particles passing through matter will undergo inelastic Coulomb interactions with atomic electrons, resulting in energy loss of the incident particle. For particles such as pions, kaons and protons, their mass can be approximated as being infinitely heavy relative to the atomic electrons. In this limit, the energy loss of the incident particle is related only to its velocity and hence simultaneous measurement of momentum and energy loss allows for identification of charged particles according to their mass.

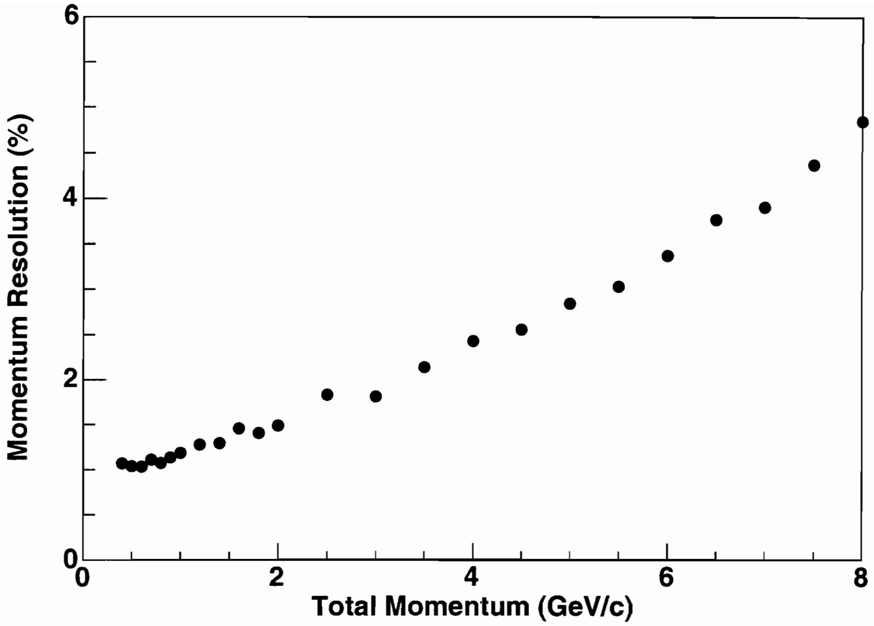

Figure 4-7: Momentum resolution achieved by the complete tracking procedure.

4.3.1 Mean Energy Loss; Bethe-Bloch Formula

The quantum-mechanical calculation of the mean energy loss by a charged particle due to collisions with atomic electrons as it passes through matter is enshrined in the Bethe-Bloch formula [10]:

$$-\frac{dE}{dx} = Kz^2\frac{Z}{A}\frac{1}{\beta^2}\left[\frac{1}{2}\ln\frac{2m_ec^2\beta^2\gamma^2T_{max}}{I^2} - \beta^2 - \frac{\delta}{2}\right] \tag{4.4}$$

where $K = 4\pi N_A r_e^2 m_e c^2 = 0.307$ MeV cm^2

$N_A = 6.022 \times 10^{23}$ mol^{-1} is Avogadro's number

$r_e = e^2/4\pi\varepsilon_0 m_e c^2 = 2.817$ fm is the classical electron radius

$m_e c^2 = 0.510$ MeV is the electron mass

ze is the charge of incident particle

Z is the atomic number of absorbing material

A is the atomic mass of absorbing material in g mol^{-1}

$\beta = v/c$ where v is the velocity of incident particle

$\gamma = 1/\sqrt{1-\beta^2}$

T_{max} is the maximum kinetic energy which can be imparted to a free electron in a single collision

I is the mean excitation energy in eV

and δ is the density correction.

The density correction arises because the electric field of a high-energy incident charged particle will polarize the atoms in the absorbing material. This induced polarisation screens the charge of the incident particle, so the energy loss due to electromagnetic interactions is less than it would otherwise be. The screening effect will clearly be greater if there is a greater density of atoms in the absorbing material.

At non-relativistic velocities, the energy loss is dominated by the $1/\beta^2$ term and dE/dx decreases with momentum. A minimum is reached at $\beta \approx 0.96$, where particles are said to be *minimum ionizing.* Above this, the $1/\beta^2$ term is approximately constant and dE/dx starts to rise again logarithmically. However, for solid-state materials, this *relativistic rise* is largely cancelled out by the density correction, and all higher-energy particles deposit basically the same amount of energy, which is called the *Minimum Ionizing Peak* (MIP). For the silicon sensors in the PHOBOS Spectrometer, this MIP value is approximately 80 keV.

4.3.2 Fluctuations in Energy Loss; Truncated Mean dE/dx

The Bethe-Bloch formula works well for describing the mean energy loss for a charged particle of a given velocity, but fluctuations about this mean are also important, and can be large. This is often called *straggling.* Straggling arises because of statistical fluctuations in the number of collisions between the charged particle and the atomic electrons, and also in the energy transferred per collision.

The nature of the energy loss fluctuations depends on the thickness of the absorbing material. For a relatively thick absorber, the number of collisions will be large; from the Central Limit Theorem of statistics, therefore, fluctuations in energy loss will have a Gaussian form with a calculable width. The 300 μm-thick PHOBOS silicon sensors, however, are best treated by the thin absorber approximation, where the mean energy loss is much less than the maximum allowable energy loss T_{max}.

This case was treated by Landau, making the approximation that $T_{max} \rightarrow \infty$ and assuming that the atomic electrons are essentially free. For an absorber of thickness x, with a mean energy loss Δ (obtained from the Bethe-Bloch formula), the Landau distribution[96] is:

$$f(x, \Delta) = \phi(\lambda)/\xi \tag{4.5}$$

where

$$\phi(\lambda) = \frac{1}{\pi}\int_0^\infty \exp(-u \ln u - u\lambda) \sin \pi u \, du \tag{4.6}$$

$$\lambda = \frac{1}{\xi}[\Delta - \xi(\ln \xi - \ln \epsilon + 1 - C)] \tag{4.7}$$

$$\ln \epsilon = \ln \frac{(1-\beta^2)I^2}{2mc^2\beta^2} + \beta^2 \tag{4.8}$$

and C = 0.577... is Euler's constant.

ϵ is essentially the minimum energy transfer allowed in this approximation. The function $\phi(\lambda)$ has to be evaluated numerically for a given value of λ.

The Landau distribution is highly asymmetric, with a large tail extending to high energies. As a result, the mean energy loss is always larger than the most probable value; these two quantities parametrize the distribution. Including the density effect, the most probable energy loss is given by:

$$\Delta_{mp} = \xi[\ln(\xi/\epsilon) + 0.198 - \delta] \tag{4.9}$$

The Landau treatment has been generalised by Vavilov to consider all scenarios for energy loss fluctuations, reducing to the Gaussian case for thick absorbers. Shulek has produced a further correction due to atomic electrons being bound in shells, not truly free. References for these can be found in [96].

Truncated Mean Energy Loss

Large fluctuations in energy loss worsen the resolution of particle identification. This is somewhat alleviated by the calculation of a truncated mean energy loss for each reconstructed charged particle.

Reconstructed tracks in the PHOBOS spectrometer typically consist of 12 (sometimes 11) hits. Once the track trajectory is known, the energy of each hit on the track can be corrected for the path length travelled in the sensor. The truncated mean energy loss of this track is calculated by first discarding the four hits with the highest energy and then averaging the remainder. This reduces the effect of the high-energy tail for energy-loss fluctuations.

4.3.3 Obtaining Raw Particle Yields

The truncated mean dE/dx for reconstructed charged particles in the Spectrometer is plotted as a function of the particle momentum in Figure 4-8. Bands due to pions, kaons and protons can be seen.

Unlike previous PHOBOS analyses of identified particles for antiparticle/particle ratios ([97, 54, 98, 99]), we do not attempt to identify particles individually, rather we just count the total number of particles of each species (within a momentum bin) in a statistical fashion. This allows a larger p_T range for identification.

The dE/dx data is binned in total momentum. For each bin, we extract the yields of each species by fitting the theoretical 'line-shapes' for pions, kaons and protons to the observed dE/dx histogram.

The theoretical dE/dx is taken from a form of the Bethe-Bloch equation:

$$\frac{dE}{dx} = \frac{E_0}{\beta^2}(b + 2\ln\gamma - \beta^2) \tag{4.10}$$

where $E_0 = 0.039$ is an overall constant, and $b = 20$ characterizes the relative strength of the logarithmic rise.

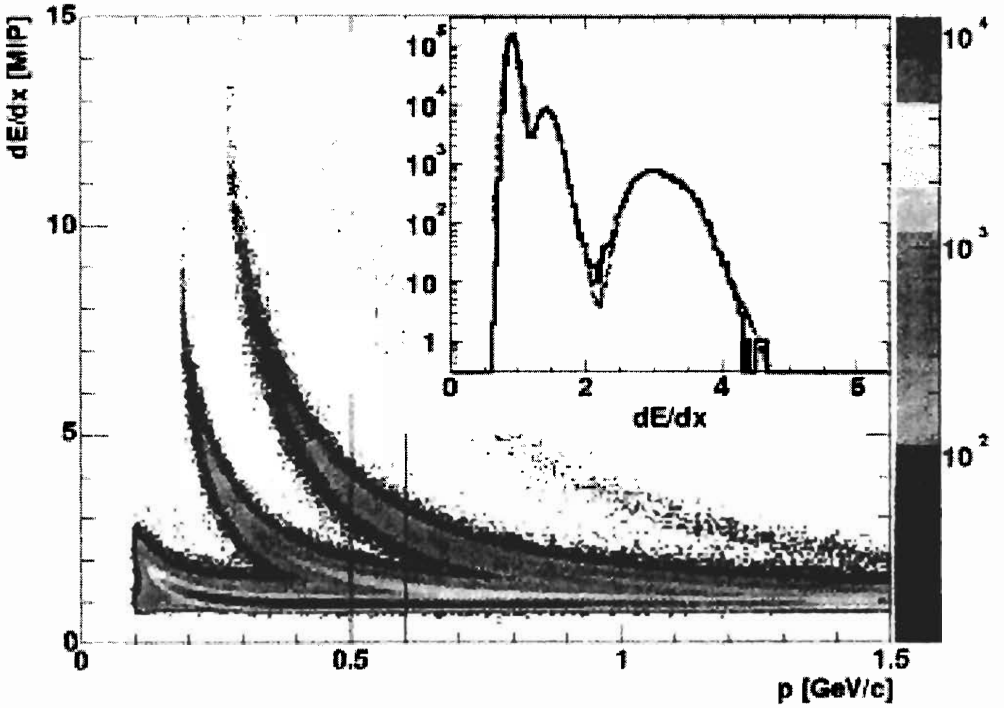

Figure 4-8: Specific energy loss (dE/dx) of particles in the silicon spectrometer as a function of their momentum.

The theoretical dE/dx is convoluted with a Gaussian to mimic the effect of the detector energy resolution. This resolution is found to be $\sim$ 0.07 MIP and to vary according to: $\Delta dE \propto dE^{0.8}$. In fact, two Gaussians are used in practice for each lineshape – this gives a better fit and accounts for the asymmetric dE/dx fluctuations. The proper line-shape also requires using the corrections to the reconstructed momentum for kaons and pions as described in Section 4.2.5.

With the pion, kaon and proton lineshapes determined for each momentum bin, the yields of each species are extracted by fitting these lineshapes to the histogram of data dE/dx. A 3-parameter fit is performed with the yields of the species as the free parameters.

For $p >\sim 0.6$ GeV/c, kaons and pions cannot be separately identified. We have to make an assumption as to the pion/kaon ratio and constrain the fit using this ratio. The assumed ratio comes from interpolating between the low momentum dE/dx and

TOF data regions where the species are separable.

4.4 Particle Identification using Time-of-Flight Measurement

Two particles with the same momentum p but different masses m_1 and m_2 will travel the same distance d with a time difference given by:

$$\Delta t = \frac{d(E_1 - E_2)}{pc^2} \tag{4.11}$$

where $E_i = \sqrt{m_i^2 c^4 + p^2 c^2}$.

Simultaneous knowledge of the momentum, path length and time-of-flight of a particle can therefore be used to identify particles according to their mass.

The results presented here use only the PHOBOS Time-of-Flight wall TB (see Section 2.6.3) because of low statistics in the other wall. TOF wall TB is at a distance of 5.4 m from the nominal interaction point and has a timing resolution of $\sim$ 140 ps – this should allow a 2σ separation between pions and protons up to roughly $p_T \sim 3.5$ GeV/c.

4.4.1 Making TOF Hits

After calibration as described in Section 3.3, the signal in each TOF sensor is checked to see if it has good timing characteristics and if sufficient energy was deposited to indicate a true charged particle detection. The timing and pulse-height signals from both photo-multiplier tubes at the top and bottom of the sensor are checked for consistency.

If the signal in the sensor passes these tests, then the start time obtained from the T0 detectors is subtracted, producing the transit time of the particle from the event vertex to the TOF sensor.

The T0 start time is simply obtained from the average of all T0 counters which fire for the event, after the signals have been calibrated.

4.4.2 Extrapolation of Spectrometer Tracks to TOF Walls

TOF particle identification requires association of TOF hits with Spectrometer tracks. Reconstructed tracks from the negative Spectrometer arm ('SpecN') are obtained from the tracking package, then re-fit without using the covariance matrices (see Section 4.2.5). This re-fitting gives a better estimate of the track momentum as it exits the Spectrometer, which makes for a more accurate extrapolation of the track out towards the TOF walls. However, the original momentum value obtained using the proper covariance matrix fit is still retained, as this is the quantity which is most important for physics, since it is the momentum of the track at its vertex position.

Tracks with the re-fitted momentum vector are then propagated outwards from the Spectrometer using the same Runge-Kutta algorithm as in the tracking procedure to determine their trajectory, because the magnetic field in the region between the Spectrometer and TOF walls is small but not zero. The step size for this extrapolation was chosen to be 1 cm.

For each TOF hit, the minimum distance to the extrapolated track is calculated; if this residual is less than a defined cut-value, then the TOF hit is associated with this Spectrometer track and a new 'TOFTrack' is made. The distribution of residuals is plotted in Figure 4-9. It is seen that even though the TOF wall TB is over 4 m away from the last layer of the Spectrometer, we are able to match TOF hits to reconstructed tracks with a resolution of better than 2 cm. In the analysis, only TOF hits which are matched to a track with a residual less than 4 cm are used. This prevents incorrect associations that might arise.

The path length travelled by the particle to the TOF wall is obtained from the extrapolated trajectory.

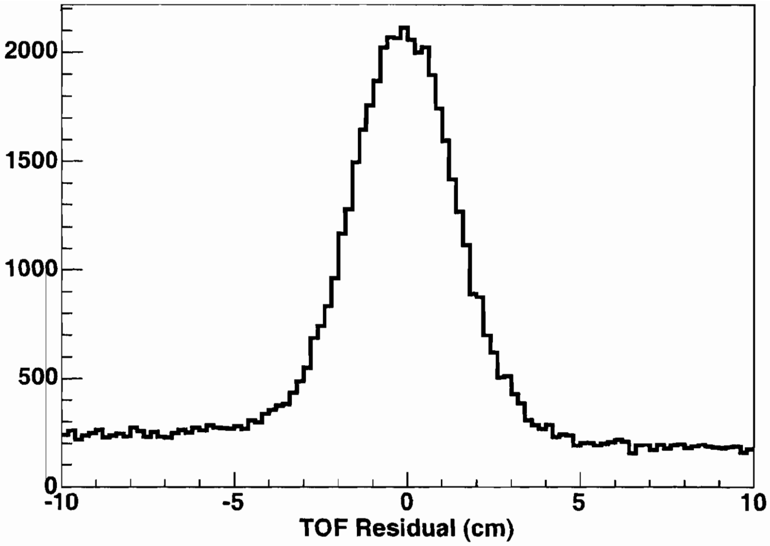

Figure 4-9: Residuals from matching extrapolated Spectrometer tracks to hits in the TOF walls. A cut on residual < ±4 cm is used in the analysis.

4.4.3 Obtaining Raw Particle Yields

For Time-of-Flight particle identification, it is common to use the *inverse velocity*:

$$\frac{1}{v} = \frac{t}{d} \tag{4.12}$$

The inverse velocity of TOF-tracks is plotted versus momentum in Figure 4-10, where one can see the separation of particle-species into bands of pions, kaons and protons.

Similar to the process used for dE/dx particle yields, the TOF data is partitioned into bins of total momentum and in each bin the inverse velocity of all particles is histogrammed. To obtain the yields for each species, one must construct the theoretical line-shape of pions, kaons and protons for this momentum bin and fit these lineshapes to the observed inverse velocity histogram.

For a particle of mass m and momentum p, the theoretical inverse velocity is:

$$\frac{1}{v} = \frac{t}{d} = \frac{E}{p} = \frac{\sqrt{m^2 + p^2}}{p} \tag{4.13}$$

This is convoluted with a Gaussian to represent the finite timing resolution of the detector. In practice, for each particle species, three Gaussians with the same mean but different amplitudes and widths are used to give the best fit. Physically, this is because the timing resolution can vary from run to run, since the calibration procedure is not perfect.

The pion, kaon and proton line-shapes are now simultaneously fit to the data $1/v$ histogram for each momentum. The free parameters are just the yields of the 3 species. The 3-parameter fit is performed using the MINUIT method implemented in the ROOT software framework; the fit attempts to minimise the difference between the area of the data histogram and the area of the 3 fitted line-shapes.

At a total momentum of $p \sim 2$ GeV/c, pions and kaons merge and become indistinguishable. At this point, the pion and kaon yields cannot be determined independently and we have to constrain the fit with an assumed pion/kaon ratio. This ratio is obtained by extrapolating from the region where the particles can be separated. By varying this ratio we determine the error on the proton yields that arises from having to make such an assumption.

Figure 4-11 shows an example of the fitting procedure applied to a TOF momentum bin.

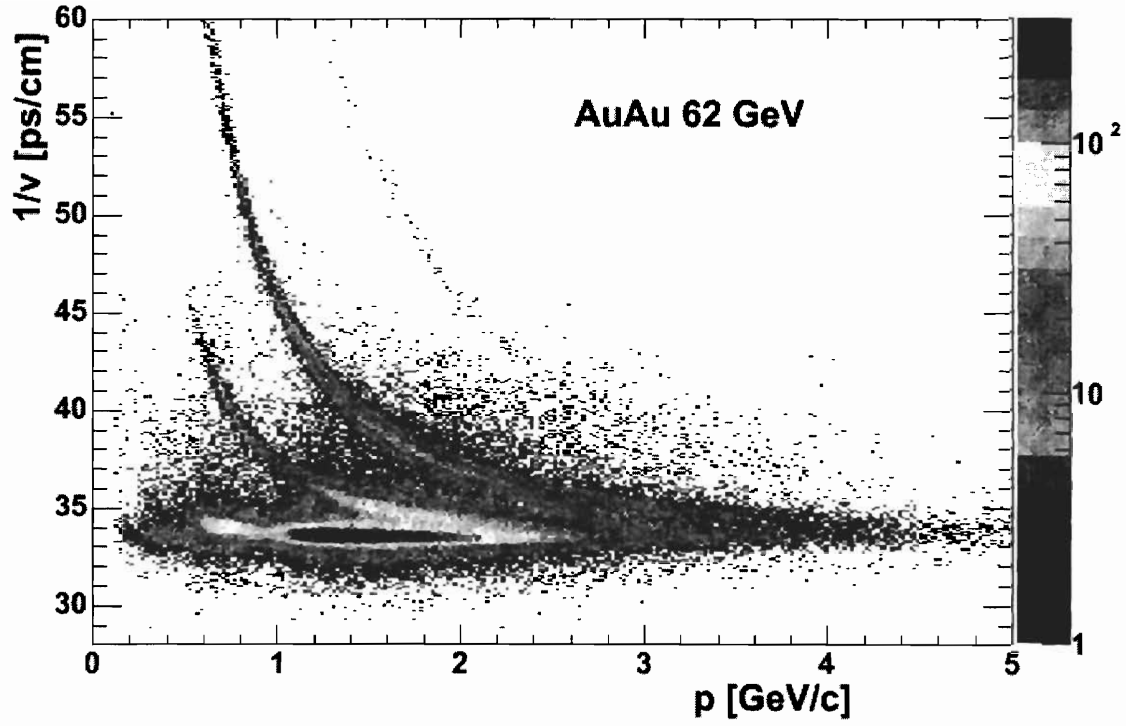

Figure 4-10: Inverse velocity as measured by Time-of-Flight detector, versus particle momentum, obtained from Spectrometer, for Au+Au collisions at 62.4 GeV.

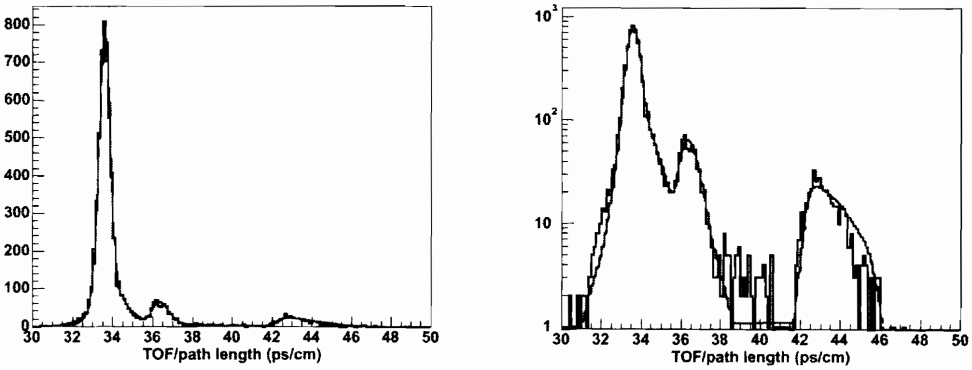

Figure 4-11: Example of the fit to extract the pion, kaon and proton yields from Time-of-Flight data in a momentum bin, shown with linear and log scales.

Chapter 5

Obtaining Transverse Momentum Distributions

The raw particle yields as a function of transverse momentum which emerge from the identification process need to be corrected for detector-dependent effects in order to obtain the true primary particle yields.

The largest single correction is due to the geometric acceptance of the detector and the efficiency of reconstructing particle tracks, which is momentum-dependent. The p_T distributions are also corrected for tracking 'ghosts', momentum resolution, secondary particles, feed-down from weak decays and the effect of dead channels.

Finally, the results obtained independently from the Spectrometer and TOF detectors are synthesized into an overall result.

5.1 Geometrical Acceptance and Tracking Efficiency

The most important corrections to the raw particle yields are to account for the geometrical acceptance of the detector and the momentum-dependence of the efficiency of the track-reconstruction procedures.

These two corrections are combined by generating individual Monte Carlo particles over a range of phase space, 'swimming' them through a full GEANT simulation of the PHOBOS detector. These events are then put through the complete track-

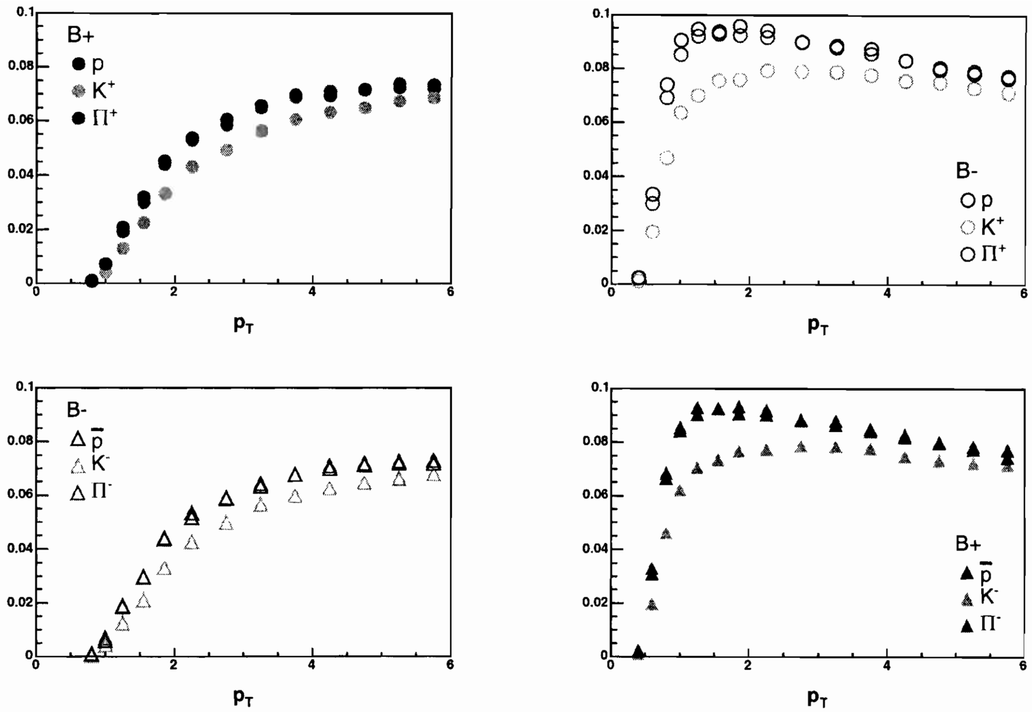

Figure 5-1: Combined tracking efficiency and geometrical acceptance correction as a function of transverse momentum p_T for Time-of-Flight tracks. Blue symbols are (anti)protons, red are pions and green are kaons. The four permutations of charge-sign and magnet polarity are shown separately.

reconstruction package, including the final track selection cuts on fit-probability and distance-of-closest-approach to the event vertex. For TOF tracks, the simulation also includes extrapolation to the Time-of-Flight walls and the hit-matching procedure. This gives the probability, as a function of transverse momentum, that an individual particle will leave a findable track in the Spectrometer and be properly reconstructed by the tracking procedure; this is then used to correct the raw spectra. This procedure is performed for all species separately.

Examples of these correction functions for TOF data are shown in Figure 5-1. Kaons have a lower overall efficiency than other species, because they can decay in flight before reaching the TOF wall – kaons in their rest frame have a lifetime of 1.237×10^{-8} s, which for a 500 MeV particle for example equates to a mean decay length of roughly 5 m.

5.1.1 Occupancy Correction

The procedure described above provides an appropriate correction for events containing only single tracks. A more realistic correction is obtained by embedding and reconstructing individual Monte Carlo tracks in real data events. This realistic correction is found to have a similar momentum-dependence to the single-track correction, but also an overall reduction in efficiency which varies with the density of hits in the Spectrometer. For the most central 62.4 GeV Au+Au collisions, this is about a 10% reduction from the single-track case.

5.1.2 Ghost Correction

The other part of the tracking-based correction procedure is to account for the 'purity' of the reconstructed tracks, since it is possible that the tracking package will produce spurious tracks from the available hits. Such tracks are commonly called 'ghosts'.

The ghost fraction will in general depend on both the track momentum and the hit-density. It is studied by reconstructing Monte Carlo events where the tracking output can be compared to the known input tracks. The p_T-dependence of the ghost contribution to the observed particle yields was found to be:

$$g(p_T) = Ae^{-0.83p_T} \tag{5.1}$$

The constant A depends on the density of hits in the Spectrometer. For the 62.4GeV Au+Au analysis, this constant ranges from 2% for the most peripheral bin to 5% for the most central.

No species-dependence was found for the ghost correction.

5.1.3 Momentum Resolution

As illustrated in Figure 4-7, the momentum resolution of the track-reconstruction package becomes inherently worse as the particle momentum increases. When combined with a steeply-falling p_T spectrum, this results in a smearing effect and requires correction. This has been done for previous unidentified charged hadron spectra [65, 69, 100], where the momentum range covered extends out to ≈ 5 GeV/c. For this analysis, however, the momentum range is restricted to $p_T < 3.5$ GeV/c and the momentum resolution correction to the particle p_T spectra has therefore been found to be very small (less than 2-3%). It is accounted for implicitly by the method used to produce the geometrical acceptance and tracking efficiency correction, and does not require a special correction of its own.

5.2 Feed-down from Weak Decays

Many of the particles produced in heavy-ion collisions are short-lived resonance states. Strongly-interacting resonances will decay on time-scales of the order of 10^{-24} s and the decay products are usually considered as part of the primary particle distribution. Products of weak decays, on the other hand, are typically generated $\sim 10^{-10}$ s after the interaction. Since the goal is to present the transverse momentum distributions of primary particles that come directly from the collision system, one must correct for contributions due to 'feed-down' from weak decays. Perhaps even more importantly, the contributions of feed-down particles to the measured spectra are detector-dpendent and therefore comparing results from different experiments is much less reliable if feed-down corrections have not been applied.

5.2.1 Λ and Σ Simulations

The major feed-down contribution to the proton yields comes from the Λ decay: $\Lambda \rightarrow p + \pi^-$. This process has a branching ratio of 63.9% and a lifetime expressed as $c\tau = 7.9$ cm; the daughter particles have a momentum of 101 MeV/c in the centre-of-mass frame.

The Λ, being neutral, leaves no trace as it passes through the silicon detectors. The PHOBOS tracking procedures are such that the daughter proton will only be reconstructed if the decay happens before the first spectrometer layer - since this layer is only roughly 10 cm from the nominal interaction point, the PHOBOS experiment has good sensitivity for distinguishing between primary and feed-down protons.

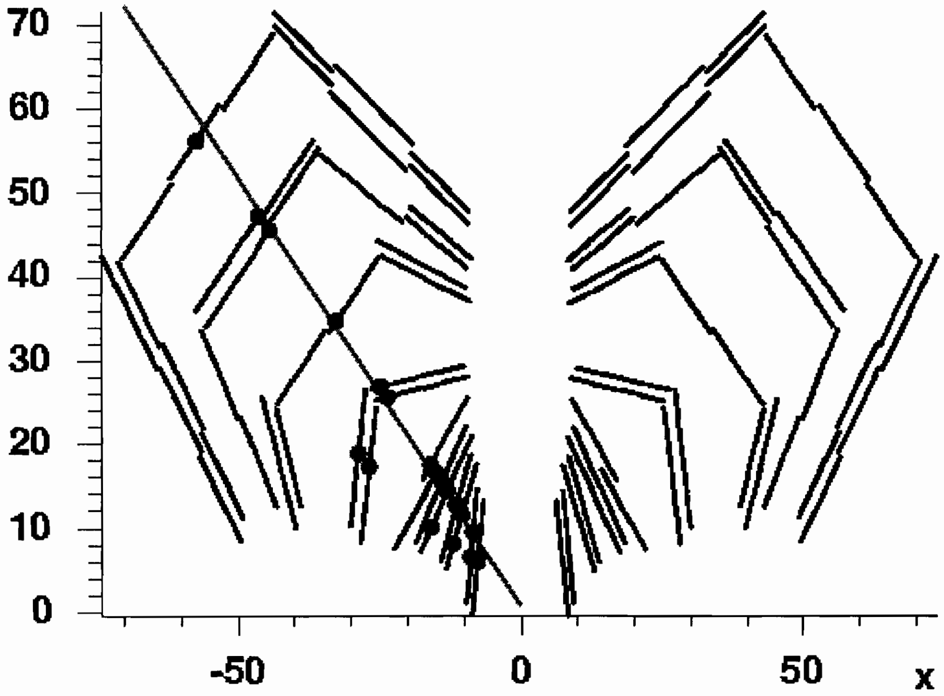

Figure 5-2: Monte Carlo simulation of a Λ decay in the spectrometer. The green line is the path that the neutral Λ would have taken if it did not decay; blue is the daughter proton and red the daughter pion; purple are electrons from scattering events. In this instance, the Λ decays before the first spectrometer layer, resulting in a findable daughter proton track.

The feed-down correction seeks to answer the question: for a given reconstructed proton, what is the probability that it is a primary particle and not the product of a weak decay? There are two pieces of information required for the answer: the relative yields of primary Λs to protons; and the relative efficiency of reconstructing primary protons and protons from Λ decays.

The GEANT Monte Carlo package was used to simulate Λ decays in the PHOBOS spectrometer. These decays are illustrated in Figures 5-2 and 5-3. Λs were generated with realistic transverse momentum distributions, and the efficiency of reconstructing the daughter proton at a given momentum can then be compared to the efficiency for primary proton reconstruction at the same momentum. This result can then be scaled according to the Λ/p ratio to determine, as a function of transverse momentum, the

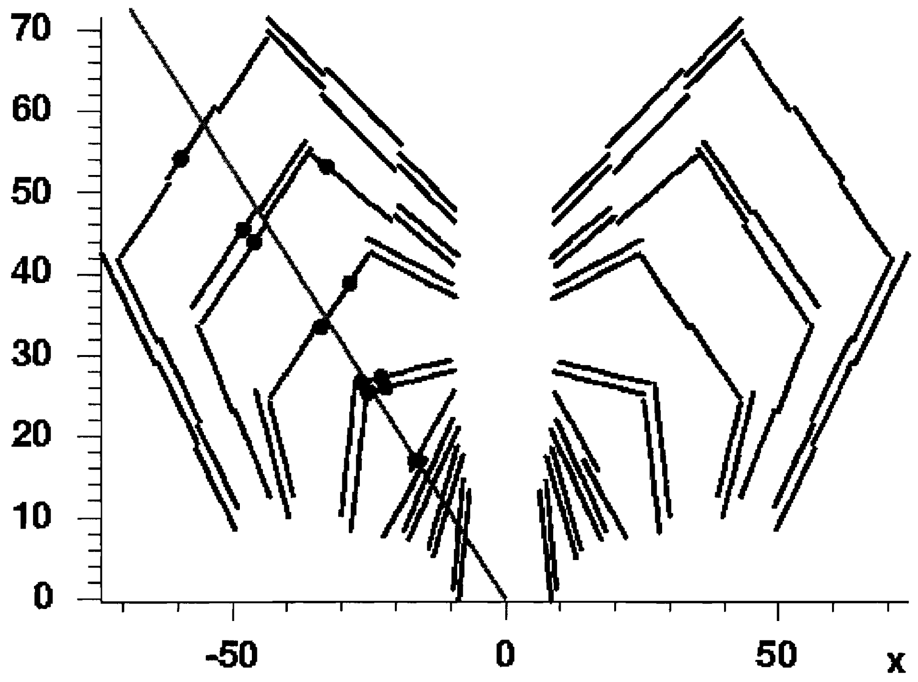

Figure 5-3: Monte Carlo simulation of a Λ decay in the spectrometer. The green line is the path that the neutral Λ would have taken if it did not decay, and the blue and red dots are the hits left by the daughter proton and pion respectively. On this occasion, the Λ decay happens inside the spectrometer and the daughter proton cannot be reconstructed.

fraction of observed protons which are expected to actually originate from Λ decays.

This method was tested on Monte Carlo events produced by the HIJING event generator and incorporated into a GEANT simulation of the PHOBOS detector. When the Λ/p ratio was known, this method was able to correctly describe the contribution from feed-down protons.

The PHENIX collaboration has measured $\Lambda/p = 0.89 \pm 0.07$ in Au+Au collisions at $\sqrt{s_{NN}} = 130$ GeV [101]. The value at 62.4 GeV is not yet known, which leads to an uncertainty in the feed-down correction. Based on the Monte Carlo simulations, the expected fraction of observed protons which are really feed-down particles is shown in Figure 5-4, for different values of the ratio in the range $0.7 \leq \Lambda/p \leq 1.1$.

Feed-down protons can also originate from the decay $\Sigma^+ \rightarrow p + \pi^0$, which has a branching ratio of 51.6% and $c\tau = 2.4$ cm. The Σ/p ratio has not been measured for Au+Au collisions at RHIC. The HIJING event generator predicts $\Sigma/p \approx 0.3$, and

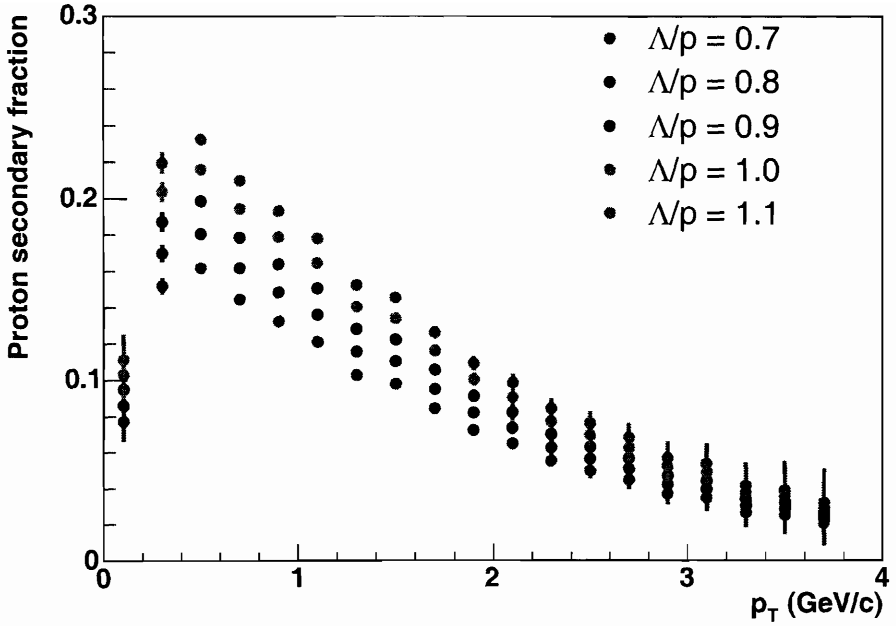

Figure 5-4: Expected fraction of observed protons which actually originated from Λ decays, for different values of the Λ/p ratio.

measurements from p+p collisions at similar energies have found a value of around 0.5. Σ decays were studied in the PHOBOS detector using the same techniques as for Λs; the results for the expected feed-down contribution are shown in Figure 5-5 for $0.1 \leq \Sigma/p \leq 0.9$.

5.2.2 Distance-of-Closest Approach to Event Vertex

As discussed in Section 4.2.6 and illustrated in Figure 4-5, the distance of closest approach of a reconstructed track to the event vertex (DCA) is a quantity which has some sensitivity for distinguishing between primary particles and those which did not originate from the event vertex. This is particularly useful for the study of feed-down particles.

From the Monte Carlo simulations described above, the DCA distributions for primary and feed-down protons were obtained. In the case of primaries, the DCA distribution is narrow, reflecting the fact that these particles really did originate at

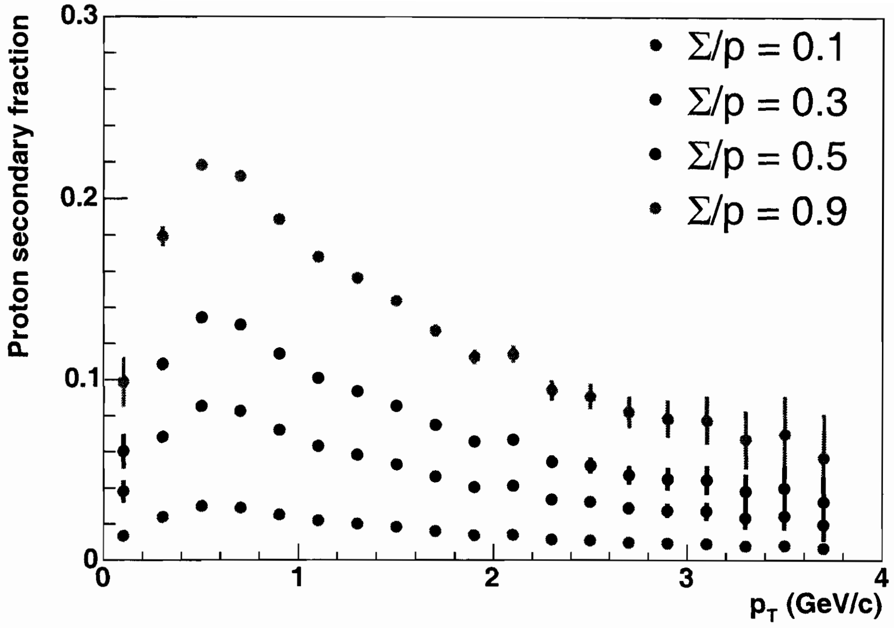

Figure 5-5: Expected fraction of observed protons which actually originated from Σ decays, for different values of the Σ/p ratio.

the event vertex. The distribution for feed-down particles has a tail which extends to much higher values of DCA. The cut on DCA < 0.35 cm used as part of track selection was found to remove $\sim$ 25% of the feed-down protons but $\ll$ 1% of the primaries.

These primary and feed-down DCA distributions are used in an attempt to reproduce the observed distribution for data protons. It was found that a feed-down contribution in the range of 25-30% gave the best consistency, and that less than 20% or greater than 35% seemed to be inconsistent with the data.

5.2.3 Proton Feed-down Correction

The final feed-down correction for protons is chosen to be the sum of the simulated contributions from Λ and Σ decays, assuming $\Lambda/p = 0.9$ and $\Sigma/p = 0.3$. This gives a correction which is consistent with the data-driven feed-down estimates from the DCA distributions. The relatively large systematic uncertainty on this correction comes from plausible variations in the Λ/p and Σ/p ratios (which have not yet been measured for 62.4 GeV Au+Au collisions) and are also consistent with the bounds obtained from the DCA analysis.

The function used for the fraction of observed protons which are feed-down products is:

$$f(p_T) = 0.39e^{-0.57p_T} + (0.22p_T - 0.16)e^{-5p_T^2} \tag{5.2}$$

with upper and lower bounds defined by $(1.22 + 0.1p_T)f(p_T)$ and $(0.78 - 0.1p_T)f(p_T)$ respectively.

5.2.4 Antiproton Feed-down

We determine the antiproton feed-down correction in relation to that for protons, by making the reasonable assumption that antiproton feed-down is dominated by $\bar{\Lambda}$ decays and therefore the main difference between proton and antiproton feed-down comes from differences in the $\bar{\Lambda}/\bar{p}$ and Λ/p ratios.

The value of these ratios in Au+Au collisions at $\sqrt{s_{NN}}$ =130 GeV has been measured [101] to be:

$$\frac{\Lambda}{p} = 0.89 \pm 0.07(\text{stat})$$

$$\frac{\bar{\Lambda}}{\bar{p}} = 0.95 \pm 0.09(\text{stat})$$

By considering the quark content of these states, one can postulate that they should be related by:

$$\frac{\bar{\Lambda}}{\bar{p}} = \frac{\Lambda}{p} \times \frac{K^+}{K^-} \tag{5.3}$$

This relationship was found to hold true for 130 GeV Au+Au collisions, where the kaon ratio was measured [97] to be: $K^-/K^+ = 0.91 \pm 0.07(\text{stat}) \pm 0.06(\text{sys})$.

We assume it holds at 62.4 GeV too, where we have measured $K^-/K^+ \approx 0.85$ (see Section A). Thus we take the antiproton feed-down correction to be roughly $1.0/0.85 = 1.18$ times the correction for protons.

5.2.5 Feed-down to Kaons and Pions

There are no weak processes which produce kaons as the final state thus there is no feed-down correction for the kaon yields. This expectation from theoretical grounds is verified by the measured DCA distribution for reconstructed kaons, shown in Figure 5-6. The data agree well with the distribution for primary kaons obtained from Monte Carlo simulations, with no indication of a tail that would correspond to feed-down contributions.

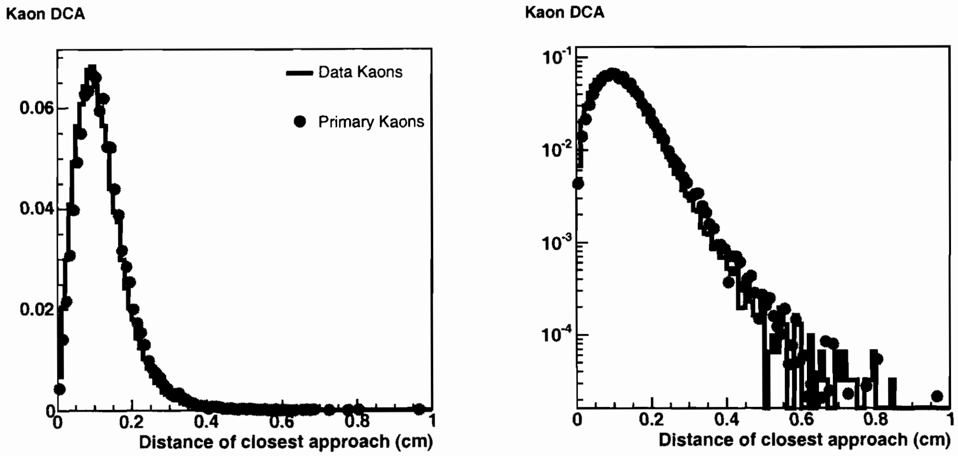

Figure 5-6: Distance-of-closest-approach (DCA) for kaons, shown on both linear and log scales. The data (black lines) agree well with the distribution obtained for simulated primary kaons (red dots), demonstrating the absence of any feed-down contribution to the kaon yields.

The DCA distribution for pions from data was also studied and compared to that for simulated primary pions. The feed-down contribution to the pion yields after applying the cut of DCA $<$ 0.35 cm was estimated to be less than 1% and was therefore neglected, but this was included as a 1% contribution to the overall systematic uncertainty on the pion spectra.

5.3 Secondary Particle Correction

Secondary particles are here defined as those that do not originate directly from the collision system but are not the products of weak decays. Thus, the main source of secondaries is scattering events as the primary particles pass through the beryllium beam-pipe and detector material.

Secondary particle production is studied using a GEANT simulation of the PHOBOS detector, with the HIJING event generator as the source of primary particles. On average, secondary particles tend to be produced with low momentum. This means that they are unlikely to be detected in our Spectrometer because they tend to bend out of the acceptance and also suffer more multiple scattering, making their hit pattern less likely to be reconstructed.

It was found that there *is* considerable contribution to the proton yield from secondaries, but only for $p_T < 200$ MeV. Since we do not extend this low in transverse momentum, the secondary correction to our measured proton yields is actually less than 1%. We therefore choose to neglect this correction, but make an additional contribution of 1% to the systematic uncertainty of the spectra.

The secondary contribution to kaon and pion yields was found to be negligible.

5.4 Spectrometer and TOF Dead Channels

Dead channels in the Spectrometer are not included in the initial geometrical acceptance correction. The identification of dead channels is described in Section 3.4.2. Since the tracking algorithms require hits in 11 or 12 Spectrometer layers, even a small fraction of dead channels in each layer can add up to a considerable contribution to the overall track yield.

The dead channel mask is applied to Monte Carlo track simulations and to a sample of the data. The difference in track yields between the masked and not-masked cases is then used to obtain the final correction. Although in principle the spatial distribution of dead channels can produce a momentum-dependent effect, it was found that, when averaged over the vertex range $-10 < v_z < 10$ cm used in this analysis, the dead-channel correction to the track yields is independent of p_T.

The final correction for Spectrometer dead channels is $8 \pm 5\%$. This is applied to both Spectrometer and TOF data, since TOF results require the extrapolation of

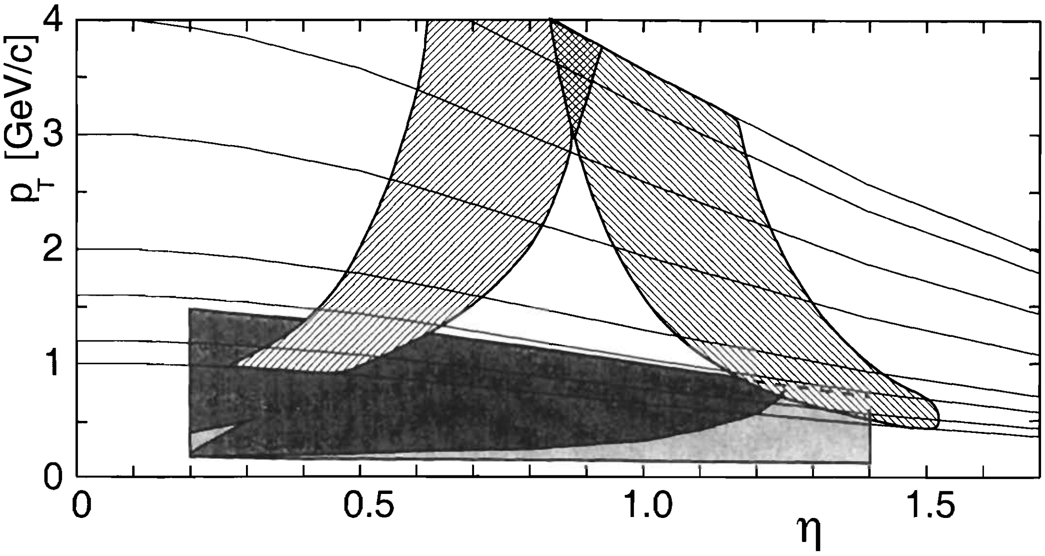

Figure 5-7: Region in transverse momentum p_T and pseudorapidity η that is covered by the Spectrometer (grey) and Time-of-Flight (hatched) particle identification techniques. The two bands for each detector represent the two different bending-directions for charged particles in the magnetic field. The lines on the plot are lines of total momentum, $p = p_T \cosh \eta$.

tracks from the Spectrometer.

There are also dead channels in the Time-of-Flight walls. This is determined to be a $5.1 \pm 2\%$ correction to the TOF particle yields.

5.5 Synthesis of Spectrometer and TOF Data

Figure 5-7 shows the acceptance of the Time-of-Flight and Spectrometer detectors for particle identification in (p_T, η)-space. To provide for an easier comparison of our data to theoretical models or other experiments, we synthesize the results from these different acceptances to generate the p_T spectrum for each species at a constant rapidity.

The Spectrometer data is divided in three bins in pseudorapidity, $0.2 \leq \eta < 0.6$, $0.6 \leq \eta < 1.0$ and $1.0 \leq \eta \leq 1.4$. Particle identification is performed in bins of total momentum p; the rapidity for each species is calculated from the mean p and p_T in each PID bin, and the location of the data-points is plotted in (p_T, y)-space. An example is shown for protons in Figure 5-8. From these plots a common rapidity

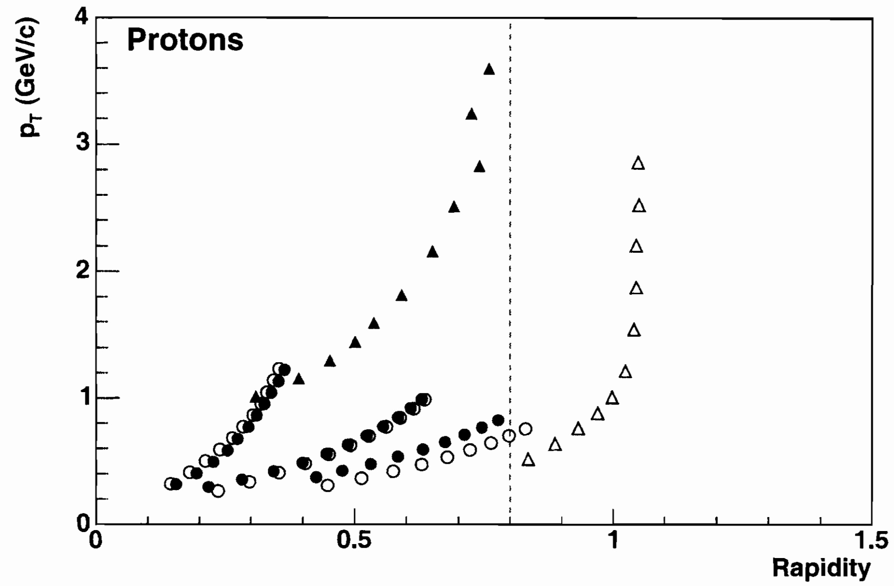

Figure 5-8: Location of PID data-points in (p_T, y)-space, for protons from the most central bin. The red dashed line indicates the chosen common rapidity point at which all invariant particle yields will be evaluated.

value is chosen; the invariant yield at this rapidity will be evaluated for the whole p_T range.

The synthesis procedure begins by choosing data-points which are close in p_T. Then the invariant yield $\frac{1}{2\pi p_T}\frac{d^2N}{dp_T dy}$ is plotted versus rapidity for these points. The simplest assumption is that the rapidity dependence of the invariant yields is constant over this short rapidity interval near mid-rapidity. We therefore take the best constant fit to these points as the value of the invariant yield at this p_T. For comparison, we also fit a straight line to the points and the difference between the constant fit and this straight line evaluated at the common rapidity point is taken as a measure of the systematic error introduced by this assumption. The statistical error on the synthesized invariant yield is the propagated error from the individual points.

This process is illustrated for a single p_T 'slice' in Figure 5-9.

A technical consideration is that though the data-points have similar p_T values, they are not identical, so it is necessary to account for the strong p_T-dependence of

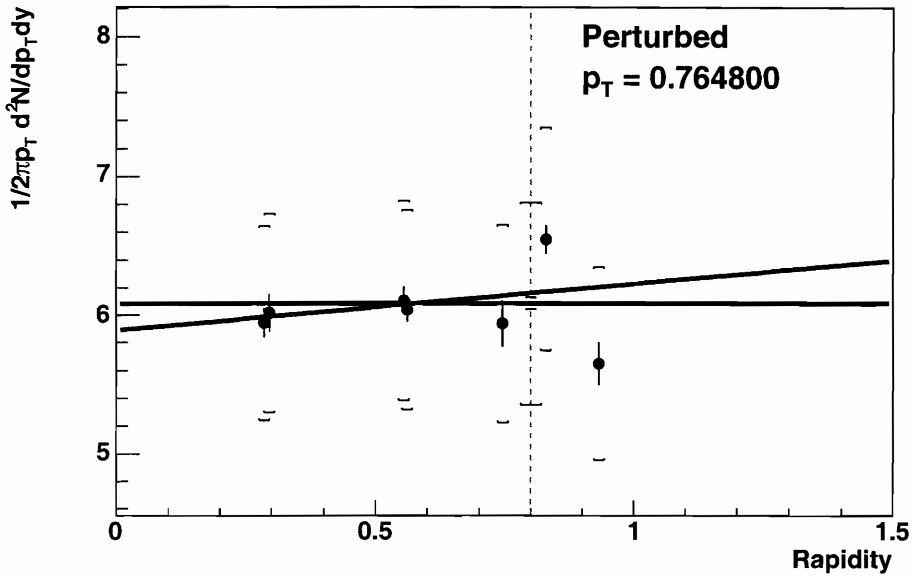

Figure 5-9: Example of the synthesis of TOF and dE/dx data. The invariant yield is plotted versus rapidity for protons with $p_T = 0.765$. The statistical errors on each point are shown, and the brackets represent the systematic errors. The black lines show the best constant and straight-line fits to these points. The red dashed line is the common rapidity value; the brackets on this line represent the total systematic error on the synthesized invariant yield.

the particle yields. This is done using a Taylor expansion:

$$f(p_T + \Delta p_T) \approx f(p_T) + \Delta p_T \frac{df(p_T)}{dp_T} \tag{5.4}$$

All data-points to be combined are 'perturbed' in this way to the same p_T value before synthesizing. The derivative is obtained by fitting a function to the data-points.

5.6 Integrating p_T Spectra

Proton and antiproton d^2N/dp_Tdy transverse momentum distributions can be integrated to obtain the total particle yield dN/dy.

We integrate over the measured data-points and extrapolate over the unmeasured low-p_T region. Because the p_T spectrum falls so sharply, the high-p_T region beyond the measured points makes a negligible contribution to the total yield and is not included. The low-p_T extrapolation uses a simple straight-line from zero to the first data-point and for comparison, a variety of physically-motivated fit functions. An example of a d^2N/dp_Tdy spectrum and a fit to it is shown in Figure 5-10.

Statistical errors on the sum turn out to be negligible. Systematic errors on the total yield come from propagation of the errors on the individual data-points, plus additional uncertainty which arises as a result of the extrapolation over the unmeasured low-p_T region.

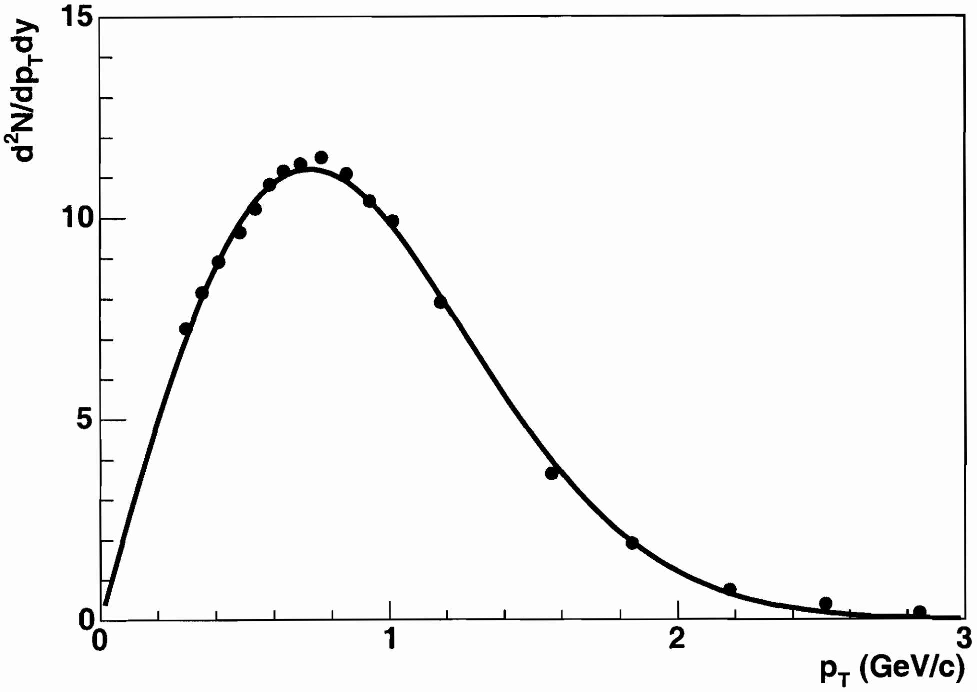

Figure 5-10: Example of d^2N/dp_Tdy distribution for central (0-15%) protons, with a fitted function used for estimating systematic errors.

5.7 Systematic Error Analysis

The estimated systematic uncertainties on all corrections that are applied to the data are listed in Table 5.1. These are convoluted to obtain the overall systematic error

Correction	Uncertainty (%)	Applied To
Acceptance and Efficiency	$5 + 2(\lvert p_T - 2\rvert)$	All
Occupancy	3	All
Ghost	2	All
Feed-down	See Section 5.2.3	Protons
Feed-down	1	Pions
Secondaries	1	Protons
Spec Dead Channels	5	All
TOF Dead Channels	2	TOF only

Table 5.1: Systematic uncertainties on the corrections applied to the data.

arising from corrections to the data. This error is largest for protons, varies slowly with p_T, and is in the range 10-15%.

There is an additional contribution to the systematic uncertainty on the final spectra that arises from the synthesis of the TOF and dE/dx data to a common rapidity point. This contribution is estimated by the procedure described in Section 5.5 and then added in quadrature to the systematic error from the corrections.

When considering quantities derived from the p_T-spectra, such as antiparticle/particle ratios or integrated yields dN/dy, one has to carefully consider which sources of systematic error are correlated, in order to obtain the correct uncertainty on the final derived quantity. To take a simple example, the correction for dead channels affects all spectra equally and the uncertainty on this correction should therefore cancel when considering ratios. This procedure has been done for all results described in the next chapter.

Chapter 6

Results and Conclusions

6.1 Transverse Momentum Distributions

The Lorentz-invariant particle yields $\frac{1}{2\pi p_T}\frac{d^2N}{dp_T dy}$ near mid-rapidity ($y = 0.8$) for pions, kaons, protons and antiprotons from Au+Au collisions at $\sqrt{s_{NN}} = 62.4$ GeV are presented as a function of the transverse momentum (p_T) in Figure 6-1. These data are being prepared for publication at the time of writing, but should still be considered preliminary.

These results have been corrected for the geometrical acceptance of the detector; the efficiency, purity and momentum resolution of track reconstruction; secondary particles and feed-down from weak decays.

The shapes of these p_T-spectra exhibit a clear species-dependence for $p_T < 1.5$ GeV/c, with the pion spectra having a concave form while kaons and protons are convex. When summed, the identified particle spectra are found to agree within systematic errors with the PHOBOS non-identified charged hadron transverse momentum distributions at 62.4 GeV [100].

The centrality bins used in this analysis, with the mean number of participant nucleons (N_{part}) and binary nucleon-nucleon collisions (N_{coll}) for each bin, are listed in Table 6.1.

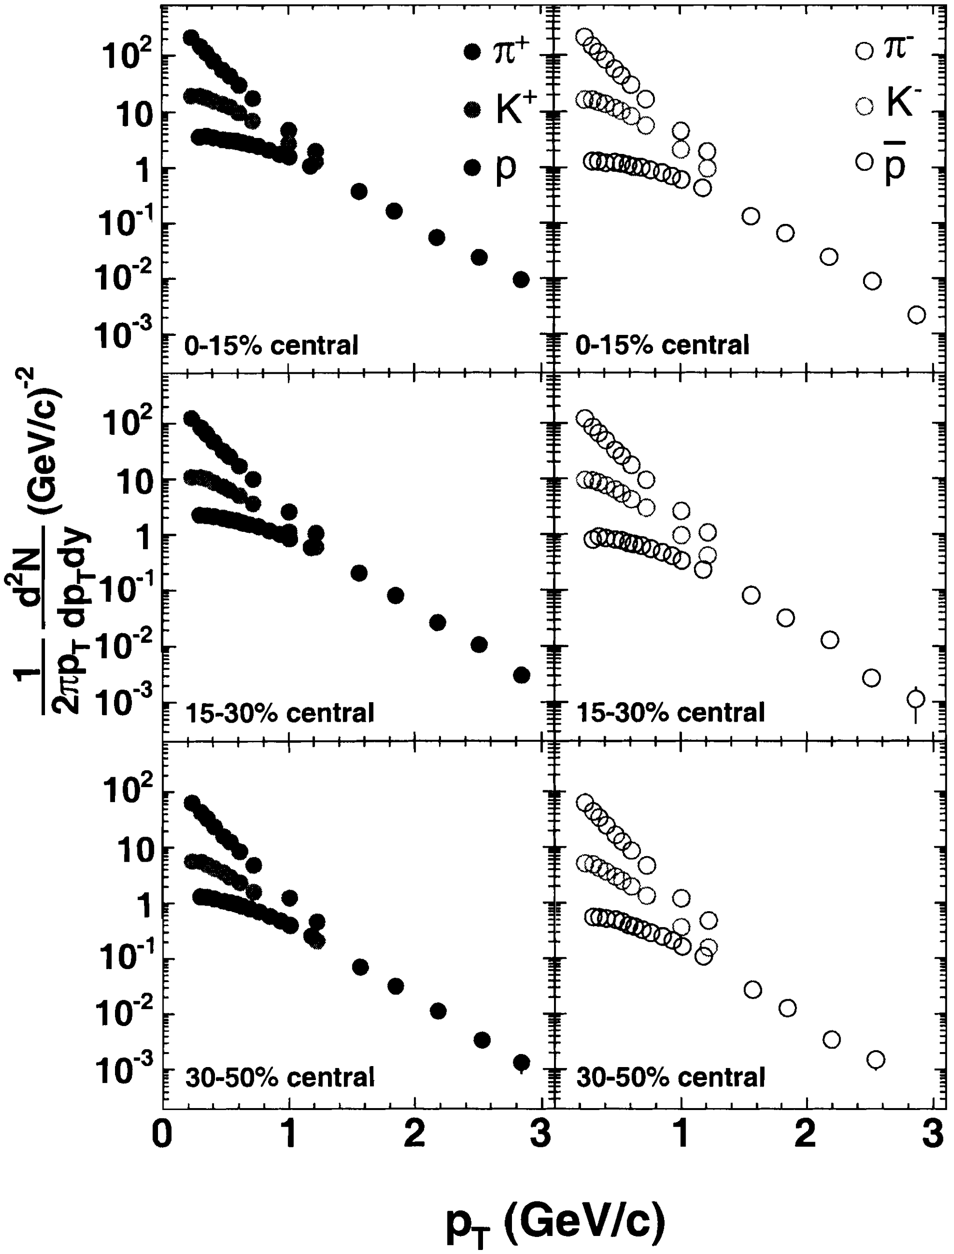

Figure 6-1: Transverse momentum distributions for pions, kaons, protons and antiprotons for all centrality bins used in this analysis. Only statistical errors are shown. Systematic errors vary from 10 to 20% and are generally smaller than the symbol size.

Centrality	$\langle N_{part} \rangle$	$\langle N_{coll} \rangle$
0-15%	294 ± 10	684 ± 23
15-30%	175 ± 9	331 ± 23
30-50%	88 ± 9	128 ± 17

Table 6.1: Centrality bins used in this analysis. The bins are defined in terms of percentages of the total inelastic Au+Au collision cross-section, with the mean number of participating nucleons and binary nucleon-nucleon collisions given for each bin.

6.1.1 Collision-Energy Dependence of Proton and Antiproton Spectra

The dependence on collision energy of the proton and antiproton spectra for central (0-15%) Au+Au collisions can be seen by comparing to data from the PHENIX collaboration at $\sqrt{s_{NN}} = 200$ GeV [102]. This is illustrated in Figure 6-2.

The magnitude of the antiproton yields are seen to increase with the energy and the spectrum becomes 'harder' – that is to say, proportionally more antiprotons are found at higher transverse momentum at the higher collision energy. This agrees with the general trend observed in charged particle p_T spectra for high-energy p+p and A+A collisions.

The transverse momentum distributions of protons, on the other hand, arise from the interplay of baryon transport from the initial colliding nuclei and particle production due to energy liberated in the collision. Qualitatively, as the collision energy increases, one would expect the effects of particle production to increase (as it is seen to do for antiprotons) while those of baryon stopping should decrease. The result is such that, although the total yield of protons increases with collision energy (see Section 6.4), there are actually more protons at low-p_T ($p_T < 1$ GeV/c) for 62.4 GeV compared to 200 GeV Au+Au collisions. Baryon stopping must therefore dominate in this region.

There is also data available on identified particle spectra from Au+Au collisions at 130 GeV [103]. However, these data have not been corrected for feed-down from weak decays and therefore cannot be directly compared to the results of this analysis, for which feed-down corrections have been applied.

Although there is also p+p data available at a collision energy of 63 GeV [104, 105], it is not clear that these data are consistent and to avoid being potentially misleading we do not compare to them in this thesis.

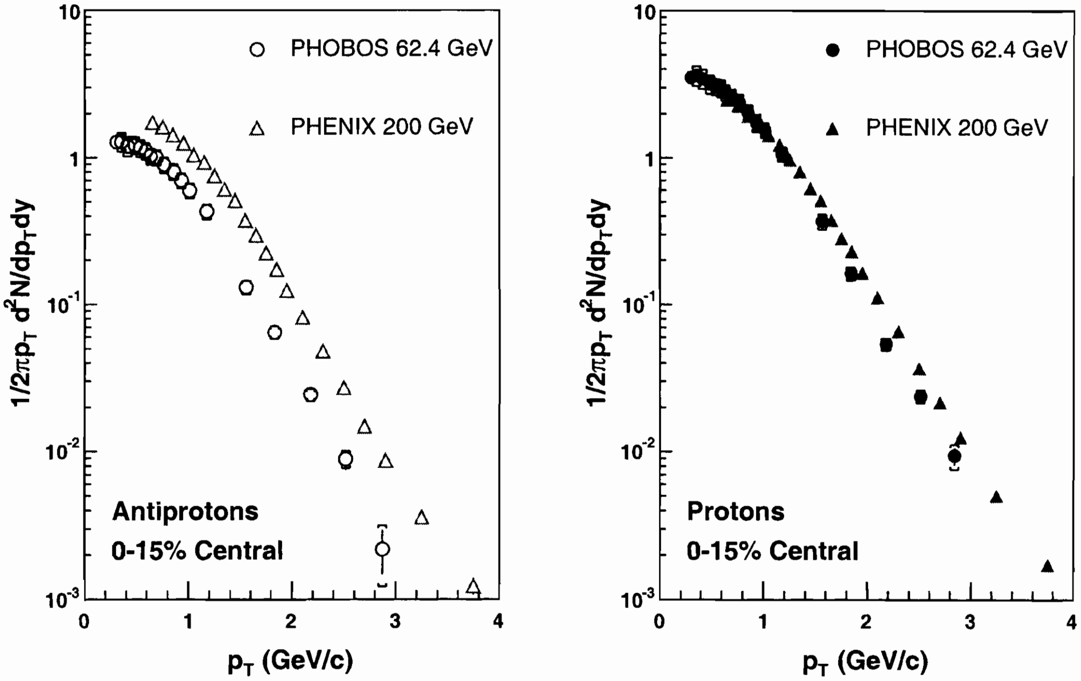

Figure 6-2: Proton and antiproton transverse momentum distributions for the 15% most central Au+Au collisions are compared to data from the PHENIX collaboration at $\sqrt{s_{NN}} = 200$ GeV [102].

6.2 Proton/Hadron Fraction

The contribution of protons and antiprotons to the total charged hadron yields as a function of transverse momentum are displayed for each centrality bin at 62.4 GeV in Figure 6-3.

Over the p_T range measured here, we see that the proton and antiproton contributions increase steadily as a function of transverse momentum until at $p_T \approx 3$ GeV/c in central collisions, protons comprise about half the total charged hadron yield and are therefore the dominant species.

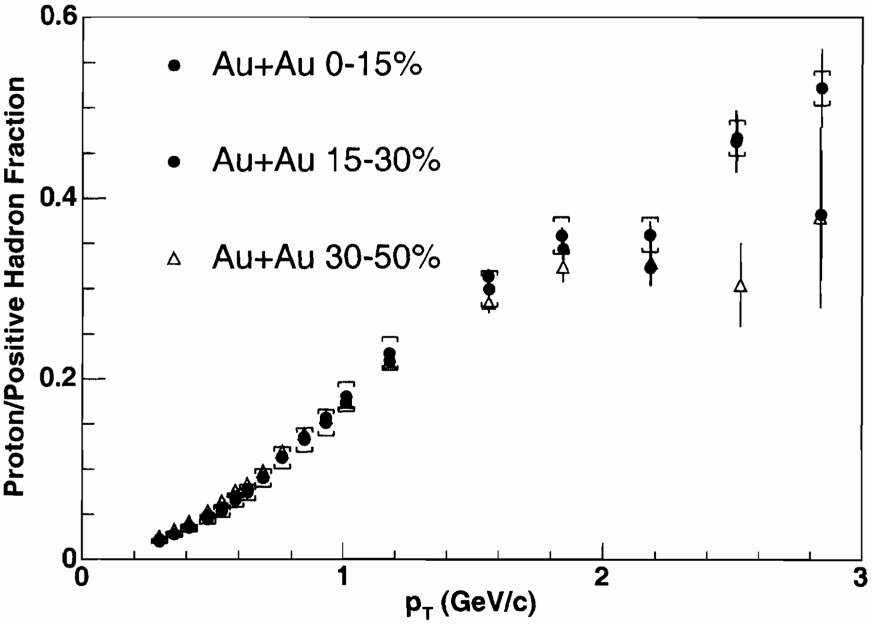

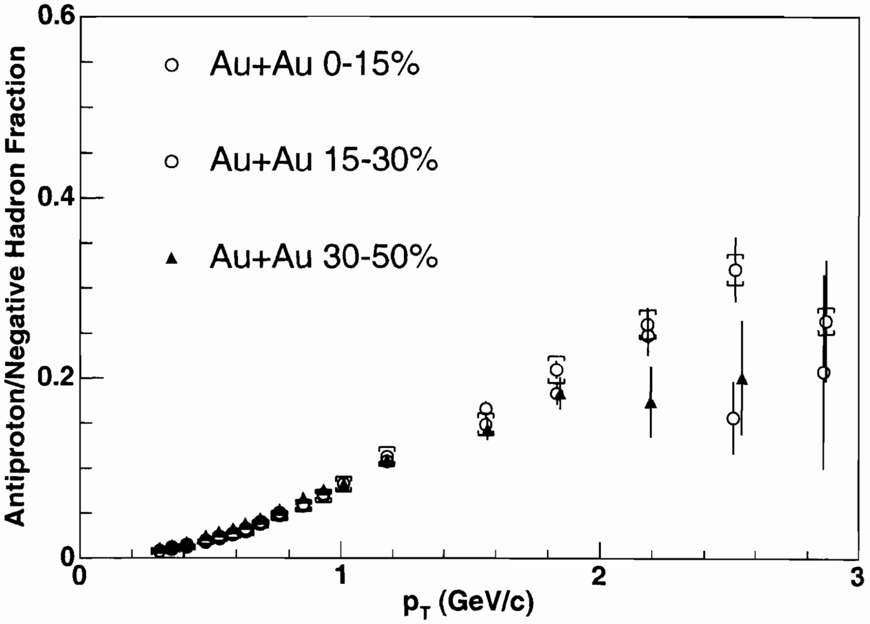

Figure 6-3: Top: Ratio of protons to all positive-charged hadrons as function of transverse momentum, for the three centrality classes used in this analysis. Bottom: Ratio of antiprotons to all negative-charged hadrons. The brackets represent systematic errors – for clarity, these are shown only for the most central data.

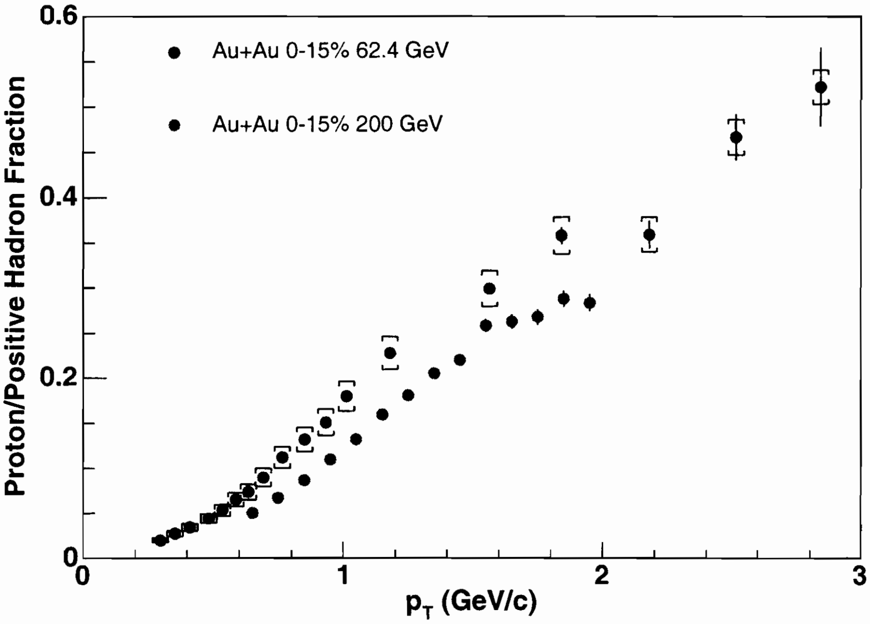

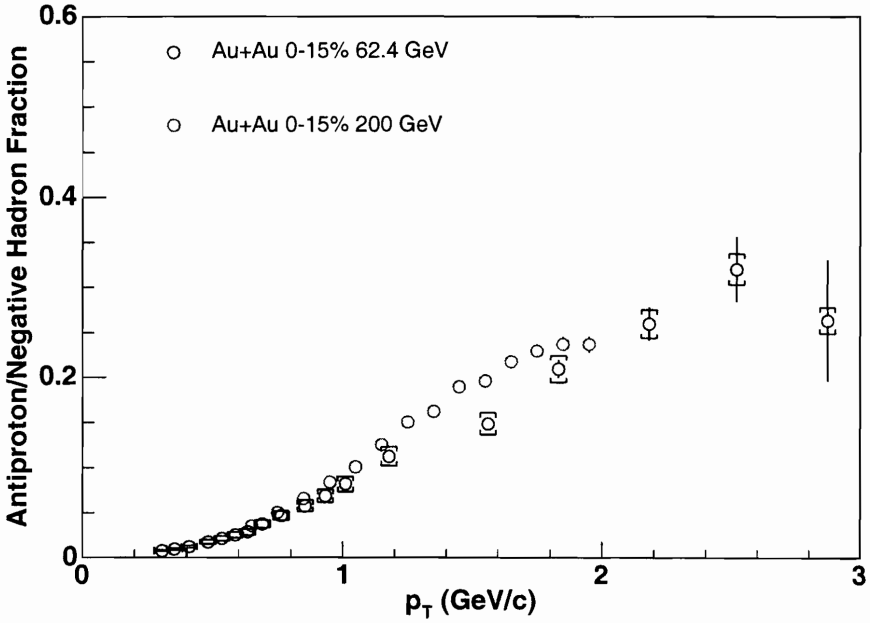

Figure 6-4: Top: Ratio of protons to all positive-charged hadrons as function of transverse momentum, for central (0-15%) Au+Au collisions at 62.4 GeV (this analysis) and 200 GeV (data from the PHENIX collaboration [102]). Bottom: Ratio of antiprotons to all negative-charged hadrons. The brackets represent systematic errors on the 62.4 GeV data.

Results for the most central (0-15%) Au+Au collisions at 62.4 and 200 GeV are compared in Figure 6-4. In the light of the interest generated by the large proton/pion ratio measured at 200 GeV (see Section 1.4.8), it is perhaps surprising to see that the proton/hadron fraction at a given p_T actually *decreases* with collision energy when going from 62.4 to 200 GeV collisions. It appears that, although the yield of protons does increase with collision energy, the contribution from stopped protons is reduced and the result is that other species increase more quickly with energy – so the proton/hadron ratio in fact decreases.

The large relative yield of protons at intermediate p_T is thus seen to be a general feature of heavy-ion collisions in the energy range 62.4 - 200 GeV and not a phenomenon confined only to the highest energies.

6.3 Evolution of p_T Spectra with Collision Centrality

The other aspect of the proton/pion puzzle at 200 GeV was the evolution of their p_T spectra with collision centrality: for $2 < p_T 4$ GeV/c, pions were found to be strongly suppressed relative to N_{coll} scaling, while protons were not.

We investigate the centrality evolution for different species by dividing the central spectrum by the spectrum for peripheral events. Figure 6-5 shows this centrality ratio at 62.4 GeV for protons, antiprotons and mesons (= pions + kaons), further scaled by the relative number of participating nucleons in the two centrality classes. We choose to plot mesons because pions and kaons can only be separately identified over a small p_T range, but their sum can be determined over the same range for which proton identification is possible.

Although statistical errors on this centrality ratio are large at the highest measured p_T, there appears to be a clear difference between the centrality evolution of mesons and (anti)protons. The meson ratio is relatively flat over the measured p_T range and could be described as approximate N_{part}-scaling (to within $\sim 25\%$) at all p_T. The

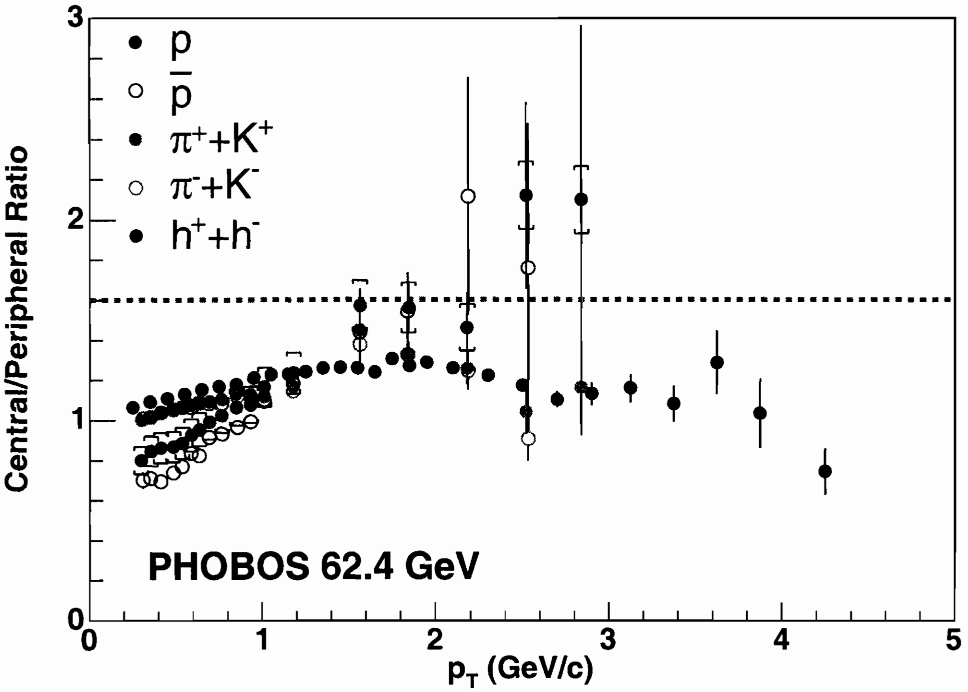

Figure 6-5: Ratio of central (0-15%) to peripheral (30-50%) spectra at 62.4 GeV, scaled by the relative number of participating nucleons, for protons (blue), antiprotons and mesons (magenta). For comparison, the equivalent ratio for non-identified charged hadron spectra from [100] is shown (black), although there the peripheral bin is 35-50%. The dashed red line represents N_{coll}-scaling. Brackets represent the systematic error on the proton ratio.

centrality evolution of protons and antiprotons, on the other hand, varies considerably as a function of transverse momentum, scaling more slowly than the mesons at low-p_T, yet increasing rapidly as p_T rises until it becomes consistent with N_{coll}-scaling at $p_T > 2$ GeV/c.

This behaviour is very similar to that observed at 200 GeV. This is illustrated in Figure 6-6, where the centrality ratios for different species at 200 GeV have been plotted for the same centrality bins as the present analysis, using data from the PHENIX collaboration [102].

We therefore conclude that the essential features of the proton/pion puzzle in 200 GeV Au+Au collisions are also present at 62.4 GeV. This observation will be important for understanding the nature of the processes which govern particle pro-

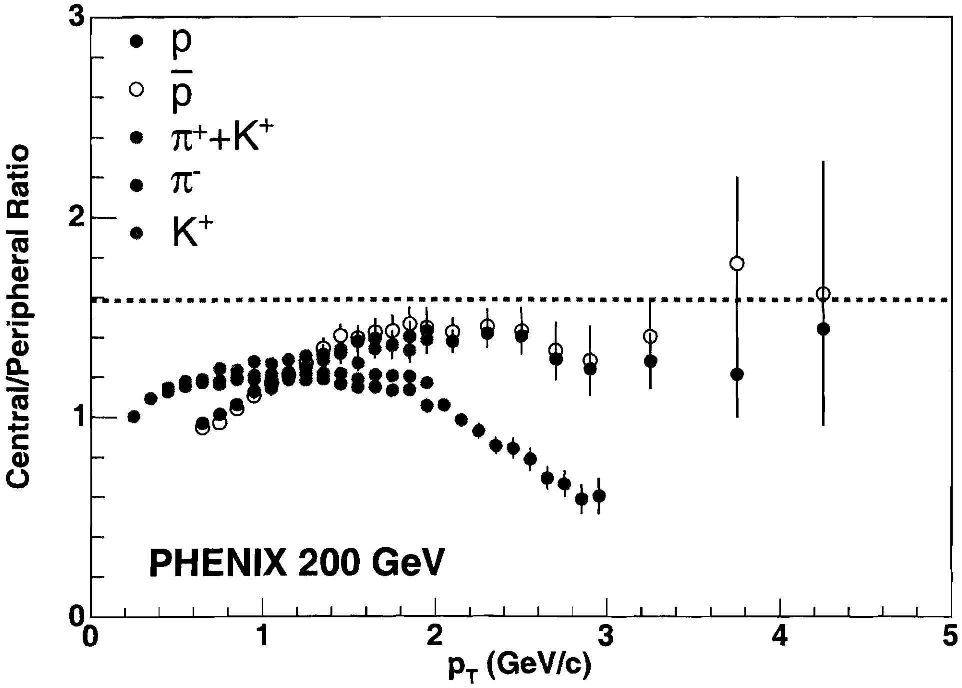

Figure 6-6: Ratio of central (0-15%) to peripheral (30-50%) spectra at 200 GeV, scaled by the relative number of participating nucleons, for protons (blue), antiprotons, mesons (magenta), kaons (green) and pions (red), using data from [102]. The dashed red line represents the expectation of N_{coll}-scaling.

duction at intermediate transverse momentum in heavy-ion collisions.

As described in Section 1.4.8, two scenarios have been proposed to explain the results at 200 GeV. One scenario [76] combines the suppression of high-p_T pions due to jet-quenching in the dense partonic medium with large proton yields that are attributed to a high degree of baryon transport from the initial nuclei, using a mechanism based on gluon junctions [59]. The other scenario proposes a 'recombination/coalescence' model of hadron formation at intermediate p_T which naturally gives rise to significant differences between baryons and mesons based on the number of consituent quarks [77, 78, 79].

Theoretical predictions for 62.4 GeV Au+Au collisions based on these models have not yet been made, so we cannot make any direct comparisons. It is not immediately clear that either mechanism would predict the effect to be so similar over a wide of range of collision energies.

It is worth noting that there is hope that future experimental results may well provide enough information to really distinguish between the two theories mentioned above, because each must simultaneously describe at least one other observable in addition to the identified particle p_T spectra: parton recombination models must also fit the elliptic flow parameter $v_2(p_T)$ for different species, while the model based on baryon junctions must correctly give the net proton $(p - \bar{p})$ yields. The transverse momentum distributions from this analysis can be used to obtain the first measurements of net proton yields near mid-rapidity for Au+Au collisions at 62.4 GeV – we present these results in Section 6.4.

Low-p_T Centrality Scaling of Protons and Antiprotons

As a brief aside, N_{part}-scaling has become such a standard base-line expectation for the centrality-dependence of heavy-ion observables that it is surprising to see that low-p_T protons and antiprotons at 62.4 GeV scale more *slowly* than the number of particpating nucleons. Protons and antiprotons at 200 GeV appear to follow a similar trend in their centrality evolution as at 62.4 GeV, but the 200 GeV spectra do not extend low enough in p_T to see the scaling fall below the N_{part} line. PHOBOS has the capability to measure identified particle yields at extremely low transverse momentum ($p_T \approx 40$ MeV for pions and ≈ 180 MeV for protons) [106]. Thus, future measurements of very low-p_T particle yields as a function of collision centrality at both 62.4 and 200 GeV should be able to further investigate this surprising low-p_T proton centrality evolution.

6.4 Net Proton Yields and Baryon Transport

The transport of baryons from the initial nuclei in a heavy-ion collision (also known as 'baryon stopping') dictates the energy available for particle production and hence determines the subsequent evolution of the system. Baryon stopping can be explored through measurement of the net proton yields $(p - \bar{p})$ as a function of rapidity.

The proton and antiproton transverse momentum spectra presented in this thesis can be integrated (as described in Section 5.6) to obtain the particle yields dN/dy near mid-rapidity. These are the first such results from Au+Au collisions at $\sqrt{s_{NN}}$ =62.4 GeV.

The proton, antiproton and net proton integrated yields for all centrality bins are given in Table 6.2 and plotted versus the number of participating nucleons (N_{part}) in Figures 6-7 & 6-8.

Centrality	$\frac{dN}{dy}(p)$	$\frac{dN}{dy}(\bar{p})$	$\frac{dN}{dy}(p-\bar{p})$
0-15 %	13.4 ± 1.9	5.1 ± 0.7	8.3 ± 1.3
15-35 %	7.7 ± 1.1	3.0 ± 0.4	4.6 ± 0.7
35-50 %	3.8 ± 0.5	1.6 ± 0.2	2.2 ± 0.3

Table 6.2: Proton, antiproton and net proton ($p-\bar{p}$) integrated yields. The errors quoted are systematic – statistical errors are negligible. Correlations in the systematic errors for protons and antiprotons are included in the systematic error on the net proton yield.

The observed linearity of the mid-rapidity net proton yields with the number of participating nucleons is interesting. Although the total net protons integrated over all rapidity is certainly expected to be proportional to N_{part}, it is not necessary for the yield at mid-rapidity to have this behaviour – one could imagine physical scenarios where the degree of stopping increases significantly as the collision impact parameter decreases, as a result of multiple collisions per participating nucleon, such that the net proton yield at mid-rapidity scales faster than N_{part}. That this does not appear to happen at mid-rapidity for high-energy heavy-ion collisions should be a constraint on models of baryon transport.

The mid-rapidity net-proton yields for 0-5% central heavy-ion collisions over a range of energies are plotted in Figure 6-9 as a function of the beam rapidity. The PHOBOS point at $\sqrt{s_{NN}}$ =62.4 GeV ($y = 4.2$) is obtained by extrapolating the net proton results from Figure 6-8 to $N_{part} \approx 350$. That the NA49 data is from Pb+Pb collisions while all the rest are Au+Au is estimated to be only about a 5% effect.

Our result fills a large gap between SPS collisions at $\sqrt{s_{NN}} = 17.2$ GeV ($y = 2.9$) and the higher RHIC energies of 130 & 200 GeV ($y = 4.9, 5.4$), and therefore provides

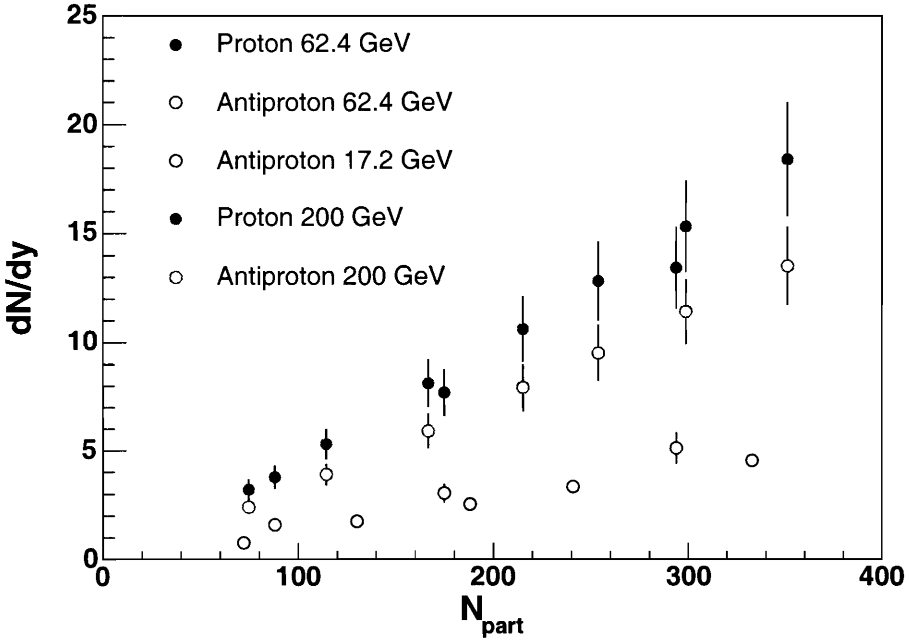

Figure 6-7: Proton and antiproton yields dN/dy near mid-rapidity as a function of the number of participating nucleons N_{part} for Au+Au collisions at 62.4 (this analysis) and 200 GeV ([102]), and Pb+Pb collisions at 17.2 GeV ([107]). Error bars represent systematic errors; statistical errors are negligible.

very useful data for comparing to models of baryon transport mechanisms.

In the near future, the Large Hadron Collider (LHC) at CERN will collide Pb+Pb at $\sqrt{s_{NN}} = 5.5$ TeV ($y = 8.7$). One can see from imagining an extension of Figure 6-9 that the net proton yield should be essentially zero at those collision energies. This is important because lattice QCD calculations encounter technical difficulties at non-zero baryon number – such difficulties should vanish at LHC energies and there is hope that the field of heavy-ion physics will then enter a period where detailed experimental results can be compared directly to first-principles QCD calculations.

6.5 Summary

Pion, kaon, proton and antiproton transverse momentum distributions at $y = 0.8$ have been presented for Au+Au collisions at $\sqrt{s_{NN}} = 62.4$ GeV. These are the first

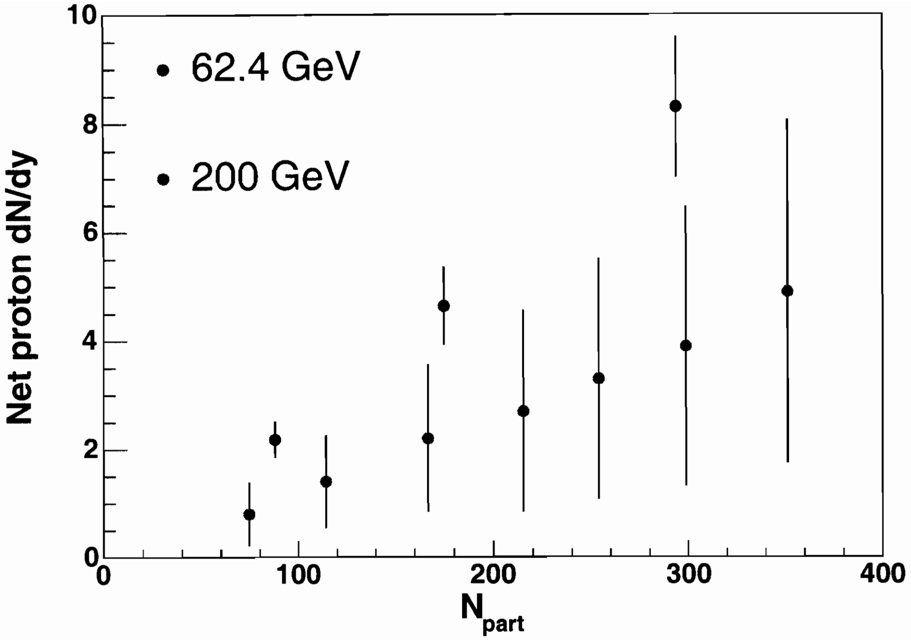

Figure 6-8: Net proton ($p-\bar{p}$) yields dN/dy near mid-rapidity as a function of the number of participating nucleons N_{part} for Au+Au collisions at 62.4 and 200 GeV. Data at 200 GeV is from the PHENIX collaboration [102]. Error bars represent systematic errors; statistical errors are negligible.

identified particle p_T spectra at this collision energy.

The contribution of protons to the total yield of charged hadrons is found to increase rapidly with p_T until at $p_T \sim 3$ GeV/c in central collisions they are the dominant species. With respect to collision centrality, in the intermediate p_T range, proton yields are consistent with N_{coll} scaling while mesons are suppressed.

This is very similar to the behaviour seen in 200 GeV collisions, which generated significant interest and prompted new physical models to try to explain it. This new observation that the proton/pion behaviour is similar at 62.4 and 200 GeV will therefore be important in constraining these proposed models and helping to understand the processes which govern particle production at intermediate p_T in the complex system formed in relativistic heavy-ion collisions.

The proton and antiproton p_T spectra are integrated and used to obtain the net proton ($p-\bar{p}$) yields near mid-rapidity in 62.4 GeV Au+Au collisions. This data

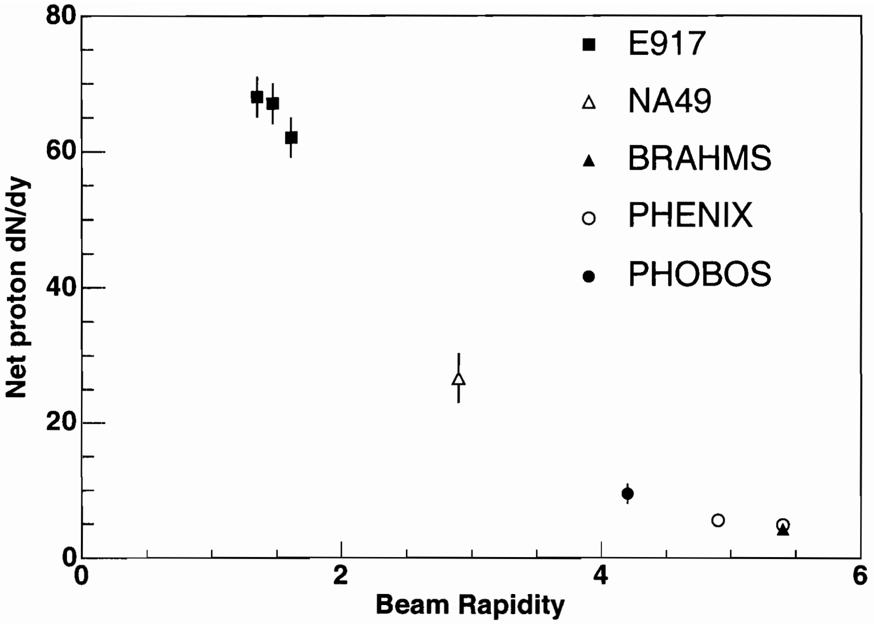

Figure 6-9: Mid-rapidity net proton dN/dy as a function of the beam rapidity for central heavy-ion collisions. Data are taken from [108, 109, 101, 102, 58].

will be useful for investigating the mechanisms for baryon transport from the initial nuclei in heavy-ion collisions.

Some other results which can be obtained from the identified particle transverse momentum spectra in Au+Au collisions at 62.4 GeV are presented in the Appendices.

Appendix A

Antiproton-to-Proton Ratio and the Baryochemical Potential

The antiproton to proton and K^-/K^+ ratios are shown as a function of p_T in Figure A-1. As at other RHIC energies of 130 and 200 GeV [97, 54], these ratios are found, within errors, to have no p_T-dependence.

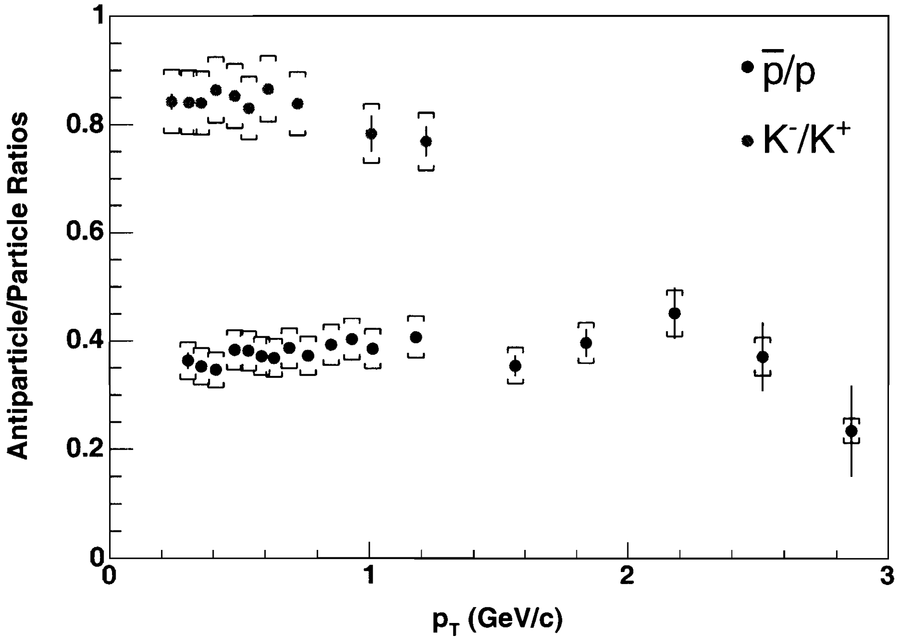

Figure A-1: Proton and kaon antiparticle/particle ratios as a function of transverse momentum for 0-15% central Au+Au collisions at 62.4 GeV.

By integrating the tranvserse momentum spectra over the measured p_T range, the antiparticle/particle ratios for 0-15% central Au+Au collisions at 62.3 GeV are found to be:

$$\frac{\bar{p}}{p} = 0.38 \pm 0.03 \text{ (sys.)}$$

$$\frac{K^-}{K^+} = 0.83 \pm 0.03 \text{ (sys.)}$$

Statistical errors are negligible. Antiparticle/particle ratios are plotted as a function of collision energy in Figure A-2; these results at 62.4 GeV are seen to fit smoothly into the energy evolution of these quantities.

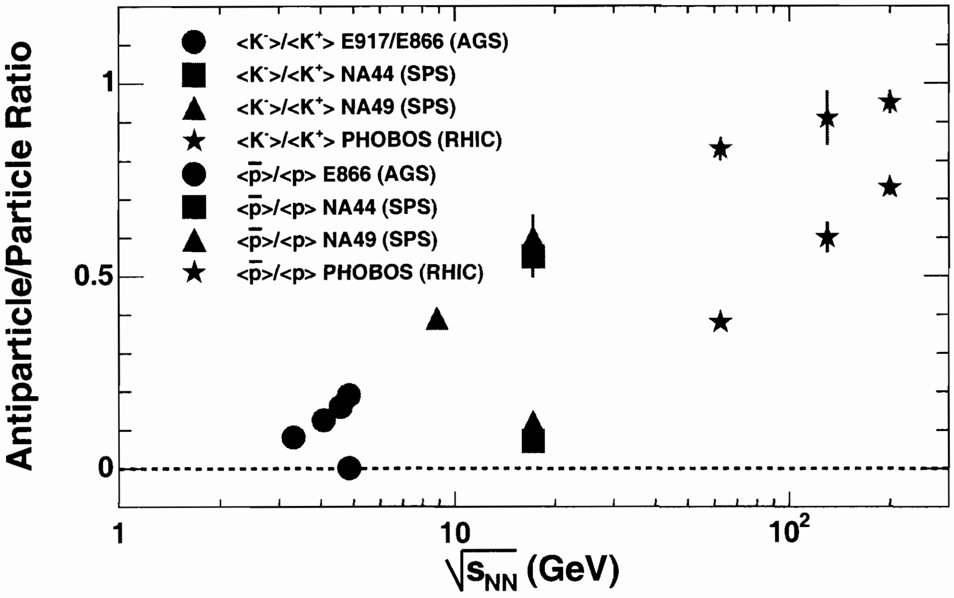

Figure A-2: Proton and kaon antiparticle/particle ratios in heavy-ion collisions as a function of collision energy. The compilation of data is taken from [94] and references therein.

The measurement of the antiproton-to-proton ratio allows a simple estimate of the baryochemical potential μ_B. Starting from the statistical distribution of particle yields in thermal equilibrium, given in Equation 1.12, and making the approximation that the temperature is large enough that quantum effects are negligible, then, since protons and antiprotons have equal mass but opposite quantum numbers, one obtains:

$$\frac{\bar{p}}{p} \approx e^{-2\mu_B/T} \Rightarrow \frac{\mu_B}{T} \approx -0.5 \ln\left(\frac{\bar{p}}{p}\right) \tag{A.1}$$

It has been shown in [94] that this simple estimate agrees to within 5% with the value obtained using the more sophisticated thermal model calculations of [55].

To obtain μ_B, a value for the temperature T has to be assumed. For consistency with the PHOBOS estimates of μ_B for central Au+Au collisions at 130 and 200 GeV [97, 54, 94], we choose a value of T = 165 MeV. The values of μ_B obtained for all three energies are displayed in Table A.1.

$\sqrt{s_{NN}}$ (GeV)	$\bar{p}/p$	μ_B (MeV)
62.4	0.38 ± 0 (stat.) ± 0.03 (sys.)	80 ± 0 (stat.)
130	0.60 ± 0.04 (stat.) ± 0.06 (sys.)	45 ± 5 (stat.)
200	0.73 ± 0.02 (stat.) ± 0.03 (sys.)	27 ± 2 (stat.)

Table A.1: Table of $\bar{p}/p$ and μ_B at different collision energies. A temperature of T = 165 MeV is assumed. The 62.4 GeV result uses the simple approximation of Equation A.1 to obtain μ_B while the 130 and 200 GeV values are taken from [94] and are the results of a more sophisticated thermal model calculation.

Appendix B

Transverse Mass Distributions

Motivated by the Color Glass Condensate description of heavy-ion collisions, the authors of [110] suggest that the transverse distributions of all particle species should follow a universal relationship in terms of the *transverse mass*, $m_T = \sqrt{m_0^2 + p_T^2}$, known as *$m_T$-scaling*.

Preliminary results from p+p [111] and d+Au [75] collisions at $\sqrt{s_{NN}} = 200$ GeV do indicate an approximate scaling of particle yields with m_T (although the kaon spectra are lower than the rest by about a factor of two, attributed to strangeness suppresion). However, measurements of particle yields extending down to very low transverse momentum show a clear violation of m_T-scaling in central Au+Au collisions at the same energy [106]. The usual interpretation of this is that the system undergoes a tranvserse expansion after colliding, and the resulting radial velocity field gives mass-dependent contributions to the transverse momentum spectra, breaking any universal scaling with m_T that may have existed.

The particle yields for all species at 62.4 GeV are shown as a function of the transverse mass in Figure B-1. Although it would be preferable to have pion/kaon identification out to higher transverse momentum in order to be more conclusive, these spectra show no indication of m_T-scaling in Au+Au collisions at this energy either.

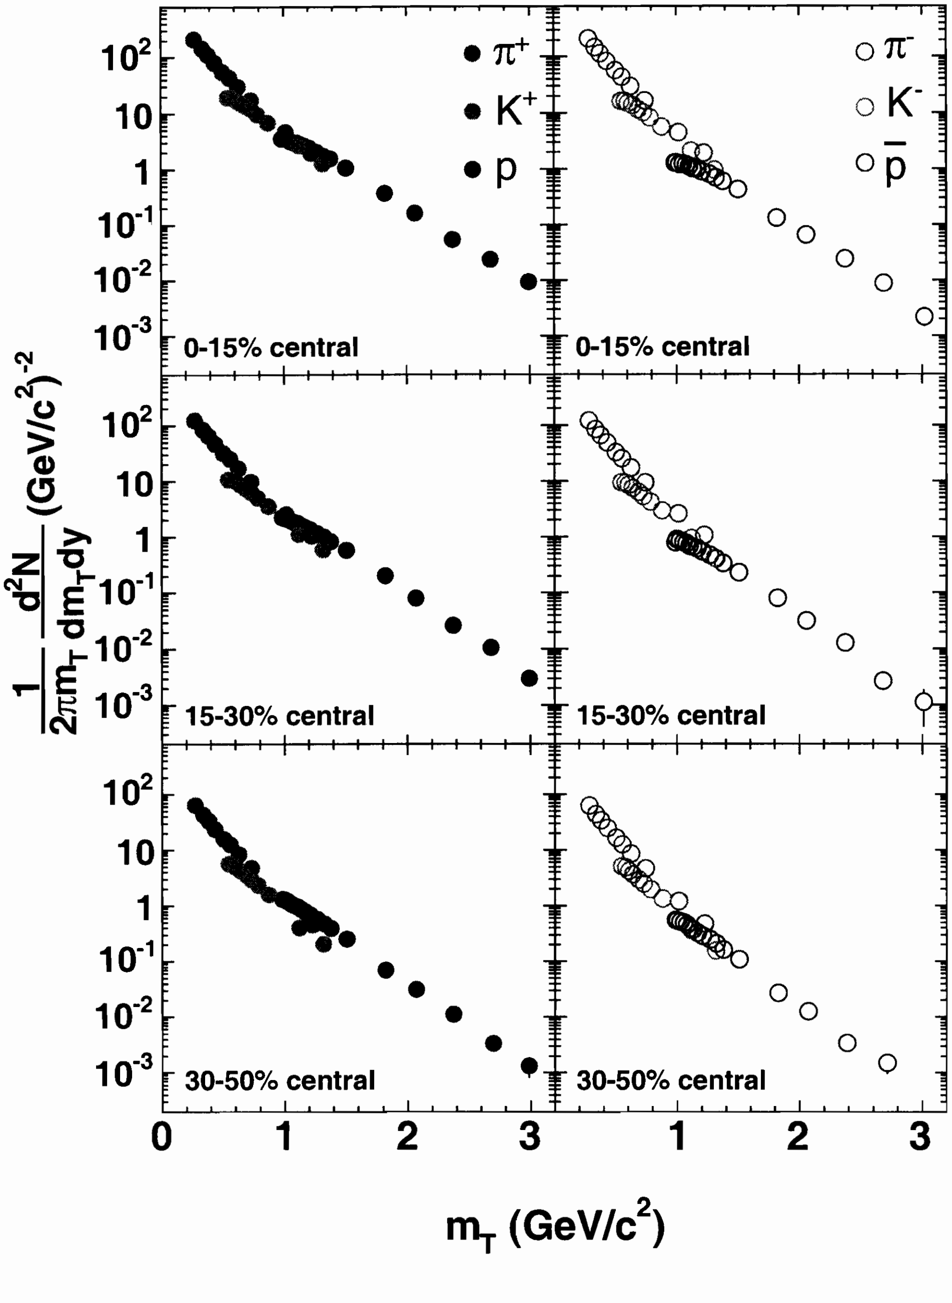

Figure B-1: Transverse mass ($m_T = \sqrt{m_0^2 + p_T^2}$) distributions for pions, kaons, protons and antiprotons from Au+Au collisions at $\sqrt{s_{NN}} = 62.4$ GeV.

Bibliography

[1] M. Gell-Mann and Y. Ne'eman. *The Eightfold Way.* Benjamin, 1964. This book contains reprints of the original articles on the classification of high-energy particles.

[2] V. E. Barnes et al. Observation of a hyperon with strangeness minus three. *Phys. Rev. Lett.*, 12:204, 1964.

[3] Murray Gell-Mann. A schematic model of baryons and mesons. *Phys. Lett.*, 8:214, 1964.

[4] G. Zweig. Quarks. CERN Report, 1964. Unpublished.

[5] J. J. Aubert et al. Experimental observation of a heavy particle J. *Phys. Rev. Lett.*, 33:1404, 1974.

[6] J.-E. Augustin et al. Discovery of a narrow resonance in e^+e^- annihilation. *Phys. Rev. Lett.*, 33:1406, 1974.

[7] S. W. Herb et al. Observation of a dimuon resonance at 9.5 GeV in 400-GeV proton-nucleus collisions. *Phys. Rev. Lett.*, 39:252, 1977.

[8] F. Abe et al. Observation of top quark production in anti-p p collisions. *Phys. Rev. Lett.*, 74:2626–2631, 1995.

[9] S. Abachi et al. Observation of the top quark. *Phys. Rev. Lett.*, 74:2632–2637, 1995.

[10] K. Hagiwara et al. Review of particle physics. *Physical Review D*, 66:010001, 2002.

[11] J. Friedman and H. Kendall. Deep inelastic electron scattering. *Ann. Rev. Nucl. Science*, 22:203, 1972.

[12] J. D. Bjorken. Asymptotic sum rules at infinite momentum. *Phys. Rev.*, 179:1547, 1969.

[13] D. B. Lichtenberg. *Unitary Symmetry and Elementary Particles.* Academic Press, 1970. This work is credited with introducing the term 'color' to describe the quantum number of the strong interaction.

[14] David J. Gross and Frank Wilczek. Ultraviolet behavior of non-abelian gauge theories. *Phys. Rev. Lett.*, 30:1343, 1973.

[15] H. David Politzer. Reliable perturbative results for strong interactions? *Phys. Rev. Lett.*, 30:1346, 1973.

[16] S. Bethke. Determination of the QCD coupling α_s. *J. Phys. G*, 26:R27, 2000.

[17] Kenneth G. Wilson. Confinement of quarks. *Phys. Rev. D*, 10:2445, 1974.

[18] John B. Kogut. Lattice gauge theory approach to quantum chromodynamics. *Rev. Mod. Phys.*, 55:775, 1983.

[19] Jeffrey Goldstone, Abdus Salam, and Steven Weinberg. Broken symmetries. *Phys. Rev.*, 127:965, 1962.

[20] Krishna Rajagopal. Mapping the QCD phase diagram. *Nucl. Phys.*, A661:150–161, 1999, hep-ph/9908360.

[21] W. Trautmann. Multifragmentation and the liquid-gas phase transition: an experimental overview. 2004, nucl-ex/0411023. To appear in the proceedings of International Nuclear Physics Conferences, INPC 2004.

[22] J. Pochodzalla. The nuclear liquid-gas phase transition: Present status and future perspectives. 1996, nucl-ex/9607004. In 'Critical phenomena and collective observables', Acicastello, 1996, pages 1-22.

[23] Robert L. Jaffe and Frank Wilczek. Diquarks and exotic spectroscopy. *Phys. Rev. Lett.*, 91:232003, 2003.

[24] Frank Wilczek. Diquarks as inspiration and as objects. 2004, hep-ph/0409168.

[25] Krishna Rajagopal and Frank Wilczek. The condensed matter physics of QCD. 2000, hep-ph/0011333. Chapter 35 in the Festschrift in honor of B. L. Ioffe, 'At the Frontier of Particle Physics / Handbook of QCD', edited by M. Shifman, World Scientific, 2001.

[26] Mark G. Alford, Krishna Rajagopal, and Frank Wilczek. Color-flavor locking and chiral symmetry breaking in high density QCD. *Nucl. Phys.*, B537:443–458, 1999.

[27] Jeffrey A. Bowers and Krishna Rajagopal. The crystallography of color superconductivity. *Phys. Rev.*, D66:065002, 2002.

[28] Jeffrey Allan Bowers. *Color superconducting phases of cold dense quark matter*. PhD thesis, Massachusetts Institute of Technology, 2003, hep-ph/0305301.

[29] Mark Alford, Chris Kouvaris, and Krishna Rajagopal. Gapless color-flavor-locked quark matter. *Phys. Rev. Lett.*, 92:222001, 2004.

[30] Sanjay Reddy. Neutron stars, supernova and phases of dense quark matter. *J. Phys. G: Nucl. Part. Phys.*, 30:S879, 2004.

[31] Z. Fodor and S. D. Katz. Critical point of QCD at finite T and μ, lattice results for physical quark masses. *JHEP*, 04:050, 2004, hep-lat/0402006.

[32] Frithjof Karsch. Lattice results on QCD thermodynamics. *Nucl.Phys.*, A698:199–208, 2002, hep-ph/0103314.

[33] Robert D. Pisarski and Frank Wilczek. Remarks on the chiral phase transition in chromodynamics. *Phys. Rev. D*, 29:338, 1984.

[34] Ágnes Mócsy, Francesco Sannino, and Kimmo Tuominen. Deconfinement and chiral symmetry restoration. *J. Phys. G: Nucl. Part. Phys.*, 30:S1255, 2004.

[35] Kenji Fukushima. Relation between colour deconfinement and chiral restoration. *J. Phys. G: Nucl. Part. Phys.*, 30:S1263, 2004.

[36] Krishna Rajagopal. The chiral phase transition in QCD: Critical phenomena and long wavelength pion oscillations. 1995, hep-ph/9504310. In 'Quark-Gluon Plasma 2', edited by R. Hwa, World Scientific, 1995.

[37] Christof Roland. Event-by-event fluctuations of particle ratios in central Pb+Pb collisions at 20-158 AGev. *J. Phys. G: Nucl. Part. Phys.*, 30:S1381, 2004.

[38] H G Ritter and X-N Wang, editors. *Quark Matter 2004: Proceedings of the 17th International Conference on Ultra-Relativistic Nucleus-Nucleus Collisions*, 2004. Published as Journal of Physics G: Nuclear and Particle Physics, volume 30.

[39] H Gutbrod, J Aichelin, and K Werner, editors. *Quark Matter 2002: Proceedings of the 16th International Conference on Ultra-Relativistic Nucleus-Nucleus Collisions*, 2003. Published as Nuclear Physics A, volume 715.

[40] T J Hallman, D E Kharzeev, J T Mitchell, and T Ullrich, editors. *Quark Matter 2001: Proceedings of the 15th International Conference on Ultra-Relativistic Nucleus-Nucleus Collisions*, 2002. Published as Nuclear Physics A, volume 698.

[41] B.B. Back et al. The PHOBOS perspective on discoveries at RHIC. 2004, nucl-ex/0410022. Submitted to Nucl. Phys. A.

[42] J Adams et al. Experimental and theoretical challenges in the search for the quark gluon plasma: The STAR collaboration's critical assessment of the evidence from RHIC collisions. 2005, nucl-ex/0501009. Submitted to Nucl. Phys. A.

[43] K Adcox et al. Formation of dense partonic matter in relativistic nucleus-nucleus collisions at RHIC: Experimental evaluation by the PHENIX collaboration. 2004, nucl-ex/0410003. Submitted to Nucl. Phys. A.

[44] I. Arsene et al. Quark-gluon plasma and the color glass condensate at RHIC? the perspective from the BRAHMS experiment. 2004, nucl-ex/0410020. Submitted to Nucl. Phys. A.

[45] R. J. Glauber and G. Matthiae. High-energy scattering of protons by nuclei. *Nucl. Phys. B*, 21:135, 1970.

[46] W. Czyz and L. C. Maximon. High energy, small angle elastic scattering of strongly interacting composite particles. *Annals Phys.*, 52:59, 1969.

[47] M. Gyulassy and X-N Wang. HIJING: A monte carlo model for multiple jet production in pp, pA, and AA collisions. *Phys. Rev. D*, 44:3501, 1991.

[48] A. Bialas, B. Bleszynski, and W Czyz. Multiplicity distributions in nucleus-nucleus collisions at high energies. *Nucl. Phys. B*, 111:461, 1976.

[49] S.S. Adler et al. Centrality dependence of direct photon production in $\sqrt{s_{NN}} =$ 200 GeV Au+Au collisions. *Phys. Rev. Lett.*, 94, 2005.

[50] B.B. Back et al. Significance of the fragmentation region in ultrarelativistic heavy ion collisions. *Phys. Rev. Lett.*, 91:052303, 2003.

[51] J. D. Bjorken. Highly relativistic nucleus-nucleus collisions: the central rapidity region. *Phys. Rev. D*, 27:140, 1983.

[52] B.B. Back et al. Centrality and pseudorapidity dependence of elliptic flow for charged hadrons in Au+Au collisions at $\sqrt{s_{NN}}$= 200 GeV. 2004, nucl-ex/0407012. Submitted to Phys. Rev. C (Rapid Comm.).

[53] P. Braun-Munzinger, K. Redlich, and J. Stachel. Particle production in heavy-ion collisions. 2003, nucl-th/0304013. To appear in Quark Gluon Plasma 3, eds. R. C. Hwa and Xin-Nian Wang, World Scientific.

[54] B.B. Back et al. Ratios of charged particles to antiparticles near mid-rapidity in Au+Au collisions at $\sqrt{s_{NN}}$= 200 GeV. *Phys. Rev. C*, 67:021901(R), 2003.

[55] F. Becattini, J. Cleymans, A. Keranen, E. Suhonen, and K. Redlich. Features of particle multiplicities and strangeness production in central heavy-ion collisions between 1.7A and 158A GeV/c. *Phys. Rev. C*, 64:024901, 2001.

[56] F. Becattini. A thermodynamical approach to hadron production in e^+e^- collisions. *Z. Phys. C*, 69:485, 1996.

[57] W. Busza and A.S. Goldhaber. Nuclear stopping power. *Phys. Lett. B*, 139:235, 1984.

[58] I.G. Bearden et al. Nuclear stopping in Au+Au collisions at $\sqrt{s_{NN}}$= 200 GeV. *Phys. Rev. Lett.*, 93:102301, 2004.

[59] D. Kharzeev. Can gluons trace baryon number? *Phys. Lett. B*, 378:238, 1996.

[60] K. Adcox et al. Suppression of hadrons with large transverse momentum in central Au+Au collisions at $\sqrt{s} = 130$ GeV. *Phys. Rev. Lett.*, 88:022301, 2002.

[61] K. Adcox et al. Centrality dependence of the high p_T charged hadron suppression in Au+Au collisions at $\sqrt{s_{NN}} = 130$ GeV. *Phys. Lett. B*, 561:82, 2003.

[62] S. S. Adler et al. High-p_T charged hadron suppression in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV. *Phys. Rev. C*, 69:034910, 2004.

[63] C. Adler et al. Centrality dependence of high p_T hadron suppression in Au+Au collisions at $\sqrt{s_{NN}}$= 130 GeV. *Phys. Rev. Lett.*, 89:202301, 2002.

[64] J. Adams et al. Transverse momentum and collision energy dependence of high p_T hadron suppression in Au+Au collisions at ultrarelativistic energies. *Phys. Rev. Lett.*, 91:172302, 2003.

[65] B.B. Back et al. Charged hadron transverse momentum distributions in Au+Au collisions at $\sqrt{s_{NN}}$= 200 GeV. *Phys. Lett. B*, 578:297, 2004.

[66] Miklos Gyulassy, Ivan Vitev, Xin-Nian Wang, and Ben-Wei Zhang. Jet quenching and radiative energy loss in dense nuclear matter. nucl-th/0302077. Review for: Quark Gluon Plasma 3, Editors: R.C. Hwa and X.-N. Wang, World Scientific, Singapore.

[67] D. Kharzeev, E. Levin, and L. McLerran. Parton saturation and N_{part} scaling of semi-hard processes in QCD. *Phys. Lett. B*, 561:93, 2003.

[68] D. Kharzeev and E. Levin. Manifestations of high density QCD in the first RHIC data. *Phys. Lett. B*, 523:79, 2001.

[69] B.B. Back et al. Centrality dependence of the charged hadron transverse momentum spectra in d+Au collisions at $\sqrt{s_{NN}}$= 200 GeV. *Phys. Rev. Lett.*, 91:072302, 2003.

[70] S. S. Adler et al. Absence of suppression in particle production at large transverse momentum in $\sqrt{s_{NN}}$ = 200 GeV d+Au collisions. *Phys. Rev. Lett.*, 91:072303, 2003.

[71] J. Adams et al. Evidence from d+Au measurements for final state suppresion of high-p_T hadrons in Au+Au collisions at RHIC. *Phys. Rev. Lett.*, 91:072304, 2003.

[72] I. Arsene et al. Transverse-momentum spectra in Au+Au and d+Au collisions at $\sqrt{s_{NN}}$ = 200 GeV and the pseudorapidity dependence of high-p_T suppression. *Phys. Rev. Lett.*, 91:072305, 2003.

[73] B.B. Back et al. Pseudorapidity dependence of charged hadron transverse momentum spectra in d+Au collisions at $\sqrt{s_{NN}}$= 200 GeV. *Phys. Rev. C*, 70:061901(R), 2004.

[74] S. S. Adler et al. Scaling properties of proton and anti-proton production in $\sqrt{s_{NN}}$= 200 GeV Au+Au collisions. *Phys. Rev. Lett.*, 91:172301, 2003.

[75] Gabor Veres. Identified hadron spectra from PHOBOS. *J. Phys. G: Nucl. Part. Phys.*, 30, 2004.

[76] I. Vitev and M. Gyulassy. Jet quenching and the $\bar{p} > \pi^-$ anomaly in heavy-ion collisions at relativistic energies. *Phys. Rev. C*, 65:041902(R), 2002.

[77] R. Fries, B. Muller, C. Nonaka, and S.A. Bass. Hadron production in heavy ion collisions: Fragmentation and recombination from a dense parton phase. *Phys. Rev. C*, 68:044902, 2003.

[78] R. Fries. Recombination models. *J. Phys. G*, 30:S853–S860, 2004.

[79] R. C. Hwa and C. B.Yang. Fragmentation or recombination at high p_T? *J. Phys. G*, 30:S1117, 2004.

[80] M. Harrison, T. Ludlam, and S. Ozaki. RHIC project overview. *Nucl. Instr. Meth.*, A499:235–244, 2003. This NIM issue is devoted to the RHIC facility.

[81] I. Alekseev et al. Polarized proton collider at RHIC. *Nucl. Instr. Meth.*, A499:392–414, 2003.

[82] B.B. Back et al. The PHOBOS detector at RHIC. *Nucl. Instr. Meth.*, A499:603–623, 2003.

[83] Heinz Pernegger. Layout and tests of silicon pad detectors for the PHOBOS experiment at RHIC. *Nucl. Instr. Meth.*, A419:549–555, 1998.

[84] R. Nouicer et al. Silicon pad detectors for the PHOBOS experiment at RHIC. *Nucl. Instr. Meth.*, A461:143–149, 2001.

[85] Miroslav Plesko et al. Front-end electronics for the silicon partition of the PHOBOS detector at RHIC. *IEEE Nucl. Sci. Symposium Conference Record*, 1:76–80, 2001.

[86] R. Bindel, R. Baum, E.Garcia, A.C. Mignerey, and L.P. Remsberg. Array of cherenkov radiators for PHOBOS at RHIC. *Nucl. Instr. Meth.*, A488:94, 2002.

[87] R. Bindel, E. Garcia, A. Mignerey, and L. Remsberg. Array of scintillator counters for PHOBOS at RHIC. *Nucl. Instr. Meth.*, A474:38–45, 2001.

[88] C. Adler, A. Denisov, E. Garcia, M. Murray, H. Strobele, and S. White. The RHIC zero degree calorimeters. *Nucl. Instr. Meth.*, A470:488–499, 2001.

[89] Piotr Kulinich, Pradeep Sarin, and Andrei Sukhanov. The DAQ system with a RACEway switch for the PHOBOS experiment at RHIC. *IEEE Trans. Nucl. Sci.*, 49:2455–2458, 2002.

[90] William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes in C.* Cambridge University Press, second edition, 1999. Straight-line and general least-squares fitting procedures are described in Chapter 15. The Runge-Kutta method is described in Chapter 16. Section 20.4 describes Huffman coding.

[91] B.G. Gibbard and T.G. Throwe. The RHIC computing facility. *Nucl. Instr. Meth.*, A499:814–818, 2003.

[92] R. Brun and F. Rademakers. ROOT. *Nucl. Instr. Meth.*, A389:81, 1997.

[93] Rene Brun et al. GEANT version 3.21. Detector description and simulation tool. The GEANT webpage is http://wwwasd.web.cern.ch/wwwasd/geant/.

[94] Kristjan Herlache Gulbrandsen. *Relative Yields of Antiparticles to Particles in Au+Au Collisions at 130 and 200 GeV per Nucleon Pair.* PhD thesis, Massachusetts Institute of Technology, 2004.

[95] Jay Lawrence Kane. *Charged Hadron Transverse Momentum Spectra in Au+Au and d+Au Collisions at 200 GeV per nucleon pair.* PhD thesis, Massachusetts Institute of Technology, 2005.

[96] W.R. Leo. *Techniques for Nuclear and Particle Physics Experiments.* Springer-Verlag, second edition, 1994.

[97] B.B. Back et al. Ratios of charged particles to antiparticles near mid-rapidity in Au+Au collisions at $\sqrt{s_{NN}}$= 130 GeV. *Phys. Rev. Lett.*, 87:102301, 2001.

[98] B.B. Back et al. Centrality dependence of charged antiparticle to particle ratios near mid-rapidity in d+Au collisions at $\sqrt{s_{NN}}$ = 200 GeV. *Phys. Rev. C*, 70:011901(R), 2004.

[99] B.B. Back et al. Ratios of charged particles to antiparticles near mid-rapidity in p+p collisions at $\sqrt{s_{NN}}$= 200 GeV. *Phys. Rev. C*, 71:021901(R), 2005.

[100] B.B. Back et al. Centrality dependence of charged hadron transverse momentum spectra in Au+Au collisions from $\sqrt{s_{NN}}$ = 62.4 to 200 GeV. *Phys. Rev. Lett.*, 94:082304, 2005.

[101] K. Adcox et al. Measurement of the Λ and $\bar{\Lambda}$ particles in Au+Au collisions at $\sqrt{s_{NN}}$ = 130 GeV. *Phys. Rev. Lett.*, 89:092302, 2002. This paper also includes feed-down corrected proton and antiproton dN/dy values, based on the measured Λ distributions.

[102] S. S. Adler et al. Identified charged particle spectra and yields in Au+Au collisions at $\sqrt{s_{NN}}$= 200 GeV. *Phys. Rev. C*, 69:034909, 2004.

[103] K. Adcox et al. Single identified hadron spectra from $\sqrt{s_{NN}}$ = 130 GeV Au+Au collisions. *Phys. Rev. C*, 69:024904, 2004.

[104] B. Alper et al. The production of charged particles with high transverse momentum in proton-proton collisions at the CERN ISR. *Nucl. Phys. B*, 87:19, 1975.

[105] K. Guettler et al. Inclusive production of low-momentum charged pions, kaons and protons at x=0 at the CERN interacting storage rings. *Nucl. Phys. B*, 116:77, 1976.

[106] B. B. Back et al. Particle production at very low transverse momenta in Au+Au collisions at $\sqrt{s_{NN}}$= 200 GeV. *Phys. Rev. C*, 70:051901(R), 2004.

[107] Gabor Veres. Antiproton production in nuclear collisions at 158 AGeV/c. *Nucl. Phys. A*, 661:383c, 1999.

[108] B. B. Back et al. Baryon rapidity loss in relativistic Au+Au collisions. *Phys. Rev. Lett.*, 82:2471, 1999.

[109] H. Appelshauser et al. Baryon stopping and charged particle distributions in central Pb+Pb collisions at 158 GeV per nucleon. *Phys. Rev. Lett.*, 82:2471, 1999.

[110] Jurgen Schaffner-Bielich, Dima Kharzeev, Larry McLerran, and Raju Venugopalan. Scaling properties of the transverse mass spectra. 2002, nucl-th/0202054.

[111] Richard Witt. $< p_T >$ systematics and m_T-scaling. 2004, nucl-ex/0403021.

List of Figures

List of Tables

Acknowledgements

There are many people whom I wish to thank for their role in making this thesis possible.

First and foremost is my advisor Gunther Roland. He has been a terrific role model for a young physicist and I have learnt a huge amount from working under his supervision.

I would like to thank the other members of my thesis committee, Wit Busza and Iain Stewart, for their co-operation and contributions. Wit, as head of my group and spokesperson for my experiment, has also been a large influence on me, particularly in his approach to understanding high-energy physics.

The MIT heavy-ion group has been an excellent environment for doing physics. Bolek Wyslouch, George Stephans, Christof Roland, Gerrit van Nieuwenhuizen, Maarten Ballintijn and Constantin Loizides have been great senior colleagues (as well as challenges to my spelling abilities). My fellow students Patrick Decowski, Pradeep Sarin, Kris Gulbrandsen, Carla Vale, Jay Kane, Corey Reed, Ed Wenger, Sergei Vaurynovich, Burak Alver and Li Wei have made it an interesting, multicultural and fun place to work. Jay has also been a great friend, even if we were comically mismatched as work-colleagues. I have worked especially closely with Gabor Veres on many projects during my PhD and it has been a great experience – I owe him a large debt of gratitude for his advice and mentoring.

I have thoroughly enjoyed being a member of PHOBOS, a collaboration which steadfastly refuses to let its small size curb its grand ambitions. Just about every single member deserves my thanks in some way, so I feel bad about singling-out only a few. Still, here goes . . . Mark Baker was a great presence, Peter Steinberg ensured collaboration discussions were always interesting, Burt Holzman was ridiculously good at helping with computer problems. No PHOBOS analysis would be possible without the sterling work of Andrei Sukhanov on the DAQ, the UIC crew on triggering and event selection, and Marguerite Belt Tonjes (and before her, Nigel George) on data production. My analyis of Time-of-Flight data could not have happened without the extensive help of Frank Wolfs and Erik Johnson. George Stephans, Gerrit, Marguerite and Dave Hofman served as the review committee for my analysis, and I thank them for all their efforts. Anna Convertino and Lois Caligiuri have made my life considerably easier through their excellent administrative work.

MIT is a great and inspiring place, but it can be tough. The American Jiu-Jitsu club and the THDA helped make it easier for me. Ulrich Becker gave me good academic guidance and Brian Canavan always answered my foolish questions patiently. The John F Kennedy Memorial Trust supported me during my first year at MIT, and Anna Mason has been a great resource.

I thank Lisa, Dave, Rab, Quinn, Colm and Gerdy for being great friends and I have to thank my whole family for being a tremendous bedrock. I am very proud to be a Henderson.

Finally, and most importantly, I owe everything to my parents Joe and Nuala. I will be eternally grateful to them for all their love and support, and for teaching me so much. I dedicate this thesis to them.

Biography

Conor Henderson was born in Lurgan, Northern Ireland on 14th August 1977. He attended Queen's University, Belfast, Northern Ireland, graduating with an MSci in Physics and Applied Mathematics, and won the inaugural John Geddes Prize for best physics research project. He came to MIT in 1999 as a John F Kennedy Memorial Scholar. He is very pleased that you are reading his thesis.