

Search For Rare Events in $\sqrt{s_{NN}} = 200$ GeV Au+Au
PHOBOS Data

by

Alexander Mott

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

BACHELOR OF SCIENCE

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

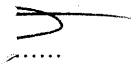
[June 2009]

May, 2008

©2009 ALEXANDER MOTT

All Rights Reserved

The author hereby grants to MIT permission to reproduce and to distribute
publicly paper and electronic copies of this thesis document in whole or in

Signature of Author


Department of Physics
May 11, 2009

Certified By

George Stephans
Thesis Supervisor, Department of Physics

Accepted by

Professor David E. Pritchard
Senior Thesis Coordinator, Department of Physics



ARCHIVES

Contents

1	Introduction	4
1.1	The Relativistic Heavy Ion Collider	4
1.2	Physics and Terminology	6
1.3	PHOBOS Detector	8
1.3.1	Multiplicity Array	10
1.3.2	Triggering System	11
1.4	Events	13
1.5	Types of Rare Events	15
1.5.1	Black Holes	16
1.5.2	DCCs	18
2	Event Selection	21
2.1	Preliminary Event Selection	21
2.2	Pileup Removal	23
3	Experimental Procedure	27
3.1	Binning and Data Structure	27
3.2	Validation Toy Model	30
3.3	Analysis Chain	31
3.3.1	Calculating Average Multiplicity	31
3.3.2	Covariance Matrices	32
3.3.3	Calculating χ^2_ν	35
4	Toy Monte Carlo	41
4.1	A More Advanced Toy Model	41
4.2	Spike Model	44
4.3	What is Missing From the Toy Models	45
5	Results	48
5.1	Statistical Event Distribution	48
5.2	Rare Events	52
6	Conclusion	63
7	Appendix	67
7.1	Report Card Examples	67

Search For Rare Events in $\sqrt{s_{NN}} = 200$ GeV Au+Au PHOBOS Data

by

Alexander Mott

Submitted to the Department of Physics
on May 8, 2009 in partial fulfillment of the
requirements for the degree of Bachelor of Science in
Physics

Abstract

In this analysis, we set an upper bound on the rate of rare events: events whose $\frac{dN}{d\eta}$ distribution deviates more than statistically from the ensemble average $\frac{dN}{d\eta}$ distribution for $\sqrt{s_{NN}} = 200$ GeV Au+Au collisions in PHOBOS data. We carefully remove events that may exhibit non-statistical fluctuations due to other effects, such as event pileup and detector pathology. We also use very fine binning in the z and y vertex positions to eliminate fluctuations due to different event vertices. We eliminate global correlations within the $\frac{dN}{d\eta}$ distributions by using a covariance matrix in the analysis, which is used to scale out correlations between difference pseudorapidity (η) regions.

In the end we produce a value of χ^2 per degree of freedom (χ^2_ν) for each event, which reflects how well the event agrees with the ensemble average. We compare this distribution with the distribution we would expect for a model using uncorrelated random variables with the same degrees of freedom of our system, and find remarkable agreement between our data and this random distribution. This allows us to conclude that most events in our data set are statistical fluctuations about the ensemble average. We are further able to determine that there is a signal of non-statistical events with $\chi^2_\nu > 2.22$, and that the rate of these events in the PHOBOS data is less than $(1.97 \pm 0.4(stat) \pm 0.1(sys)) \times 10^{-5} \frac{\text{rare events}}{\text{event}}$.

Thesis Supervisor: George Stephans

Title: Senior Research Scientist

1 Introduction

1.1 The Relativistic Heavy Ion Collider

The Relativistic Heavy Ion Collider (RHIC) is a large particle collider located at Brookhaven National Laboratory in Upton, NY. It is capable of accelerating two counter-propagating elliptical beams of heavy ions (Au^{79+}) up to a center of mass energy per nucleon pair of $\sqrt{s_{NN}} = 200$ GeV with 6 crossing points for interactions. There were 4 experiments (BRAHMS, PHENIX, PHOBOS, and STAR) running during the physics runs from 2000-2005, each located at one of these crossing points.

Figure 1 shows the layout of the RHIC experimental facility. Au^{1-} ions are fed into the Van de Graaf generator, which accelerates and strips electrons from the gold ions producing a beam of Au^{32+} ions with a total energy of roughly 1 GeV each [1]. The Au^{32+} ions are fed into a booster ring where they are accelerated up to 1.03 GeV/nucleon and have 45 additional e^- stripped from them, leaving a beam of Au^{77+} ions. From there, they are fed into the AGS and accelerated up to 9.79 GeV/nucleon and stored. The AGS feeds multiple experiments, one of which is RHIC.

The output from the AGS is fed into RHIC, and accelerated up to 100 GeV per nucleon by RF cavities arranged along the beamline. Superconducting magnets curve the beam along the circular beam line, and into the 6 interaction points spaced evenly along the beam path [1]. At 100 GeV/nucleon, these Au^{79+} ions are highly relativistic, so we expect them to lose energy due to synchrotron radiation. We can calculate the loss due to synchrotron




Figure 1: Schematic of the RHIC experimental facility (*Source: LLNL.gov*)

radiation in one full revolution using a formula from [2] (in which $\hbar = c = 1$):

$$\Delta E = -\frac{4\pi\alpha}{3R} \left(\frac{E}{m}\right)^4 \quad (1)$$

Evaluating this for the RHIC accelerator (the radius is $R \approx 610$ m) we find that $\Delta E \approx 1$ keV/revolution ≈ 78 GeV/s is the energy loss for a single gold ion. This is a relatively small amount of energy, which we must add back into the system (done using RF cavities to boost and quadrupole magnets to

focus the beam along the beamline).

1.2 Physics and Terminology

The vast energy of the heavy ion collision is localized into a very small physical area, and occurs in a very short amount of time, so the energy density is very high. The particles produced from this region are generally vastly more numerous than those that went into the collision (roughly a factor of 10 for the most central collisions).

One traditionally parametrizes scattering and collision processes using the impact parameter, b , of the collision, which is defined as the transverse distance between the centers of the colliding ions. In relativistic heavy ion collisions, however, this parameter is not particularly useful to characterize the events and understand the output [3]. The more useful parameters, which we shall use throughout are N_{coll} and N_{part} . N_{part} is defined as the number of nucleons that participate in the collision, while N_{coll} is the number of nucleon-nucleon collisions in the event. These are related to the impact parameter, but encode more useful information. The complement of N_{part} is the number of spectators nucleons, which is defined as the number of nucleons which don't interact in the event. One can see that if $b = 0$, then $N_{part} \approx N_{nucleons}$ and $N_{spectators} \approx 0$. Obviously, these relations will not be exact in the real world, but in theory this would be observed.

We would also like to encode the impact parameter of the events in some easy to use variable, which we call the centrality of the event. More specifically, we would like to have certain bins of centrality which encompass a small range of impact parameters, so that we can easily address groups of

similar data. We call these centrality bins; there are 17 of them, and we create them by dividing the data into bins of fractional cross-section. In practice, we bin data in paddle signal because, as we will discuss later, this is directly related to centrality. For instance, the 17th centrality bin (the most central bin) contains the top 3% of events in centrality as measured by the paddles. This should correspond closely to the lowest 3% in centrality and impact parameter, but will include a small tail of events outside the desired region.

The distributions used in this analysis will be measured in pseudorapidity, which is a commonly used parameter of relativistic systems. It is an angular parameter which describes the angle of a particle with respect to the beamline. If we define θ as the angle between the particle's momentum and the beamline, then the pseudorapidity (η) is:

$$\eta = -\log(\tan(\theta/2)) \quad (2)$$

One can see that $\eta = 0 \implies \theta = 90^\circ$, so the center of the pseudorapidity region is the same as the center of the region in polar coordinates. When we say “mid-rapidity” we mean the area centered at $\eta = 0$. The phrase pseudorapidity suggests that there is something called rapidity; this is usually denoted y , and is given by:

$$y = \tanh^{-1}\beta \quad (3)$$

where $\beta = \frac{v}{c}$. Rapidity is an important variable in special relativity, the

simplified form of a Lorentz boost is given by [4]

$$\Lambda(y) \begin{pmatrix} ct \\ x \\ y \\ z \end{pmatrix} = \begin{pmatrix} \cosh(y) & -\sinh(y) & 0 & 0 \\ -\sinh(y) & \cosh(y) & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} ct \\ x \\ y \\ z \end{pmatrix} \quad (4)$$

Rapidity is, in general, the more important for physical phenomena, but pseudorapidity is much easier to measure in the lab (the angle of the particle is much easier to determine than its velocity). Fortunately, in the highly relativistic limit, where particles in this experiment live, the two are very closely related. This can be most easily seen by writing [4]:

$$\eta = \frac{1}{2} \log \left(\frac{|\vec{p}| + p_{\parallel}}{|\vec{p}| - p_{\parallel}} \right) \quad (5)$$

$$y = \frac{1}{2} \log \left(\frac{E + p_{\parallel}}{E - p_{\parallel}} \right) \quad (6)$$

Here p_{\parallel} is the component of momentum along the beamline. One can see that in the limit where $\frac{m}{c^2} \ll \frac{p}{c}$ then $E \approx p$, so in the limit rapidity and pseudorapidity are equal.

1.3 PHOBOS Detector

The PHOBOS detector (shown in Figure 2) consists of 4 primary systems: the multiplicity array, the Vertex detector, the Spectrometer, and the trigger system. The detector is covered extensively in Ref [5], so only the level of detail necessary for this analysis will be provided here. The multiplicity array provides large solid angle coverage for counting the number and angular

distribution of the particles emitted in the collision. The Vertex detector uses finely segmented Si sensors to record particle hits, from which tracks can be derived and the primary interaction vertex determined. The Spectrometer is used to track, identify, and measure the momentum of a small fraction of the particles; this functionality, however, is not used in this analysis. The trigger system is designed to pick out potentially interesting and well-formed events and initiate event recording; it is also responsible for determining the approximate impact parameter of the collision. The Spectrometer is situated in a 2.18T magnetic field, generated by a double dipole magnet in order to bend the charged species which are produced in the interaction.




Figure 2: The PHOBOS detector. The relevant detector subsystems are described in detail in the text (From Ref [3])

1.3.1 Multiplicity Array

The multiplicity array detects the number and angular distribution of particles produced in the collisions in the pseudorapidity range $-5.4 < \eta < 5.4$ and angular range $0 < \phi < 2\pi$ [5]. The array consists of an octagonal detector which covers $-3.2 < \eta < 3.2$ and 3 pairs of rings cover the remaining region. Apart from 4 small holes where Si wafers are removed from the Octagon (2 to accommodate the Spectrometer, and 2 to accommodate the Vertex detector) near mid-rapidity, the multiplicity array covers the full available region outside the beam pipe. Hits measured by the Spectrometer and Vertex detectors are used to partially fill these holes, so the effective missing region is very small.

The Octagon is centered around mid-rapidity, and extends 0.55 m in each direction along the beam pipe; the diameter is 90 mm, measured from the center of each opposing face [5]. The detector contains 92 silicon pad wafers arranged in 8 rows (giving it an octagonal shape), which are divided further into 120 pads (rows of 30 in the beam direction). The individual pads cover an eta range of 0.06 near mid-rapidity to 0.005 near the edge of the Octagon. The total energy deposited in each pad is read out individually for each event and stored.

The remainder of the η coverage is provided by the ring detectors, which are divided into 8 trapezoidal sensors each consisting of 64 pads [5]. The pad sizes increase with distance from the beam pipe, and the sizes are chosen so that each pad subtends a pseudorapidity range of $\Delta\eta \approx 0.1$. There are three pairs of rings which, together cover the region $3.2 < |\eta| < 5.4$; ring 1 covers $3.4 < |\eta| < 4$, ring 2 covers $4 < |\eta| < 4.6$ and ring 3 covers $4.6 < |\eta| < 5.4$.

1.3.2 Triggering System

Primary triggering is provided by the paddle trigger counters [1]. These cover the range $3 < |\eta| < 4.5$, and are each constructed of 16 scintillators located ± 3.21 m from the center of the interaction region. Primary triggering is done by requiring nearly simultaneous hits in both the positive and negative paddles within 76 ns of each other. This condition is virtually guaranteed to be true during any collision, but can also be satisfied by a variety of backgrounds, which are rejected by other online and offline methods (some of which will be discussed in Section 3).

The paddles are also used to determine the centrality of an event, based on the total number of charged particle hits in the scintillators. As discussed in Section 1.2, the number of participant nucleons in the collision is related to the impact parameter of the collision, since a large impact parameter means most of the nucleons miss each other without interacting. It is separately observed that N_{part} is proportional to the number of particles produced in the collision, and hence with the total energy deposited in the paddles. The paddles have a very fast integration time, since they are scintillators, so they provide a very good estimation of the centrality of the event.

Rather than deal directly with the impact parameter, we instead sort the events into centrality bins corresponding to ranges of paddle signal, which is a measurable experimental variable correlated with the centrality of the event. In this analysis, we deal mostly with the 10% of the events which have the highest paddle signal: the 10% most central events. These are organized into 3 bins: 0-3%, 3%-6% and 6%-10% most central events. Figure 3 shows the mean of the signals in the two paddle detectors for all events, corrected

for inefficiency for peripheral events. The top 10% of that distribution is all of the events with a mean paddle signal above approximately 1500.

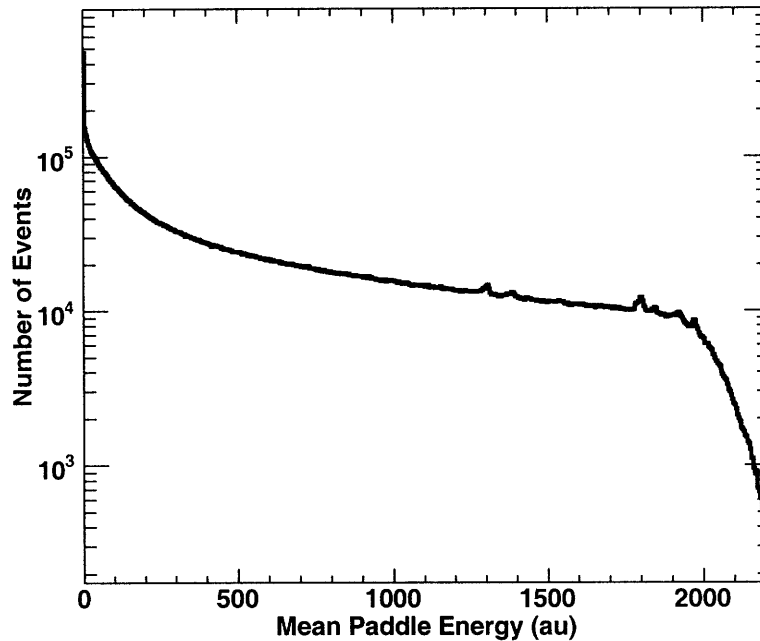


Figure 3: Distribution of the mean signal in the paddles

The final elements of the triggering system are the zero-degree calorimeters (ZDC), which are designed to measure the neutron component of the spectator nucleons in the collision. The ZDCs are situated at ± 18 m from the interaction vertex directly along the beamline. Magnets situated between the interaction point and the ZDCs will bend any charged species away, so the ZDCs will detect only the neutral component of the spectator nucleons. The ZDCs are used for triggering by requiring near coincident neutron hits in the forward and backward ZDCs.

The ZDCs are also strongly inversely correlated with centrality for the top 50% most central events. For the most central events, most of the nucleons will interact, so the number of spectators will be very low. This will lead

to a small signal in the ZDC for very central events. For very peripheral events, we expect a small number of participant nucleons (since the impact parameter is very big), and hence a large number of spectator neutrons. This will lead to a large signal in the ZDC for high impact parameter. At very large impact parameters, there are a large enough number of spectator nucleons to allow the spectator neutrons and protons to clump into small nuclei (alpha particles, deuterium, etc.). Since these small nuclei are charged, the magnets will bend them away, and they will not be registered by the ZDCs. This will have the effect of decreasing the ZDC signal for very peripheral (low paddle) events. Figure 4 shows the distributions of Negative ZDC energy and Mean Paddle Energy (correlated to centrality), and shows this trend in actual PHOBOS data. We can also see the clustering take over in the peripheral region, driving the ZDC signal down. In the extreme of lowest paddle signal, we expect a very glancing hit, so most of the spectator neutrons will be bound up and we see very small ZDC signal.

1.4 Events

PHOBOS was able to achieve very good azimuthal coverage for charged particle, and so was able to measure the density of charged particles ($\left(\frac{dN}{d\eta}\right)_{ch}$) out to very large η values compared to other experiments at RHIC. The fully corrected mean $\frac{dN}{d\eta}$ distribution is shown in Figure 5 [3]. This distribution agrees well with the predicted shape of the $\frac{dN}{d\eta}$ found in Kharzeev *et al.* [6], as well as HIJING Monte Carlo simulations [1].

It is also important to note that the statistical errors on Figure 5 are very small indeed. In fact, the statistical error bars on the figure are actually

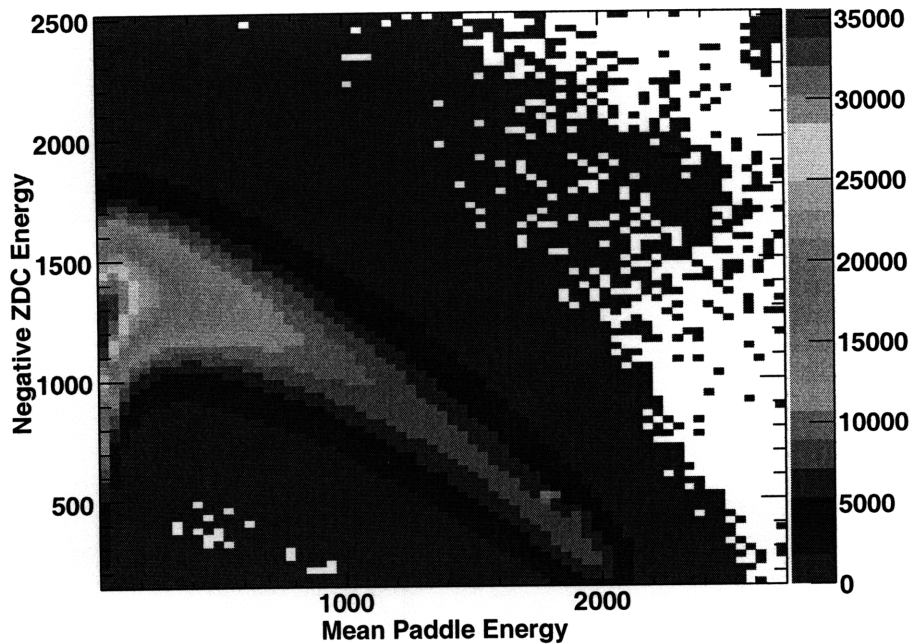


Figure 4: Distribution of events by Negative ZDC Energy and Mean Paddle Energy. Note the downward correlation of Paddle Mean Energy and ZDC energy in the high paddle region.

smaller than the symbol used to draw the points. This allows us to conclude that the RMS of the events is sufficiently small to allow very good determination of the mean. This will be important for our analysis, because we will be using an ensemble average, so we need this average to be well determined. The original in [3] also includes systematic uncertainties, which can be large, but in this paper we are concerned only with the statistical deviation of the events, which is very small.

Figure 5 shows the fully corrected $\frac{dN}{d\eta}$ distribution for PHOBOS events, but the actual output of the detectors looks very different. Figure 6 is the mean raw detector output of 0-3% central Au+Au events. Again, the statistical errors are entirely hidden by the data points. To reproduce Figure 5 from




Figure 5: Corrected $\frac{dN}{d\eta}$ distribution for 0-6% central 200 GeV Au+Au Collisions at PHOBOS. Adapted from Figure 1 in [3].

Figure 6, one must correct for a variety of physical parameters of the system, including the geometry of the detector, the acceptance of the pads, saturation of certain pads relative to others, and background particle removal, which can be very large for some pads. All of these procedures introduce some potential for error, so for this analysis we use the uncorrected, raw pad outputs.

1.5 Types of Rare Events

The goal of this search is to put an upper limit on the rate of occurrence of certain types of events in the PHOBOS data set, so it is worthwhile to discuss the types of events which might occur, and to which this analysis is sensitive. This analysis will be sensitive to events which effect the total number of par-




Figure 6: Raw $\frac{dN}{d\eta}$ distribution for a sample of 0-3% central Au+Au events.

ticles produced (N_{hits}), or the shape of the $\frac{dN}{d\eta}$ distribution. Our sensitivity to small deviations, which fall within the range of statistical fluctuations is poor, but our sensitivity to fluctuations in the number of high N_{hits} and very unusual shape fluctuations is very good. Two interesting events will be singled out for discussion as illustration: miniature black hole formation, and disoriented chiral condensate (DCC) production. It is important to stress, however, that this analysis will not be able to detect whether or not these events exist in the data, merely a limit on the rate at which they occurred.

1.5.1 Black Holes

While considered very unlikely, the production of a black hole is theoretically possible in RHIC collisions [7] [8]. If a black hole is produced, it would radiate its energy away very quickly due to Hawking radiation, and decay,

which would produce a characteristic signature in the detectors.

A PYTHIA simulation of a black hole event inside the ATLAS detector at the LHC is shown in [8]. Figure 7 shows the $\frac{dN}{d\eta}$ distribution for a black hole event versus a regular event (labeled as “initial radiation” in that paper). There is clearly a pronounced difference between the two events, with a great deal of enhancement at midrapidity and suppression at larger rapidities for the black hole event [8]. This is due to the preferential absorption of particles with low p_T by the black hole, and re-emission with a more isotropic distribution.

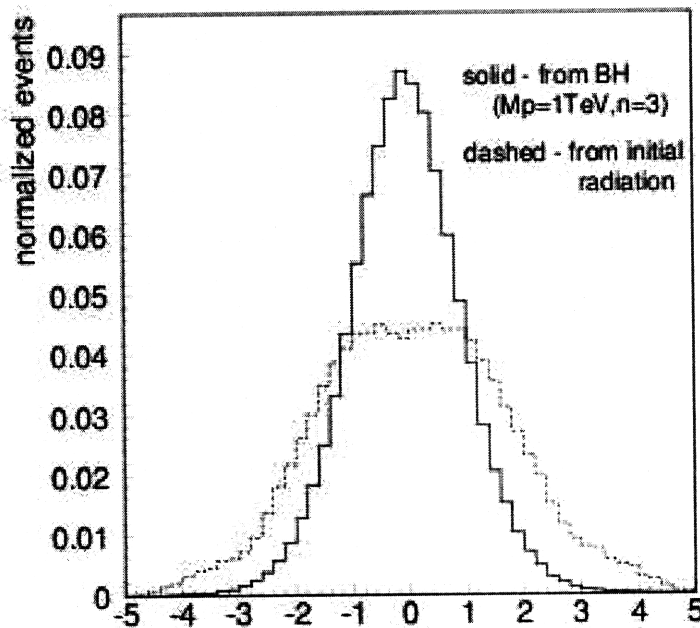


Figure 7: A PYTHIA simulation of the expected $\frac{dN}{d\eta}$ for a black hole event compared to that for a regular event. This figure is a reproduction of Figure 8 in Tanaka *et al.* [8].

Such black hole events would exhibit not only the shape deviations de-

scribed above, but also N_{hits} deviations. Figure 8 shows a comparison of the total multiplicity distribution for a regular event and for an event with a black hole decay [9]. The overlap of the distributions for the regular event and for the black hole makes this a less precise channel to exclude black hole production, since we cannot discriminate black holes in the overlap region. Coupled with the shape fluctuations, however, this analysis is expected to very strongly constrain the rate of production of sufficiently large black holes in our data.

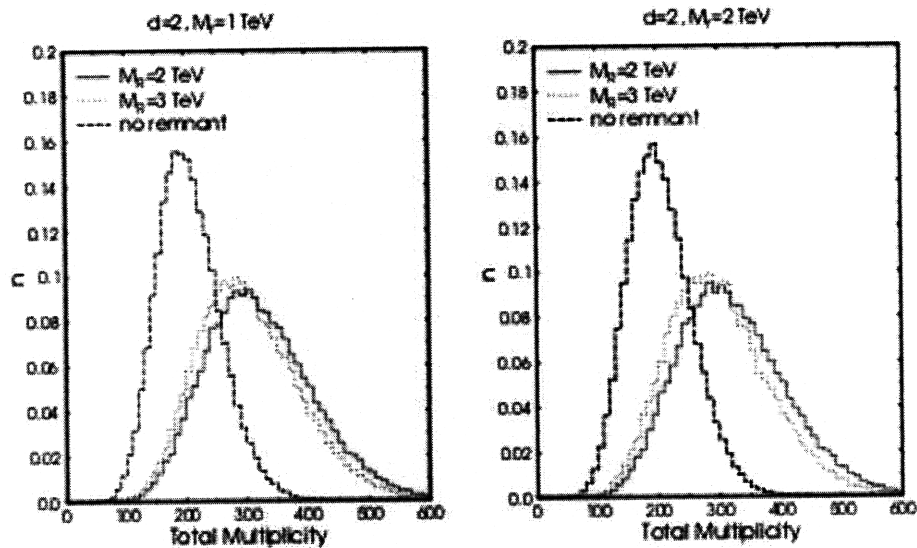


Figure 8: The expected total multiplicity distribution for a black hole event versus that for a normal event. This figure is a reproduction of Figure 5 in Stocker [9].

1.5.2 DCCs

A Disoriented Chiral Condensate (DCC) is a possible state of the vacuum corresponding to the restoration and re-breaking of the chiral symmetry between the up and down quark. According to Ref. [10], the Lagrangian for

the fermionic part of the QCD Lagrangian is

$$\mathcal{L} = \bar{u}_L i \not{D} u_L + \bar{d}_L i \not{D} d_L + \bar{u}_R i \not{D} u_R + \bar{d}_R i \not{D} d_R \quad (7)$$

where the u 's and d 's represent up and down quarks and the L and R represent the chirality. Equation 7 is separately invariant under unitary transformations of the vectors $\begin{pmatrix} u \\ d \end{pmatrix}_L$ and $\begin{pmatrix} u \\ d \end{pmatrix}_R$. At high temperature, we can have free quarks and so this symmetry is realized in the high T region. As the system cools, however, the symmetry is spontaneously broken by binding of quark-antiquark pairs. When the quarks begin pairing, the quark-antiquark doublet gains a non-zero vacuum expectation value, which breaks the full SU(2) symmetry. The pseudo-Goldstone bosons corresponding to this symmetry breaking are the three flavors of pion.

The vacuum which condensed from the big bang (that is to say, the one we live in) is oriented orthogonal to the three π directions (in the σ direction using the sigma model) [11]. There is no *a priori* reason that a system restored to the symmetric state would re-decay into the sigma state, so in general an arbitrary symmetry breaking state will have some component in one or more of the π directions. This state is a DCC; disoriented because its orientation is different from the main vacuum. According to [11] the interaction of this state with the σ vacuum will produce a large number of whichever flavor pion the DCC is oriented toward.

This makes PHOBOS sensitive to certain types of DCC, and this analysis capable of discriminating them. PHOBOS is a charged particle detector, so we will not see a DCC oriented along π^0 (an anti-Centauro type), but we can

see a charged pion DCC (Centauro type), which will appear as a spike in the number of charged particles detected in a narrow η region [12]. The DCC signal is less pronounced than the signal from the black hole, but this analysis could place production limits as a function of the size of the Centauro type DCC.

2 Event Selection

2.1 Preliminary Event Selection

This analysis is searching for fluctuations due to physically interesting events in the detector, so all sources of unphysical fluctuation must be removed before this search can occur. We first ignore any event which occurred in one of a half-dozen bad runs where various elements of the detector were miscalibrated or not working properly. We further remove events that occurred during a drift in the y position of the beam. While the events may not all be bad, they exhibit a different $\frac{dN}{d\eta}$ pattern than non-drift events, and since the statistics in this region are low, we remove them rather than trying to correct for them.

The raw hits from the detector must first be merged into a format which is readable for the analysis. The details of this are not crucial to this analysis, but can be found in [1]. We ignore pads within 0.2 in η of the end of the Octagon's eta coverage, because these receive hits in less than 50% of the events and create fluctuations.

We impose conditions on the vertex of the events to ensure that it is not too far from the average vertex, and to ensure that the vertex reconstruction worked properly for the event. We require that the x vertex of the event be in the range $-0.3 \text{ cm} < x < 0.3 \text{ cm}$ where 0 is the center of the interaction region, and require that the z vertex be in the range $-10 \text{ cm} < z < 10 \text{ cm}$. If these conditions are not met, the event will exhibit a strange $\frac{dN}{d\eta}$ due only to its vertex position and not any physically interesting event in the detector.

The vertex position of the event is crucial to our track reconstruction and errors in this determination will lead to errors in the $\frac{dN}{d\eta}$ distribution. We

determine the vertex position in 5 different ways, and use the three most reliable to cross-check that all systems worked correctly for the specific event in question. The Vertex detection subsystem computes a main vertex for the event, to which we compare the vertex positions reported by other subsystems. We also measure the vertex using the Spectrometer; the Spectrometer measures charged particle tracks through its silicon layers, so it is possible to compute the initial starting point of the event. The third method is to use the paddle time difference to compute the z vertex of the event, which is very reliable since all particles travel very close to c . For our event selection, we require that the Spectrometer and paddle vertices are within 3σ of the vertex determined by the Vertex detector. Here σ is the standard deviation of the event vertex from the mean for all events, and is roughly 0.03 cm. If this is not the case, we conclude that there was some problem in the electronics, perhaps noise, perhaps a malfunction in a subsystem, but we exclude the event to avoid introducing that error into our final calculation.

These quality cuts undoubtedly remove many perfectly valid events from consideration, but are necessary to ensure sensitivity to fluctuations. We expect to be looking for a very small signal of interesting non-statistical events on top of a very large number of statistical events [13]. We must, therefore, exclude any events which even have a chance to include unphysical fluctuations, otherwise most of the observed unusual events will simply be events with detection errors. It is preferable, therefore, to eliminate many good events rather than keep a few bad events, as long as the selection is done to eliminate sources of error, rather than specific events which deviate from the norm.

2.2 Pileup Removal

Pileup occurs when two collisions occur in the detector within the resolution time of the whole detector, and are counted as a single event. Pileup events will have a very different $\frac{dN}{d\eta}$ distribution from regular events, but are obviously not physically interesting processes, so we would like to eliminate them. We can imagine three possible cases for pileup events: two very central collisions, two very peripheral collisions and one central plus one peripheral collision. We need to be able to discriminate all three of these in order to effectively remove pileup.

The paddles have a very short integration time, because scintillator signals are intrinsically very fast, so for most pileup events they will only register one of the events. Since we only consider events which the paddles tag as highest centrality, we can eliminate the events where two peripheral events pileup, since the paddles will detect one of these and never tag it as central. It is possible to have two events occur within the integration time of the paddles, but the rate is much lower, and we will eliminate those in other ways. We can further eliminate pileup by using the difference in integration time between the paddles and multiplicity array to look for a characteristic pileup signature.

The paddles and the rings cover similar η ranges, but have very different integration times; we can use this difference to look for pileup. Since the paddles and rings detect, in general, the same particles, we expect their energy values to be very strongly correlated for normal events. In a pileup, however, the paddle will only detect one of the events, while the ring will detect both, so we expect the ring energy to be greater than the paddle

energy. This is the only reasonable situation in which this will happen, so by eliminating events which look like this, we can remove pileup events which the paddle discriminates, but the rings don't. Figure 9 shows the total mean paddle energy versus the total ring energy for our events. The tail of events with high ring energy but intermediate paddle energy is exactly the noise signal we expect: two events are registered in the rings, but only one is counted by the paddle. We cut along the diagonal line to remove these events, but not eliminate the events which have both high ring energy and high paddle (this might be pileup, but they might also be legitimate signal).

Paddle Mean Value versus Ring Energy




Figure 9: Paddle Mean versus Ring Energy for 0-3% central events. The solid line indicates the position of the pileup cut (we exclude events to the right of that line).

We can eliminate pileup events not caught by the paddle vs ring cut by appealing to the inverse relationship between the paddle and the ZDC

discussed earlier and shown in Figure 4. In a normal central event, we expect a low ZDC signal, since the number of spectators is low. In a pileup event, where there is one central event and one more peripheral event, however, we would expect to see a large signal in the ZDC from the spectators in the peripheral collision and a large signal in the paddles from the central event. Because the integration time of the ZDC is matched much more closely to the multiplicity array, we use it rather than the paddles, but the same relationship should hold for the ZDC and total energy as holds for the ZDC and paddles. Figure 10 shows the relation between the ZDCs (at ± 18 m) and the total energy measured by the multiplicity array. The negative correlation is obvious in these plots, and the similarity between these plots and Figure 4 is also evident. We can cut along a line which mirrors this negative slope in order to eliminate events with high ZDC and total energy signals, which we believe are pileup.

In principle, the negative and positive ZDC signals should be identical, over a large number of events. The two plots in Figure 10 differ because differences in the physical setup of the two ZDCs lead to slight differences in their calibration. While there is a calibration factor to make the signals look more similar, it is not sufficiently well understood and reliable to apply for this analysis. The positive ZDC is also known to be less well calibrated than the negative ZDC, so we impose a much less severe cut on the positive ZDC, and use the negative ZDC to do most of the work.

While it is not possible to exactly show that all the pileup events are removed, if we could do that we could eliminate all pileup events, we can look at some other expected features of pileup events and see how our cuts reduce those. One would expect, for example, that pileup events would produce more

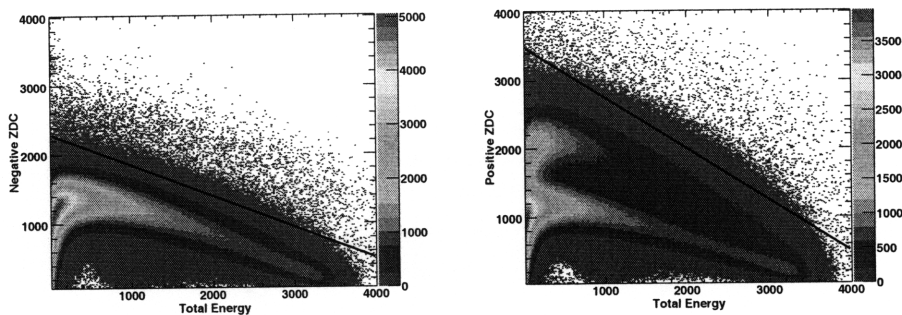


Figure 10: Positive (left) and Negative (right) position ZDC energies versus Total energy detected by the multiplicity array.

total particles than regular events, this is encoded in a variable called N_{hits} . While we cannot directly cut on N_{hits} , since we would like to see interesting physics events that might affect the total multiplicity, it is possible to see how the N_{hits} distribution is affected by the cuts. Figure 11 shows the N_{hits} distribution of our data set both before and after pileup cuts are applied. From the figure, we can see the effect that the pileup cuts have on the N_{hits} distribution; the pileup cuts remove most of the high N_{hits} events, where we expect pileup events to exist.

Since many of our pileup cuts will overlap in terms of the events that they eliminate, one can also look at how one pileup cut affects the other. Figure 12 shows the paddle energy versus ring energy plot after the ZDC cuts have been applied. We can see that the ZDC cuts effectively remove much of the trail of high ring energy and intermediate paddle energy events. Even more striking is the reduction in the high paddle and ring energy events, which are almost totally removed by the ZDC cut.

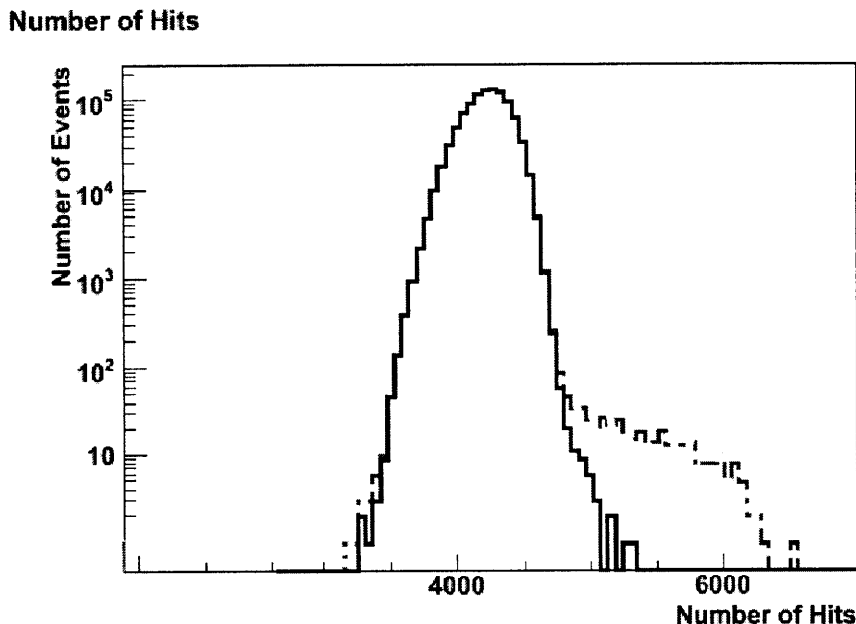


Figure 11: Distribution of the number of hits per event before (red, dashed) and after (blue, solid) event selection.

3 Experimental Procedure

3.1 Binning and Data Structure

This analysis centers around separate modules, which compute the quantities required to assign a χ^2_ν value to each event. χ^2_ν is the χ^2 value per degree of freedom, which is simple $\frac{\chi^2}{NDF}$ where NDF is the number of degrees of freedom. We begin by separating the events into 400 vertex bins, based on their y and z vertex positions in order to avoid comparing $\frac{dN}{d\eta}$ distributions that differ only due to vertex position in the final result. The z vertex (along the beamline) falls in the range $-10 \text{ cm} < z < 10 \text{ cm}$, and is separated into 200 equal bins of 0.1 cm each. The y vertex (transverse to the beamline) runs from $-0.2 \text{ cm} < y < 0.1 \text{ cm}$, and is separated in 2 bins: $[-0.2, 0)$, $[0, 0.1]$.

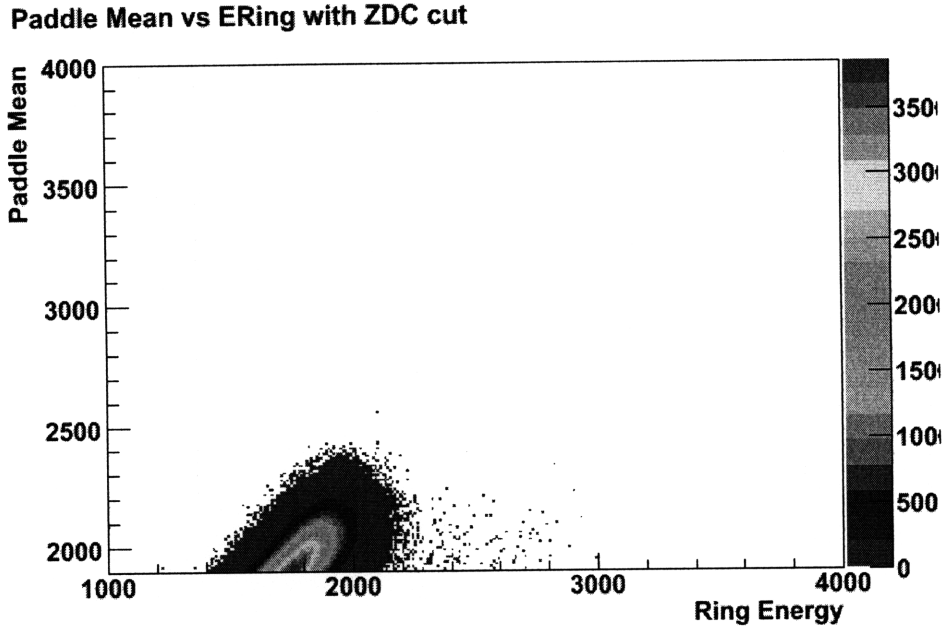


Figure 12: Ring energy versus paddle energy after the ZDC cut. Compare to Figure 9, which shows the same plot before the pileup cuts.

The bins are arranged so that the first 200 vertex bins correspond to the 200 z vertex bins in the $[-0.2, 0)$ y bin, and the next 200 bins follow the same prescription in the $[0, 0.1]$ y bin.

We use the y vertex bin because early physics runs in the data sample under analysis showed a beam drift to slightly negative mean y vertex (Figure 13). The shift in vertex position could create a slight difference in the detected $\frac{dN}{d\eta}$ for these events compared to later runs where the drift was fixed. By binning in y vertex and considering each vertex bin totally independently, we eliminate or substantially reduce any potential fluctuation from this drift. Since no drift is observed in the x position of the beam, no x bins are created, as this would only harm our final result by reducing the statistics in each bin.




Figure 13: Distribution of events by y vertex position and run number. We see that early runs had a shifted y vertex position; we bin in y vertex to mitigate this potential source of fluctuations.

Finally, we get each event's $\frac{dN}{d\eta}$ distribution by combining, as mentioned previously, the hit information of the Vertex detector, the Spectrometer, the Octagon, and the three pairs of Rings, which gives us full coverage in the η range $-5.2 < \eta < 5.2$. Each event's $\frac{dN}{d\eta}$ distribution is separated into 52 η bins, spanning the full pseudo-rapidity range, corresponding to a bin size of $\Delta\eta = 0.2$. Since we are primarily looking for shape fluctuations, we apply no correction factors or scaling to any of the outputs, instead relying on raw detector energy outputs. This allows us to get the most unbiased shape information without having to worry about the additional uncertainty introduced by scaling factors and acceptance corrections.

3.2 Validation Toy Model

All sections of the analysis chain are validated using a simple toy model with easy to understand behavior. We simulate a toy event by randomly assigning a vertex bin, and create a mock $\frac{dN}{d\eta}$ distribution where each bin is thrown like a Gaussian with mean of 100 and RMS of 10 (Figure 14). These toy events are very simple, but will be useful to validate that the analysis apparatus is working correctly, since we can say *a priori* what we expect the output of each module to look like. We generate 2 million of these events to roughly match the statistics of the real data. We will discuss the expected output, and compare it with the actual output in every module section as a means of showing that the module is doing what it is supposed to.




Figure 14: An example of the $\frac{dN}{d\eta}$ distribution for a single validation toy event. Each bin is a Gaussian deviation from the mean (100) with RMS (10).

3.3 Analysis Chain

The analysis chain consists of three modules acting on the already pre-processed data. All modules act on each vertex bin independently and so produce one output per vertex bin (except for the final module, which produces one output per event). The analysis modules are, in order: the average $\frac{dN}{d\eta}$ module, the covariance matrix module, and the χ^2_ν module. All of these modules act on $\frac{dN}{d\eta}$ distributions which have been divided up into 52 η bins, as described above.

3.3.1 Calculating Average Multiplicity

The first module computes the average $\frac{dN}{d\eta}$ distribution for the set of events in each vertex bin. It does this on a detector by detector basis as well as for the entire distribution. This average is used as the baseline, to which all events are compared to compute the χ^2_ν values. Since there are vastly more normal events than unusual events in our sample, we conclude that the presence of unusual events in the creation of this average $\frac{dN}{d\eta}$ does not significantly affect the final result.

We expect the output of this module for our toy model to produce a distribution which is flat in average $\frac{dN}{d\eta}$, and has a mean of 100. Figure 15 is an average over 100 vertex bins, which is roughly 5×10^5 events. We see that the mean is 100.002 ± 0.05 , and the RMS is 10.0002 ± 0.006 . These errors are so small because we have 2×10^6 events, and so are able to determine the mean and RMS very precisely.




Figure 15: Output of the average $\frac{dN}{d\eta}$ module for the validation toy

3.3.2 Covariance Matrices

The next step is to calculate the covariance matrix for each vertex bin. It is known that events usually contain some strong correlations, especially between adjacent or near adjacent bins. Among other things, we expect every bin to be positively correlated with N_{hits} , which leads to global positive correlations between the bins. Since these correlations are very difficult to model, we instead attempt to remove them, to look for non-global fluctuations that may characterize unusual events. We do this by computing the covariance matrix, which encodes correlation information within it.

We compute the correlation matrix by:

$$C_{i,j} = \frac{1}{N_{events}} \sum_{events} \left[\left(\frac{dN}{d\eta} \right)_I - \overline{\left(\frac{dN}{d\eta} \right)_I} \right] \left[\left(\frac{dN}{d\eta} \right)_J - \overline{\left(\frac{dN}{d\eta} \right)_J} \right] \quad (8)$$

Where $\left(\frac{dN}{d\eta}\right)_I$ is the value of the $\frac{dN}{d\eta}$ distribution in the I^{th} η bin, and $\overline{\left(\frac{dN}{d\eta}\right)}_I$ is the value of the mean $\frac{dN}{d\eta}$ distribution in the I^{th} η bin. One important case of this formula is the diagonal elements, which have the formula

$$C_{i,i} = \frac{1}{N_{events}} \sum_{events} \left(\left(\frac{dN}{d\eta}\right)_I - \overline{\left(\frac{dN}{d\eta}\right)}_I \right)^2 = \sigma^2 \quad (9)$$

σ is just the regular definition of the standard deviation, so we expect the on-diagonal entries of the covariance matrix to simply be the square of the RMS (Figure 16).




Figure 16: An example of a covariance matrix for a single vertex bin, plotted on a regular (left) and log (right) scale. Notice that all the correlations are net positive (representing the overall correlation with the total multiplicity fluctuations of the events).

In order to calculate the χ^2_ν using these covariance matrices, we must invert them. To do this, we use a built in ROOT function, which uses a Bunch-Kaufman decomposition of the symmetric matrices to transform the inversion into 52 simultaneous vector equations, and then solves those using another built in ROOT function. While the correctness of this inversion is not checked systematically by the analysis modules, we can manually check that this inversion is being done correctly (Figure 17). We can see that the inverse looks like the identity, except for order 10^{-3} off-diagonal components.

This is likely due to numerical imprecision in the variables used to handle the matrices, but should not produce a noticeable effect on the final outcome.




Figure 17: Inversion of Covariance Matrices: Covariance Matrix (left), its inverse (middle), and their matrix product (right). We see that the product of the covariance matrix and its inverse is a matrix with 1 in the diagonal positions and order 10^{-3} off-diagonal components.

The verification toy model is generated by throwing 52 independent Gaussians, so we expect no correlations between the bins, and hence no off-diagonal elements in the covariance matrix. We set the RMS of each bin to 10 to generate the toy, so we expect the diagonal entries of the covariance matrix to be 100 ($=10^2$). We can see from Figure 18 that the covariance matrix for a toy model consists of a matrix with 100 on the diagonal (the square of the RMS, as we expect), and small off-diagonal components.

We can quantify “small” by looking at the mean of the off-diagonal elements in this covariance matrix as we increase the statistics. Figure 19 shows

the trend in these off-diagonal elements. We see that the off-diagonal component goes like $\frac{1}{\sqrt{N}}$, where N is the number of events; this behavior suggests that the off-diagonal component is purely statistical, and would go to 0 in the $N \rightarrow \infty$ limit. This allows us to conclude that the covariance matrices are behaving as we expect, the square of the RMS is on the diagonal, and the off-diagonal is zero except for statistical noise.




Figure 18: Sample Covariance matrix from the validation toy. The on diagonal elements are very close to 100, and the off-diagonal entries are small compared to the RMS. Left is a log plot with the absolute value of negative values plotted to show structure, and the right is a linear plot with the on-diagonal elements set to 0.

3.3.3 Calculating χ_ν^2

χ_ν^2 is defined as the χ^2 value divided by the number of degrees of freedom, and tells us, for a given event, how much each bin deviates on average. Usually, χ_ν^2 is defined for a set of random variables X_1, \dots, X_N with standard deviations σ_i and means X_i as:

$$\chi_\nu^2 = \frac{1}{N} \sum_{i=1}^N \frac{(X_i - \bar{X}_i)^2}{\sigma_i^2} \quad (10)$$

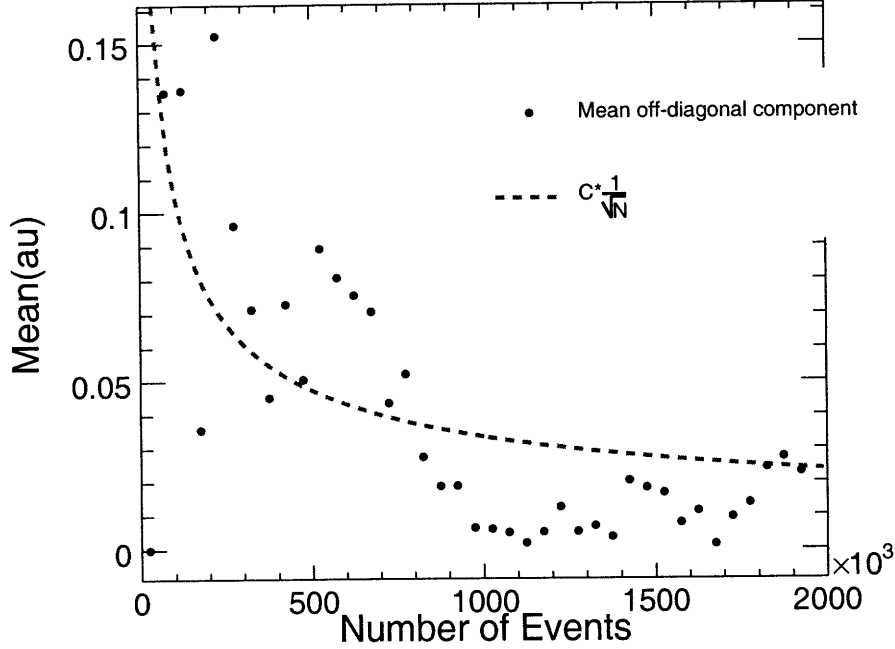


Figure 19: RMS of the off-diagonal elements of the covariance matrix as a function of the number of events. The dashed line is not a fit, but is drawn in to show the shape of a $\frac{1}{\sqrt{N}}$ distribution.

We will use a different definition to make use of our covariance matrices, but we will show that it is identical in the absence of correlations. One can see that χ_ν^2 should average to one, since the square of the variable's deviations should average to the square of σ_i . We will use this in Section 5 to justify a correction we will make to a χ_ν^2 distribution.

Using the inverted covariance matrices and the average $\frac{dN}{d\eta}$ distributions the χ_ν^2 can be calculated with global correlations removed. For each event, the χ_ν^2 is calculated as in Equation 11.

$$\chi_\nu^2 = \frac{1}{52} \sum_{I=1}^{52} \sum_{J=1}^{52} I_{I,J} \left[\left(\frac{dN}{d\eta} \right)_I - \left(\overline{\frac{dN}{d\eta}} \right)_I \right] \left[\left(\frac{dN}{d\eta} \right)_J - \left(\overline{\frac{dN}{d\eta}} \right)_J \right] \quad (11)$$

Here $I_{I,J}$ are the entries in the inverted covariance matrix for the vertex bin

of the event, and the 52s enter because there are 52 η bins.

From Equation 11, one can recover the more usual χ_ν^2 formula, for the case where there are no correlations. In the uncorrelated case, we have:

$$C_{i,j} = \begin{cases} \sigma_i^2 & I = J \\ 0 & I \neq J \end{cases}$$

and so obviously

$$I_{i,j} = \begin{cases} \frac{1}{\sigma_i^2} & i = j \\ 0 & i \neq j \end{cases}$$

In this case, Equation 11 simply becomes

$$\chi_\nu^2 = \frac{1}{52} \sum_{i=1}^{52} \frac{1}{\sigma_i^2} \left(\left(\frac{dN}{d\eta} \right)_i - \overline{\left(\frac{dN}{d\eta} \right)_i} \right)^2 \quad (12)$$

which is exactly the formula for χ_ν^2 for 52 uncorrelated random variables.

One additional thing that will be used repeatedly is the probability density function (PDF) for the χ_ν^2 of 52 uncorrelated Gaussian distributed random variables. By comparing our observed χ_ν^2 distribution to this, we will be able to estimate the rate of non-statistical fluctuations in our data, since the covariance matrix removed the correlations within our data. There is a closed form for this PDF given by [14] as:

$$PDF(\chi; r) = \frac{\chi^{r/2-1} e^{-\chi/2}}{\Gamma(r/2) 2^{r/2}} \quad (13)$$

where χ is the value whose probability we would like to know, r is the number of degrees of freedom, and $\Gamma(z)$ is the Gamma function. Figure 20 shows the distribution created by Equation 13. This plot is created by throwing 10^7

events according to Equation 13, so the fluctuations at high χ_ν^2 are purely statistical and can be ignored.




Figure 20: Random variable χ_ν^2 distribution. The fluctuations at $\chi_\nu^2 > 2$ are an artifact of the creation process.

We can produce the χ_ν^2 distribution for our validation toy model, to check how well it agrees with our random distribution. Figure 21 shows the χ_ν^2 distribution of the toy and of the random distribution. From visual inspection the shape agrees quite well, but when analyzed in detail, there is a shift in the mean of the toy distribution relative to the predicted random distribution. The random distribution has a mean of 1 (as we expect), but the toy has a mean of 1.008. This is a fairly large difference compared to the expected statistical error on the mean of this distribution (which is order 10^{-4}), and it is not clear where it comes from. We do, however, find that the agreement becomes excellent when we perform a multiplicative shift of the

toy χ^2_ν values to smaller χ^2_ν so that the mean works out to 1. This is not an ideal solution, but can be justified by looking at another distribution.

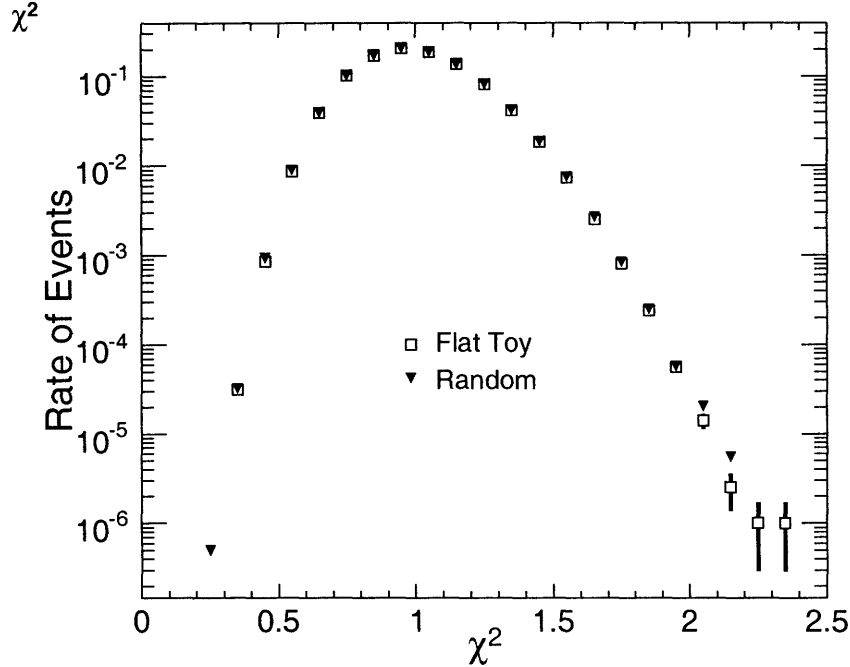


Figure 21: Flat toy model and random distribution

The χ^2_ν for an event can easily be converted to calculate the probability that that event could be a statistical fluctuation of the mean distribution. For example, an event which deviates greatly from the mean will have a large χ^2_ν value, and will hence be very unlikely to occur as a random fluctuation. On the other hand, an event which deviates only slightly from the mean will be very likely to have originated as a random fluctuation of the mean, so it will have a very high probability of occurrence. Over a large dataset, we can plot the probability distribution of the events, which shows how many events occur at each probability. For a large sample of uncorrelated statistically fluctuating events, we expect a flat probability distribution.

Figure 22 shows the probability distributions for the validation toy when

shifted as described above, and for the random model. Notice that the distribution for the random model is flat in probability, as we expect for uncorrelated statistically distributed random variables. The probability distribution for the shifted toy model is also flat, and agrees well with the random model. This is not the case for the unshifted toy (Figure 23). While this issue is being investigated, the method for this paper will be to shift the output distribution.

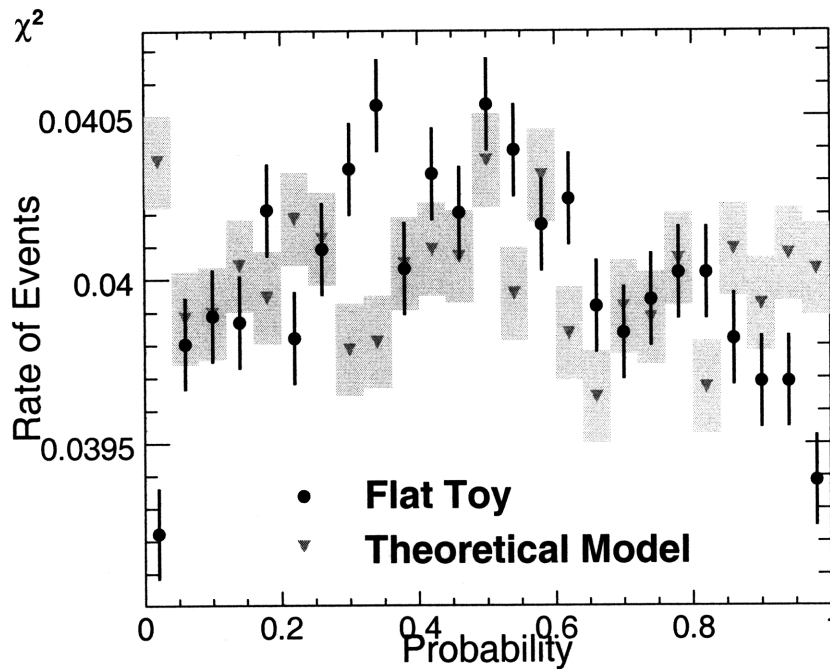


Figure 22: Probability distribution for the random (theoretical) model and the validation toy, when the χ^2_ν values are shifted so that the mean of the χ^2_ν distribution is one.




Figure 23: Probability distribution for the flat toy model without the multiplicative shift applied. The downward slope is caused by the shift to higher χ^2_ν values, but it is not clear what causes that shift.

4 Toy Monte Carlo

4.1 A More Advanced Toy Model

We have already mentioned the toy model we use to verify the the procedure, but a toy model can also be created which can be compared more closely to the data. To do this, we would like to include as many of the characteristics found in the data as possible: a $\frac{dN}{dn}$ distribution that looks like data, real vertex bin fluctuations, and real N_{hits} fluctuations. We expect this toy to behave much more like the data than the validation toy; the only things it misses which are in the data are correlations (which should be removed by the covariance matrix), and any non-statistical fluctuations. Since we

are looking for non-statistical fluctuations in the data, this toy will provide valuable insight into what is really occurring.

To generate an event in this toy model, we first pick a vertex bin according to the distribution actually seen in data. Once we have that, we generate the $\frac{dN}{d\eta}$ distribution for the event by taking the mean $\frac{dN}{d\eta}$ distribution calculated by the analysis module and, for each eta bin, offset the event from this mean by a Gaussian with RMS equal to the RMS calculated in data. Figure 24 shows an example of a toy $\frac{dN}{d\eta}$ distribution along with the $\frac{dN}{d\eta}$ distribution from data which is used to generate it. Finally, we choose an N_{hits} from the distribution in data for the chosen vertex bin and scale the event so that its total integral matches this N_{hits} value (note that Figure 24 has this scaling removed).

The covariance matrix for this toy will be more complicated than that for the flat toy used for validation, since we impose scaling on the event distributions. While we throw each of the bins using independent Gaussians, we then scale them to fixed values in total integral, which imposes correlations between the bins. Since the individual bins in the $\frac{dN}{d\eta}$ can fluctuate more than the N_{hits} distribution we have a situation where a big positive fluctuation in one bin suppresses other bins, in order to meet the total integral requirement. This actually will impose small negative correlations between the bins (small since there are big enough fluctuations to cause this effect in only a subset of the events), so we expect the off-diagonal component of the covariance matrix to be slightly negative. Figure 25 shows the covariance matrix for a single (representative) toy vertex bin. We see that the on-diagonal entries dominate the covariance matrix, which we expect as the correlations should be small. The on-diagonal entries are no longer uniform, as they were in Figure 18,




Figure 24: An example of a toy event $\frac{dN}{d\eta}$ distribution and the mean $\frac{dN}{d\eta}$ distribution from data on which it is based (toy and event are scaled to the same integral).

since the RMS values are now no longer the same for every $\frac{dN}{d\eta}$ bin. We also see in the figure the negative correlations between the different η bins, which is stronger for the bigger bins, since their variance is greater.

The χ^2_ν output of these toy events must be compared against Equation 13 and Figure 20 to fully understand what the toy is telling us. Figure 26 shows the χ^2_ν distributions for this toy and the random χ^2_ν distribution. The agreement is fairly good over most of the range, roughly $\pm 2\sigma$. Since our model includes correlations and various other dynamics which are not found in the random model, we could not expect perfect agreement.




Figure 25: Covariance matrix for a sample toy vertex bin. The full covariance matrix is shown on the left, while the same matrix with the diagonal entries set to zero is shown on the right.




Figure 26: χ^2_ν distribution for the toy derived from data and the random distribution (top) and the difference between the toy and the theoretical model divided by the standard deviation of the toy.

4.2 Spike Model

We have said that we are sensitive in this analysis to deviations in the shape of the $\frac{dN}{d\eta}$ distribution, but we would like to have a toy model that quantifies that somewhat more precisely. One possible shape fluctuation is a “spike”, or a consecutive set of η bins which exceed the mean distribution. This could be caused by any number of physical phenomena that cause bursts of particles to be detected and so we would like to understand our sensitivity to it. One can refer to Figure 6 for the average number of particles in each bin, but it averages roughly 70, with spikes to 120-130 particles. This behavior

is essentially what is reproduced in the toy model described above, so if we modify that to add in spikes, we can get a sense for our sensitivity.

We try to add spikes in a way that would best represent the data, so we introduce them only in a fraction of the events and we keep the total spike size small compared to the average bin contents. We generate 2×10^6 events and add spikes to 0.01% of them. These spikes are 3 consecutive η bins, to which we add a total of 20 hits (5 in the left and right bins, and 10 in the middle), and we randomly position these spikes in their central η . This represents an enhancement of at most 30% in the rings, and as little as 4% at the peaks in the octagon.

We run these events through the same analysis procedure, and compare them to the distribution we found for the toy model. We apply the same multiplicative shift to center the mean as before, and we are looking for an enhancement in the region of events with high χ_ν^2 . Figure 27 shows the χ_ν^2 distribution for this spike model, compared with the distribution for the toy model described above. We see a pronounced enhancement in the high χ_ν^2 region compared to the other toy. The slight suppression of the spike model relative to the regular toy is an artifact of the normalization, the crucial aspect is the enhancement in the region $\chi_\nu^2 > 2$. This plot dramatically shows that we are very sensitive to fairly small spikes in the $\frac{dN}{d\eta}$ distribution.

4.3 What is Missing From the Toy Models

The toy model is a very useful tool for understanding our data, but there are fundamental areas where it will be different. The most obvious, and most important, is in the presence of non-statistical fluctuations. Physically inter-




Figure 27: χ^2_ν distribution for the toy model based on data from section 4.1 and the toy model with spikes added.

esting events will be characterized by some sort of non-statistical fluctuation in the data; since our toy models have no physics beyond the N_{hits} fluctuations built into them, we will never see any events like that come into the toy.

Another element that is missing from the toy is more subtle, and has to do with the effectiveness of our covariance matrix in removing correlations. The covariance matrix is a very effective tool for removing global correlations, which show up in the majority of events. It removes from the data the overall correlation with N_{hits} , and other systematic correlation within the detector. What it cannot do, however, is remove correlation which show up in a minority of the events. One could imagine electrical noise causing some artifact in a small subset of a run, for example. This would create

correlations in a small number of events within the data set, but would not be visible within the global data set. There could be very slight differences between the runs themselves, which would show up as a global correlation for a specific run, but not for the data set as a whole.

5 Results

5.1 Statistical Event Distribution

The ultimate goal of this analysis is to put an upper bound on the rate of anomalous events occurring in the data set. To do this, we first will compare the data to our random model and our toy models to understand how well the the non-anomalous events conform to the statistical expectations. Figure 28 shows the χ_ν^2 distribution for 0-3% central data and the χ_ν^2 for the random events, governed by Equation 13. We can see that the bulk of the data falls in the region $0.3 < \chi_\nu^2 < 2$, and in that region the data follows a very similar shape as the random distribution. This suggests that, once correlations are discounted, most of the fluctuations observed in data are statistical. This is, in itself, an important result, because it says that we really can ask the question how often non-statistical fluctuations occur in the data. One could imagine the data having a χ_ν^2 distribution which agreed poorly, or not at all, with a statistical distribution, which make this question much harder to answer.

In order to better illuminate the match between data and the random model, we can examine the difference between the two throughout the range. Figure 29 shows the difference between the χ_ν^2 distribution from data and that for the random distribution in units of the statistical error on the data. Comparing Figures 28 and 29, we see an enhancement in the region $1.8 < \chi_\nu^2 < 2.4$, where we would like to look for rare events, but a much bigger effect in the enhancement and suppression in the region $0.3 < \chi_\nu^2 < 1.8$. Before we can understand the region that may correspond with physically interesting events, we must focus on the region $0.3 < \chi_\nu^2 < 1.8$, where the statistical




Figure 28: χ^2_ν distribution for 0-3% central events and that for the random distribution. Note that error bars are purely statistical.

events live.

Figure 29 shows that most of the random distribution falls within $1 - 2\sigma$ of the data distribution, except for some areas that deviate more. One also notices that there is a clear trend of enhancement of the event distribution over the random on the low side of $\chi^2_\nu \approx 1$, and suppression on the high side. This suggests that there is some shift between the two distributions, with the peak of the event distribution falling at slightly lower χ^2_ν than the peak of the random distribution. While it is unclear what causes this shift, we do observe that the mean of the data distribution is 0.99811 and the mean of the random distribution is 0.99996. If we manually correct this shift, by simply multiplying the χ^2_ν values of the data events by the ratio of the means, we get a χ^2_ν distribution which is visually indistinguishable from Figure 28, but




Figure 29: The difference between the data and the random distributions in units of the standard deviation of the data.

a difference distribution which looks quite different.

Figure 30 shows the difference plot for this shifted χ_ν^2 value and the random distribution, and shows that the agreement is somewhat better. The enhancement in the $1.8 < \chi_\nu^2 < 2.4$ region is still evident, but the differences in the region $0.3 < \chi_\nu^2 < 1.8$ have been suppressed. The events now fall much more in the $\pm 1\sigma$ range, suggesting agreement between the data and the random model. The residual pattern suggests that the widths of the distributions may not match perfectly, but the agreement is still better than without the shift. This supports the notion that there is some small, but systematic shift in the χ_ν^2 values of the data. It is not clear what is causing this shift, but it appears to be a non-negligible effect.

Comparing Figures 29 and 30, we also notice that the enhancement in

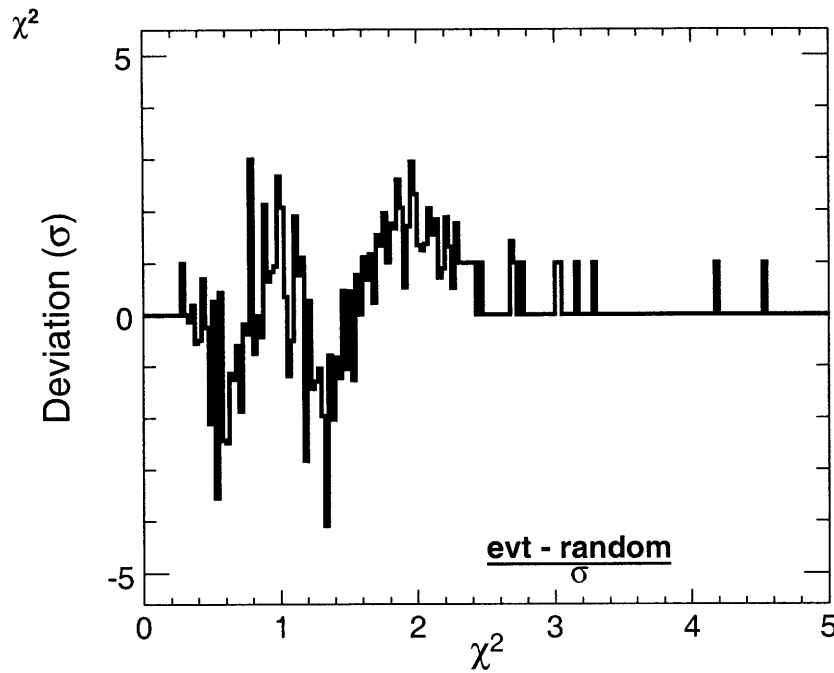


Figure 30: The difference between the shifted data and the random distributions in units of the standard deviation of the data.

the region $1.8 < \chi_\nu^2 < 2.4$ is not substantially affected by the shift. This is encouraging, since it suggests that this enhancement is a real effect, and not arising as the result of some offset in the χ_ν^2 values. The shape of the deviation in this region is also largely unaffected, suggesting that there is some real difference between the events in the two regions. Also notice that the region of $\chi_\nu^2 > 2.2$ is unaffected by the shift, and there is some clear excess in that region. This again suggests that we are seeing a signal on top of the random distribution.

Figure 31 shows the probability distributions for the data (with the mean shifted as described above) and the random distribution. We can see that the random distribution does, in fact, have a flat probability distribution as expected. The data follows a nearly flat distribution for most of its range,

consistent with its good agreement with the random distribution, but deviates sharply for low probability events. This corresponds to the deviation we see in the $1.8 < \chi^2_\nu$ region. This indicates that our data does contain a non-statistical number of events which have a very low probability of occurring. These are the events we would like to further understand and put an upper bound on the rate of.




Figure 31: Probability distributions for (shifted) data and random distributions. The random is scaled to total integral 1, and the data is scaled to have the same integral in bins 2 to 100 as the random. This is done to allow easier visual comparison.

5.2 Rare Events

Having seen the good agreement between the shifted data and the random model, we can conclude that the bulk of the data behaves like statistical fluctuations of the mean distribution. We can, therefore, endeavor to understand

the signal that arises as a deviation from the random distribution. Figure 32 shows the probability distribution from data for small probabilities. We see that it is still very flat all the way down to a probability of 1×10^{-6} , which corresponds to $\chi^2_{\nu} \approx 2.22$. We run very quickly into statistics problems when we try to reduce beyond this limit, so we will define a “rare event” as any event which has $\chi^2_{\nu} > 2.22$. We can see immediately from Figure 32 that there are roughly 28 such events, although we will be more precise about this shortly.




Figure 32: Probability distribution for low probability shifted data events.

By successively zooming the plot to smaller probability values, we can more easily see where our rare events live. Figures 33 and 34 show the probability distributions in the ranges $(0, 2 \times 10^{-3})$ and $(0, 3 \times 10^{-4})$ respectively. We see that, even at lower probability than Figure 31 can resolve, the probability distribution is flat, except for an enhancement in the lowest bin.

Comparing Figures 33 and 34, we see that the enhancement in the low bin in Figure 33 is larger than the enhancement in the low bin of Figure 34. These extra events show up in the small enhancement in the bottom 7 bins in Figure 34 (each bin in the low 15% of the first figure is 7 bins in the second). Indeed, if we integrate the first 7 bins, we find 84 events, which corresponds exactly with what we expect from the previous figure. This discussion is simply to show that the figures are consistent, the real message of these two figures is that there is a true signal of events at very low probabilities.




Figure 33: Probability distribution for data in the range $(0, 2 \times 10^{-3})$, with 100 bins. The distribution is (within error bars) flat, except for the bins at lowest probability.

To make a quantitative estimation of the excess of low probability events, we want to extrapolate the trend in the higher probability events to the data. Figure 35 shows the full probability distribution (as in Figure 31), with 2 fits




Figure 34: Probability distribution for data in the range $(0, 3 \times 10^{-4})$, with 100 bins. This represents the bottom 15% of Figure 33.

applied, which we will use to estimate the number of events expected in the range $0 < prob < 10^{-6}$. The first fit is a constant fit, which predicts 1.238 ± 0.001 events in this range. Based on the random model, we expect the probability distribution to be flat, so a flat fit should agree with the statistical events on theoretical grounds.

We also do a quadratic fit to the data in Figure 35, which is motivated by looking at the shape of the probability distribution, and seeing that there is some overall curvature. While this is not very well understood, and does not agree with our models, if we simply fit the observed data, we can get another estimate of 1.221 ± 0.003 . Even though we fit a quadratic, the parameters of the fit show that the linear and quadratic part of the fit are rough 8% of the constant part, so we are still dominated by the constant part, which is

at least partially consistent with the random model.




Figure 35: Number of events per probability bin (bin size: 10^{-2}), with a constant and quadratic fits to bins 2 through 100.

We perform a similar fit procedure to the lowest 1% of the distribution in order to get additional estimations of the expected rate. Figure 36 shows the 1% of the events with lowest probability, along with a constant fit to all but the lowest bins. Figure 32 suggests that most of this excess actually occurs in the very lowest portion of the first bin, showing 28 ± 5 events in the region $0 < prob < 10^{-6}$. Using our fit, and scaling appropriately, the fit predicts 1.26 ± 0.01 events in the range $0 < prob < 10^{-6}$.

Averaging the three different predicted values, we get 1.239 ± 0.016 expected events in the probability region $0 < prob < 10^{-6}$. The data shows 28 ± 5 events, which is more than 5σ different from the expected value. It is worth noting that this is a fairly conservative estimate for the expected




Figure 36: Number of low probability events per probability bin (bin size: 10^{-4}), with a constant fit to bins 2 through 100.

number of events. Figure 34, for example, suggests that we have a slight increase in our baseline number of events going toward low probability. We are, however, looking for an upper bound on the overall rate, so we use a conservative estimate here, and anticipate that our final rate will be somewhat of an overestimation of the true rare event rate. We will convert all of this into an event rate shortly, but this shows that there is a signal of low probability events that is clearly non-statistical. We will include the error on this prediction due to the multiple different fits, as a systematic error on the final excess rate.

Our data sample consists of 1.26×10^6 0-3% central Au+Au events that survive all the event selection and pileup cuts. Out of this sample we see 28 events that qualify as rare events under our definition. Table 1 summarizes

the number of events above various χ_ν^2 cutoff values, as well as the range this covers in probability. We can see a clear set of core events remaining down to very small probability values, once again underscoring the non-statistical nature of these rare events. The table also shows the rate of the events in the data set with statistical errors. As discussed above, we expect 1.239 ± 0.016 events from statistical fluctuations, so we must subtract this from the number of events at $\chi_\nu^2 > 2.22$ to get the excess event rate; this leaves us with an excess rate of 26.76 events. We can estimate the additional systematic error introduced by the shift in the mean of the χ_ν^2 distribution by comparing the number of events in each range before and after the shift, and seeing how this changes the rate. Since the χ_ν^2 shift is a systematic correction on our data, we include it as a systematic error on the final rate of rare events in the data and arrive at a value for the rate of rare events: $(2.13 \pm 0.4(stat) \pm 0.1(sys)) \times 10^{-5}$, which we will endeavor to improve upon.

χ_ν^2	Probability range	# of Events	Excess Rate ($\times 10^{-5}$)
2.0	$[0, 2.5 \times 10^{-5})$	92	7.2 ± 0.8
2.2	$[0, 1.4 \times 10^{-6})$	28	2.2 ± 0.4
2.22	$[0, 1.0 \times 10^{-6})$	28	2.2 ± 0.4
2.4	$[0, 6.4 \times 10^{-8})$	13	0.9 ± 0.3
2.6	$[0, 2.5 \times 10^{-9})$	10	0.7 ± 0.2
2.8	$[0, 8.4 \times 10^{-11})$	7	0.6 ± 0.2
3.0	$[0, 2.5 \times 10^{-12})$	7	0.6 ± 0.2
4.0	$[0, 1.5 \times 10^{-20})$	3	0.1 ± 0.1

Table 1: Number of Events exceeding each listed χ_ν^2 value

Before looking at the individual events remaining in the sample, it is useful to get a sense of how big the deviation in these rare events is. Recall from Section 4.2 that we generate a toy model with spikes added on top of some of the events to simulate rare events like $\frac{dN}{d\eta}$ fluctuations. The param-

eters we used were that 0.01% of events got a spike of 20 particles spread over 3 η bins. We can compare this to what we find in the data to see how our rate and average fluctuation size compares to this. Figure 37 shows the χ^2_ν distributions for data and the toy model. One can see that the spike model actually overestimates both the rate and the size of the deviations observed. We end up with roughly an order of magnitude less events than in the spike model, and the deviation of the events in data are generally smaller than those in the spike model.




Figure 37: χ^2_ν distribution for data compared to that for the spike model (toy model with spikes added)

We can also compare the probability distributions for the data and toy model, to see the relative size of the low probability peak. Figure 38 shows that our model does indeed dramatically overestimate the number of rare

events. We see the roughly order of magnitude overestimation evident in Figure 37. We also see fairly good agreement between the spike model and our data in all other bins, even mirroring the slight upward trend in data at lower probabilities. In our toy model we added spikes to 0.01% of the events, or 1 in 10^{-4} events. We see that this is actually an order of magnitude higher than the rate of rare events we see in the data so our spike model actually matches the data very well. There are 198 events in the low bin of Figure 38, which implies that virtually every event with a spike is being detected as a rare event. This suggests that we are very sensitive to spikes of around 20 particles. Furthermore, in Figure 37 we can see that the rate of events in the spike model with $\chi^2_\nu > 3.2$ is much greater than that observed in data, so if there are spikes in data of around 20 particles, we would actually expect a slightly more extended tail than we see. Based on this, we can conclude that we are sensitive to spikes of slightly less than 20 particles in data.

The number of rare events found is a small enough number to enable an event by event analysis of the rare events, which we carry out to determine if there are additional exclusions we can make or additional conclusions about the nature of the events. Looking at the raw detector output for these events it is immediately apparent that for some events there are simple detector malfunctions that influence the output. In one of the events, for example, the negative rings failed to record any hits. An example event display for this type of event is included in the appendix in Figure 39. In total, there are 2 events which suffer from this blatant detector failure.

While we will explore other detector effects which might cause these rare events, these obvious failures are the only ones which we can exclude to a high degree of certainty. Both of the detector failure events had $\chi^2_\nu \gg 3$




Figure 38: Probability distribution for data compared to that for the spike model

so, excluding these events, we have 26 rare events, which leads to an excess rate of rare events:

$$(1.97 \pm 0.4(stat) \pm 0.1(sys)) \times 10^{-5} \frac{\text{rare events}}{\text{event}}$$

We have been quite conservative here in our choice of which events to eliminate, and our estimation of the baseline rate, so we expect this rate to be an overestimation. We are looking for an upper bound on the actual rate, so this is our final result, but it is also interesting to take a brief look at one way we could be more aggressive in this rate.

One type of event, which shows up often in the rare events signal, is an

event which has individual pad subdetectors with high signal in the rings. These events are characterized by having every pad in the subdetector hit, while other pads show only mild or no enhanced signal. The hit merging routine uses a dead channel map to normalize the signal in a pad detector based on the energy value of unhit pads. It is possible that, if there are no unhit pads, the dead channel correction does not work, which could lead to a deviation which has no physical significance. An example of the readout in this type of event is found in the appendix in Figure 40. We cannot, however, determine that these events are actual malfunctions with the same level of rigor as the clear failures, so we do not remove these events from the final rate calculation. There are 6 of these events in the rare events data set, so if we could remove them, the excess rate of rare events would be reduced to $(1.49 \pm 0.5(stat) \pm 0.1(sys)) \times 10^{-5}$ rare events per event.

Finally, Figure 41 shows an event that has no detector malfunction or subdetector pathology. This event does exhibit a spike in the rings, but the problem of Figure 40 does not occur. This is the primary contribution to the high χ^2_ν value, but we also see some significant deviations in other bins, especially around the positive boundary between the octagon and the rings. This is the sort of event which could actually encode some interesting phenomena, and is the type of event we are seeking in this analysis.

6 Conclusion

This paper has endeavored to put an upper limit on the production of non-statistically fluctuating events in the PHOBOS 0-3% central Au+Au data. We have seen two examples of types of physical phenomena whose rate might be constrained by this analysis, but there are surely many other interesting and exotic ideas that could be tested with this analysis. It is not the purpose of this paper to discuss the production rate of specific phenomena, but rather to make a general statement about the production rate of a class of phenomena.

The cuts used in this analysis were selected to balance removing non-physical fluctuation with avoiding accidentally removing potential physics. The event quality selections were left as lenient as possible, only removing those events that belonged to known problems. The pileup removal was necessarily a more subtle process, but was ultimately very successful in removing events whose signatures strongly suggested that they were pileup. None of these cuts were designed to specifically remove events with high χ_ν^2 values, only those which matched a well understood picture of what pileup should look like.

The analysis chain was very successful in removing the large correlations that exist in the data, evidenced by the close agreement between the data and the random distribution. Somewhere in the analysis, however, a shift in the mean of the output χ_ν^2 distribution occurred, made more perplexing by the fact that it shifted the means of the toy model and the data in different directions. It could be that the shift in the data is a real effect: some non-global correlation actually shifts events more towards the mean. This would

be a very surprising result, and it would require some reinterpretation of our results, but there is nothing to suggest that the shift is, in fact, real. The origin of this shift is not well understood but, when it is removed, the data and toy agree exceptionally well with the random distribution.

Accepting this correction, we are able to arrive at two very important conclusions about the data: that most of the events look like statistical fluctuations about the mean, and a cutoff and rate for rare events. The agreement between the data and the random distribution in the $0.3 < \chi_\nu^2 < 1.8$ region strongly indicates that most events in the dataset simply look like statistical fluctuations about their mean. This effect is masked by the existence of correlations in the data set, but once those are eliminated the effect becomes obvious.

With this realization one is able to identify where the data begins to deviate from the random distribution and then use that to establish a rate. We see good agreement up until $\chi_\nu^2 \approx 2.22$, so we define this as the cutoff for a rare event. We then eliminate events that have clear detector pathology, and are able to come up with a value of the rate of rare events in the detector: $(1.97 \pm 0.4(stat) \pm 0.1(sys)) \times 10^{-5} \frac{\text{rare events}}{\text{event}}$. Our analysis with the spike model overestimates this rate by an order of magnitude, which is very close to the excess rate at which 20 particle spikes were added to the data. Based on this, we can conclude that this analysis is sensitive to spikes of slightly less than 20 hits.

This rate is a useful metric to test theories of phenomenology concerning the PHOBOS detector and RHIC events in general. While this analysis is not sensitive to all possible fluctuation events that may occur in the detector, it is sensitive to a substantial fraction of potential events. It will also be

very interesting to see how this value compares to the rate that might be found at future experiments at the LHC, for example. A similar analysis for other existing experiments might provide a very interesting trend for how anomalous event production scales with center-of-mass energy.

References

- [1] P. Sarin, *Measurement of charged particle multiplicity distributions in Au+Au collisions up to 200GeV*, PhD thesis, MIT, 2003.
- [2] S. Povh, Rith and Zetsche, *Particles and Nuclei* (Springer-Verlag, 2008).
- [3] B.Back *et al.*, Nuc. Phys. A **757**, 28 (2005).
- [4] A. P. French, *An Introduction to Quantum Physics* (W. W. Norton & Co., New York, New York, 1978).
- [5] B. Back *et al.*, Nuclear Instruments and Methods in Physics Research A **499**, 603 (2003).
- [6] D. Kharzeev *et al.*, Phys. Lett. B , 121 (2001).
- [7] R. Jaffe, W.Busza, J. Sandweiss, and F. Wilczek, Rev. Mod. Phys. **72**, 1125 (2000).
- [8] J. Tanaka *et al.*, Eur. Phys. J. **C41**, 19 (2005).
- [9] H. Stoecker, Int. J. Mod. Phys. **D16**, 185 (2007).
- [10] M. Peskin and D. Schroeder, *An Introduction to Quantum Field Theory* (Westview Press, 1995).
- [11] K. Rajagopal and F. Wilczek, Nucl. Phys. B **404**, 577 (1993).
- [12] T. Nakamura, Study of isospin fluctuations at RHIC-PHENIX, Quark Matter Proceedings, 2002.
- [13] B. Alver *et al.*, Nucl. Phys. A **774**, 113 (2006).
- [14] K. Kenney and E. Keeping, *Mathematics of Statistics, Pt. 2*, 2 ed. (Van Nostrand, 1951).

7 Appendix

7.1 Report Card Examples

Report cards are a combination of various single value parameters for the event, the raw hit arrays for the octagon and rings, and portions of the $\frac{dN}{d\eta}$ distribution from various detectors. Figures 39 to 41 show some report cards for a few events. The box on the top left contains information about the energy measured by different detectors for the event. The top middle box contains the vertex position as measured by different detectors. The paddle vertex detection is often very different from the others, but if other vertex positions disagree there is probably a problem that wasn't picked up by the event selection for some reason. The final box contains the χ^2_ν value and some information about the time when the event occurred.

The 2D plots show the raw hit array from the negative rings (left) the octagon (center) and the positive rings (right). One can see the holes in the octagon for the spectrometer and vertex in the center plot. The plots in the bottom layer are the ϕ integrals of these plots, the $\frac{dN}{d\eta}$ distributions from different detectors. On the right of the bottom row is the event plotted on top of the mean $\frac{dN}{d\eta}$ for the corresponding vertex bin. In the very bottom row in the center is a plot of the number of sigma that the event deviates from the mean in units of the standard deviation.

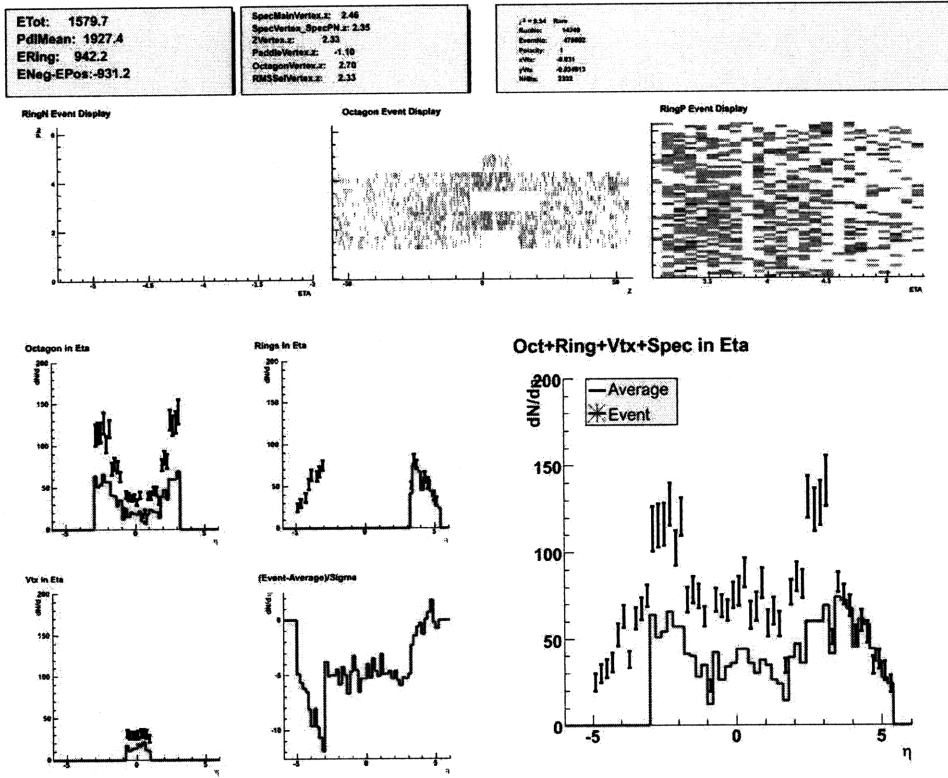


Figure 39: Report Card for a detector failure event. Notice that the negative ring and several layers of the octagon failed to register any hits. There is no physical reason for an event to look like this, so this is clearly a failure in the readout system.

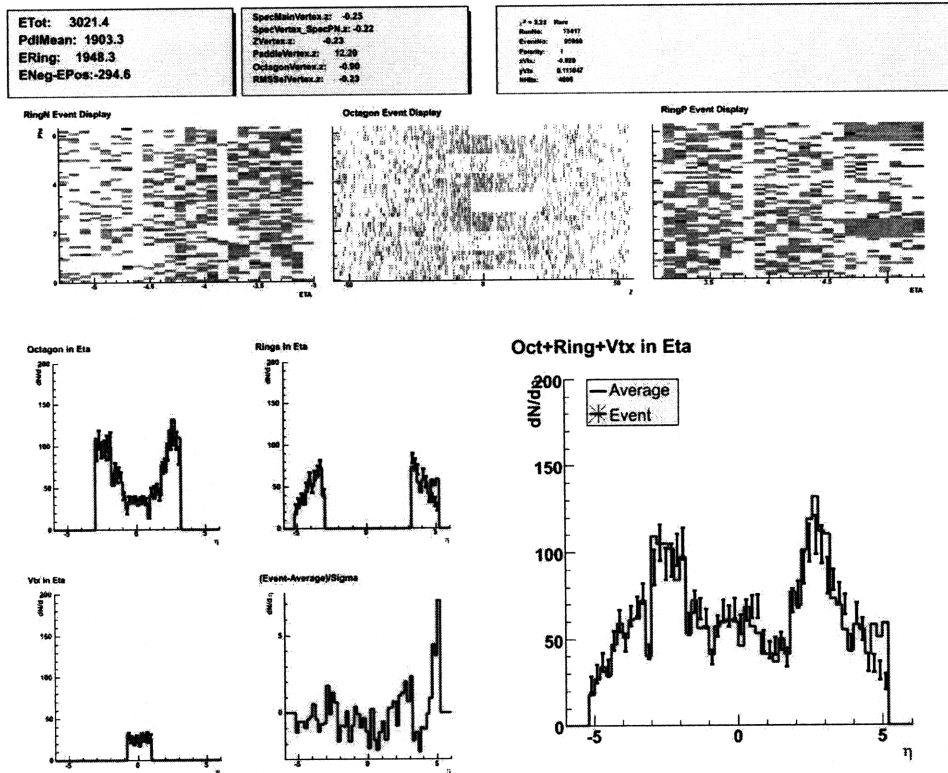


Figure 40: Report Card for possible pad saturation. One notices the spike in the event in the region $4.6 < \eta < 5.4$ which is evident in the events $\frac{dN}{d\eta}$. The $\eta - \phi$ plot shows that 2 wafers are hit in every pad in this region. We suspect that this may cause a problem with our subdetector normalization, but we cannot rule out some physics explanation.

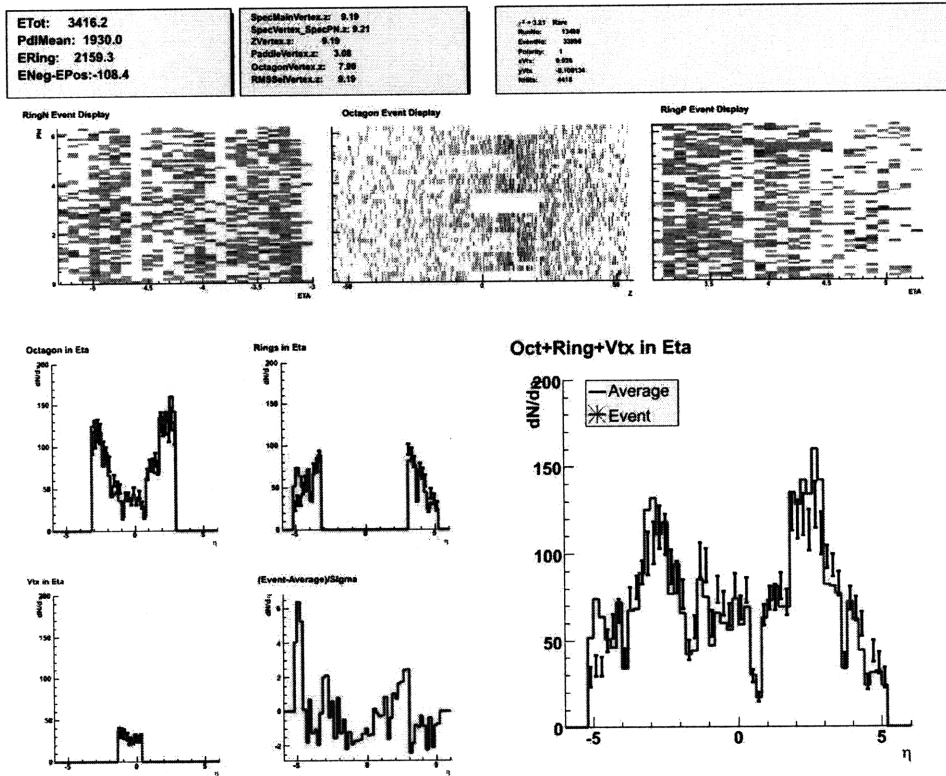


Figure 41: A rare event which shows no obvious detector pathology. There is a spike in the negative rings, but no obviously lit up subdetector indicating a malfunction. This event could be a real non-statistical fluctuation.