

UCLA

UCLA Electronic Theses and Dissertations

Title

Exploring the QCD Phase Diagram at STAR: Measurement of Proton Fluctuations in Azimuthal Partitions

Permalink

<https://escholarship.org/uc/item/Ojq1h2s1>

Author

Neff, Dylan

Publication Date

2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Los Angeles

Exploring the QCD Phase Diagram at STAR:
Measurement of Proton Fluctuations in Azimuthal Partitions

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Physics

by

Dylan McKenzie Neff

2023

© Copyright by
Dylan McKenzie Neff
2023

ABSTRACT OF THE DISSERTATION

Exploring the QCD Phase Diagram at STAR:
Measurement of Proton Fluctuations in Azimuthal Partitions

by

Dylan McKenzie Neff

Doctor of Philosophy in Physics

University of California, Los Angeles, 2023

Professor Huan Z. Huang, Chair

Quantum ChromoDynamics (QCD) is a corner stone of the Standard Model. In the non-perturbative regime first principle calculation is very difficult but phenomenology is rich. The many-body dynamics of quarks interacting in a deconfined state, Quark Gluon Plasma (QGP), produces emergent phenomena which are not manifest in the theory's Lagrangian. QGP can be created in relativistic heavy ion collisions which provide an opportunity to experimentally probe this exotic QCD matter. Of particular interest is the structure of the QCD phase diagram. While the transition from QGP to hadron gas is a cross-over at high collision energies, RHIC's beam energy scan program covers the high baryon density region of the phase diagram to search for signs of a first-order transition and its accompanying critical point.

Bubbling is a general characteristic of first-order phase transitions and may lead to clusters of quarks which coalesce into final state protons in heavy ion collisions. We construct a new observable, $\Delta\sigma^2$, to search for excess clustering in the azimuthal distributions of identified proton tracks. We show that the effects of detector inefficiency and elliptic flow can be

corrected and that the observable behaves as expected in various models and simulations. $\Delta\sigma^2$ is measured in STAR's Beam Energy Scan I data set and a strong repulsive interaction among protons is observed. The repulsion is found to depend heavily on the event multiplicity, increasing dramatically in magnitude as the event multiplicity decreases. Energy dependence, absent in model calculations, is observed in STAR data, suggesting a possible scenario of an energy dependent attractive correlation beneath a repulsive background.

The dissertation of Dylan McKenzie Neff is approved.

Michail Bachtis

E.T. Tomboulis

Zhongbo Kang

Huan Z. Huang, Committee Chair

University of California, Los Angeles

2023

To my friends and family

TABLE OF CONTENTS

1	Introduction	1
1.1	The Standard Model	1
1.2	Physics Beyond the Standard Model	3
1.2.1	Unification of the Quantum Theories	3
1.2.2	Gravity	4
1.3	Quantum Chromodynamics (QCD)	5
1.3.1	The QCD Lagrangian	5
1.3.2	Perturbative QCD	6
1.3.3	Limitations of Perturbative QCD	7
1.3.4	Asymptotic Freedom and Color Confinement	8
1.3.5	Lattice QCD	9
1.3.6	Quark Gluon Plasma (QGP)	11
1.3.7	The Phase Diagram of QCD	13
1.4	Relativistic Heavy Ion Collisions	19
1.4.1	Stages of a Heavy Ion Collision	19
1.5	Experimental Signatures of a Critical Point	21
1.5.1	Critical Point Search	22
1.5.2	Signals of a First-Order Phase Transition	24
2	RHIC and the STAR Detector	26
2.1	RHIC	26
2.1.1	Vacuum Arc Ion Source	26

2.1.2	Electron Beam Ion Source (EBIS)	27
2.1.3	Accelerating Synchrotrons	28
2.1.4	Beam Energy Scan	32
2.2	The STAR Detector	32
2.2.1	Time Projection Chamber (TPC)	34
2.3	Time Of Flight (TOF)	37
2.4	Trigger and Data Acquisition (DAQ)	40
2.5	Collision Vertex Reconstruction	41
2.6	Centrality Determination	42
2.7	Pile-Up	44
3	Fluctuations in Azimuthal Partitions	46
3.1	Azimuthal Partitions	46
3.2	Demonstrating Convergence to Binomial	48
3.3	Event Resampling	56
3.4	Statistical Error Estimation	61
3.4.1	Delta Theorem	61
3.4.2	Bootstrap	61
3.4.3	Demonstrating Consistency of Statistical Error Estimation	64
3.4.4	Statistical Error Estimation Conclusion	76
3.5	Bias in Estimators	80
3.5.1	Unbiased Estimator of Variance	82
4	The $\Delta\sigma^2$ Observable	83

4.1	Variance of Multiplicity in Azimuthal Partitions	83
4.2	Defining the Observable	84
4.3	Analytical Analysis of $\Delta\sigma^2$	85
4.3.1	A Single PDF	86
4.3.2	Two Independent PDFs	92
4.4	Elliptic Flow Correction	97
4.4.1	Exact Integration	98
4.4.2	Numerical Integration	100
4.4.3	Event by Event Simulation	104
4.4.4	Elliptic Flow Closure Test Conclusion	106
4.5	Mixed Event Correction	107
4.5.1	Numerical Integration	108
4.5.2	Event by Event Simulation	111
4.5.3	Efficiency Correction Closure Test Conclusion	113
4.6	Combined Efficiency and Flow Correction	115
5	Proton Fluctuations in Azimuthal Partitions at STAR	117
5.1	Analysis of $\Delta\sigma^2$ in STAR BES-I	117
5.1.1	Event Multiplicity Dependence	125
5.1.2	Azimuthal Partition Width Dependence	129
5.1.3	STAR Analysis Conclusion	131
5.2	Gaussian Correlation Model	132
5.2.1	$\Delta\sigma^2$ Analysis	134
5.3	STAR Data Set	142

5.3.1	Event Level Cuts	142
5.3.2	Track Level Cuts	143
5.4	Elliptic Flow Measurement	149
5.4.1	Azimuthal Distribution Flattening	149
5.4.2	Measuring v_2	151
5.4.3	STAR Proton v_2 Measurement	152
5.4.4	AMPT Proton v_2 Measurement	152
5.4.5	MUSIC+FIST Proton v_2 Measurement	153
5.5	Signed Distance of Closest Approach Bad Events	154
5.6	Pile Up Rejection	158
5.7	STAR Systematics	162
6	Conclusion and Outlook	168
A	Simulations	170
A.1	Rejection Sampling	171
A.2	Homogeneous Efficiency	171
A.3	Azimuthal Dependent Efficiency	171
A.4	Flow	172
A.5	Simple Clustering Model	172
A.6	Momentum Conservation Model	174
B	Models	177
B.1	AMPT	177
B.2	MUSIC+FIST	178

LIST OF FIGURES

1.1	Constituents of the Standard Model of particle physics.	2
1.2	An illustration of the electromagnetic, weak and strong coupling constants unifying at high energy under a supersymmetric theory.	3
1.3	The running of the strong coupling constant as a function of the momentum transfer [6].	7
1.4	The phase diagram of QCD with temperature on the y-axis and baryon chemical potential on the x-axis [14]. The true structure of the phase diagram is not known, this is a simplified depiction of how it may look.	12
1.5	Phase diagram of the Ising model [14].	14
1.6	Mapping of the Ising phase diagram onto that of QCD.	16
1.7	Lattice QCD calculations for thermodynamic quantities demonstrating that the phase transition from hadron gas to QGP is a cross-over at low μ_B	18
1.8	Visualization of the stages of a heavy ion collision.	19
1.9	The normalized skewness and kurtosis of net-proton multiplicity distributions at mid-rapidity are plotted as a function of center of mass collision energy. The non-critical UrQMD and HRG models are shown along with data from peripheral collisions as baselines.	24
2.1	An aerial view of the RHIC complex.	27
2.2	Left: Diagram of the accelerator complex leading to the RHIC ring. Right: Diagram of the RHIC rings.	28
2.3	A rendering of the STAR detector as it was configured during the 2010 and 2011 Beam Energy Scan I.	33
2.4	Left: Diagram of STAR's TPC. Right: Pad row layout for one TPC sector.	35

2.5	Energy loss per centimeter measured in the TPC as a function of the rigidity (momentum divided by charge) for 39 GeV BES-I data. Theoretical Bichsel curves are drawn as black lines for each particle species. Identified pions, protons, and kaons are plotted in red, green, and blue respectively.	36
2.6	Side view of a single TOF MRPC module.	38
2.7	The inverse of measured track velocity measured in the TOF as a function of the rigidity measured by the TPC. The theoretical curves for pion, kaon and proton masses are plotted as black lines.	39
3.1	Left: Example event with 6 total proton tracks including 3 within the 120° azimuthal partition. The direction of the initial momenta of the tracks are represented as red arrows, and the gray area represents the azimuthal partition. Right: Example azimuthal partition multiplicity distribution for 500 events with 6 total protons.	47
3.2	100,000 events are generated with 15 random tracks in each. One random azimuthal partition of 120° is taken in each event. The azimuthal partition multiplicity distribution is plotted along with the expected binomial distribution at 100, 1000, and 100,000 events to illustrate the convergence. Note that 100,000 total events are generated and these three plots are snapshots of the distribution as it fills.	49
3.3	100,000 events are generated with 15 random tracks in each and one random azimuthal partition of 120° is taken in each event. The variance of the distribution is plotted in blue against the number of events in the distribution. The expected binomial variance is plotted as a red dashed line. The error bars are calculated via the Delta Theorem.	50

3.4	For each point, x events are generated with 15 random tracks in each. One random azimuthal partition of 120° is taken in each event. The variance of each independent distribution is plotted in blue against the number of events in the distribution. The error bars are calculated via the Delta Theorem and the expected binomial variance is plotted as a red dashed line. Each point is generated independently of all others.	51
3.5	For each point, x events are generated with 15 random tracks and one random azimuthal partition of 120° is taken in each event. The absolute value of the variance of each independent distribution minus the binomial variance is plotted in blue vs the number of events in the distribution. The error bars are calculated via the Delta Theorem.	52
3.6	Events are generated with 15 random tracks and one random azimuthal partition of 120° is taken. In each panel, m events are generated 300 times (“experiments”). The distribution of variances of the azimuthal partition multiplicity distribution for each experiment is displayed as a histogram. The mean, standard deviation and standard error on the mean ($\sigma/\sqrt{n_{\text{experiments}}}$) for these distributions are shown along with the expected binomial variance.	53
3.7	Events are generated with 15 random tracks and one random azimuthal partition of 120° is taken. In each panel, m events are generated 300 times (“experiments”). The distribution of the absolute difference between the variance for each experiment and the expected binomial variance is displayed as a histogram. The mean, standard deviation and standard error on the mean ($\sigma/\sqrt{n_{\text{experiments}}}$) for these distributions are shown.	54

3.8	Events are generated with 15 random tracks and one random azimuthal partition of 120° is taken. At each point on the x-axis, 300 experiments of m events are generated and the deviations of the variances of the partitioned multiplicity distributions from the expected binomial variance (Δ) are calculated. The mean of the Δ distribution at each point on the x-axis is plotted, along with a dark band representing the standard error on this mean and a light band representing the standard deviation of the Δ distribution.	55
3.9	Events are generated with 15 random tracks in each and [1,2,3] random azimuthal partitions of 120° are taken. At each point on the x-axis, 1000 experiments of m events are generated and the deviations of the variances (skewness for middle panel and kurtosis for bottom panel) of the partitioned multiplicity distributions from the expected binomial variance (Δ) are calculated. The mean of the Δ distribution at each point on the x-axis is plotted, along with a dark band representing the standard error on this mean and a light band representing the standard deviation of the Δ distribution.	59
3.10	Events are generated with 15 random tracks with azimuthal partitions of 120° . At each point on the x-axis, 1000 experiments of 250 events are generated taking x samples per event and the deviations of the moments of the partitioned multiplicity distributions from the expected binomial moment (Δ) are calculated. The mean of the Δ distribution at each point on the x-axis is plotted, along with a dark band representing the standard error on this mean.	60
3.11	Example azimuthal partition multiplicity distribution of $m = 10000$ events with $N = 15$ randomly placed tracks and 120° partitions. The x-axis is the number of tracks within the azimuthal partition and the y-axis is the number of events with this multiplicity. The blue bars are a histogram of the simulated events and the red dashes show the expected binomial distribution.	65

3.12	The variance of 100 simulated distributions plotted against the arbitrary experiment number index on the x-axis. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. The true binomial value is plotted as a dashed black line and the green error bars are calculated via the Delta Theorem.	67
3.13	Z-value histograms of the moments of 100 simulated distributions. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. The mean of these distributions is plotted as a dotted green line and the green band about it is the standard error on this mean. A standard normal distribution scaled to the number of entries in the histogram is drawn in red.	68
3.14	Distribution of bootstrap variance, skewness, and kurtosis values derived from the distribution in Figure 3.11. The dashed red line represents the mean of the distribution while the shaded red area corresponds to one standard deviation on each side of the mean. The green dashed line is the expected binomial mean. A rug plot of the individual bootstrap variance values is displayed along with a KDE to help visualize the true distribution.	70
3.15	Moments of 100 simulated distributions plotted against the arbitrary experiment number index on the x-axis. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. The true binomial value is plotted as a dashed black line, the green error bars are calculated via the Delta Theorem and the blue error bars are calculated via bootstrap. The bootstrap and Delta Theorem uncertainties tend to agree well, so the difference in the two types of error bars can be difficult to see in this plot.	71

3.16	For the variances of the experiments shown in Figure 3.15, the bootstrap uncertainty estimate minus the Delta Theorem estimate is plotted as a function of the experiment number (left). A histogram is constructed from these differences (right) with a KDE and rug plot to help visualize the true distribution shape.	72
3.17	Moments of 100 simulated distributions plotted against the arbitrary experiment number index on the x-axis. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. $M = 25$ bootstrap sets are constructed for each experiment. The true binomial value is plotted as a dashed black line, the green error bars are calculated via the Delta Theorem and the blue error bars are calculated via bootstrap. The bootstrap and Delta Theorem uncertainties tend to agree well, so the difference in the two types of error bars can be difficult to see in this plot.	73
3.18	For the variances of the experiments shown in Figure 3.15, a histogram is constructed from the bootstrap uncertainty estimates ($M = 25$ on left and $M = 2500$ on right) minus the Delta Theorem estimates. A KDE and rug plot are included to help visualize the true distribution shape.	74
3.19	Moments of 100 simulated distributions plotted against the arbitrary experiment number index on the x-axis. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. $r = 72$ samples are taken in each event and the Delta Theorem uncertainties are calculated with $r \cdot m = 72 \cdot 10000$ entries. The true binomial value is plotted as a dashed black line, the green error bars are calculated via the Delta Theorem and the blue error bars are calculated via bootstrap.	77

3.20	Moments of 100 simulated distributions plotted against the arbitrary experiment number index on the x-axis. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. $r = 72$ samples are taken in each event and the Delta Theorem uncertainties are calculated with $m = 10000$ entries. The true binomial value is plotted as a dashed black line, the green error bars are calculated via the Delta Theorem and the blue error bars are calculated via bootstrap.	78
3.21	Z-values of the moments of 100 simulated distributions calculated from the bootstrap uncertainty estimates. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. $r = 72$ samples are taken from each event. The mean of these distributions are plotted as a dotted green line and the green band about it is the standard error on this mean. A standard normal distribution scaled to the number of entries in the histogram is drawn in red.	79
3.22	Cumulants in red and k-statistics in blue are plotted as a function of the number of events in the distribution. The second order cumulant (variance) is plotted in the top panel, the third order cumulant in the middle panel, and the fourth order cumulant on the bottom panel. The true value is plotted as a horizontal dotted line in each panel.	81
4.1	The azimuthal multiplicity distribution is compared to binomial in the case of extreme repulsion on the left and extreme attraction on the right. These distributions were simulated with the Gaussian Correlation Model from Section 5.2	84
4.2	Example of azimuthal probability distribution for elliptical flow with $v_2 = 0.1$ (Section 4.4). The gray area represents the integral of this probability distribution between partition edges.	86

4.3	Probability distributions for v_n flow distributions in which only one harmonic (v_n) has a non-zero magnitude.	99
4.4	Analytical integral of $\delta p(\psi)$ in Equation 4.22 where $\tilde{p}(\phi - \psi)$ is of form Equation 4.37, corresponding to the pdfs plotted in Figure 4.3.	100
4.5	Comparison of the analytic vs numerical computation of the $\Delta\sigma^2$ integral for elliptic flow with $v_2 = 0.1$. The bottom panel shows that the difference between numerical and analytic is exceedingly small.	101
4.6	Azimuthal probability distributions for elliptic flow with $v_2 = 0.07$ and $\psi_{v_2} = \pi/3$ and the simple Gaussian clustering model with $A = 0.165$, $\sigma = 0.8$, and $\psi_{\text{gaus}} = \pi$. Also shown is the normalized product of the elliptic flow and Gaussian cluster pdfs in this particular ψ orientation.	102
4.7	$\Delta\sigma^2$ calculated via numeric integration for a simple Gaussian clustering signal with $A = 0.165$ and $\sigma = 0.8$, an elliptic flow background with $v_2 = 0.07$, and for their combination. The elliptic flow $\Delta\sigma^2$ is subtracted from the combination and this difference is found to adequately match the pure Gaussian clustering signal. The Gaussian clustering signal minus the correct value is plotted on the bottom panel along with a band corresponding to the estimated NLO contribution.	103
4.8	A simulation is run with pure clustering signal and another is run with clustering as well as elliptic flow. The contribution from the elliptic flow is calculated from Equation 4.39 with $v_2 = 0.07$ and subtracted from the combined $\langle\Delta\sigma^2\rangle$. After correction, the combined simulation matches the pure signal simulation, demonstrating that the flow correction is effective. The bottom plot shows the signal minus the corrected values. The solid error bars represent the statistical error propagated through the subtraction while the transparent error bars represent the estimated error from NLO contributions from the subtraction.	105

4.9	The Gaussian correlation model with ($A = -0.01, \sigma = 0.1$) is run with an elliptic flow of $v_2 = 0.07$ to test the flow correction. A slight over-correction is observed as a consequence of the Gaussian correlation model coupling to the flow simulation.	106
4.10	The Gaussian correlation model with ($A = -0.01, \sigma = 1$) is run with an elliptic flow of $v_2 = 0.07$ to test the flow correction. A slight over-correction is observed as a consequence of the Gaussian correlation model coupling to the flow simulation.	107
4.11	Azimuthal probability distributions for a simple Gaussian clustering signal and a ϕ dependent efficiency background. The efficiency pdf is derived from the distribution of proton tracks in 62 GeV STAR data. The simple Gaussian clustering model has an amplitude of $A = 0.165$ and $\sigma = 0.8$. Also included is the normalized product of these pdfs in this specific relative orientation.	109
4.12	$\Delta\sigma^2$ calculated via numeric integration for the simple Gaussian clustering signal and STAR efficiency separately as well as for their combination. The efficiency $\Delta\sigma^2$ is subtracted from the combination and this difference is found to adequately match the pure Gaussian clustering signal. The difference between the pure signal and the corrected value is shown in the bottom panel along with a shaded band representing the estimated uncertainty due to NLO contributions.	110
4.13	A simulation is run with pure clustering signal and another is run with clustering as well as STAR ϕ dependent efficiency. In the data set with both signal and background, the mixed correction is implemented $\langle \Delta\sigma_{\text{Single}}^2 - \Delta\sigma_{\text{Mixed}}^2 \rangle$ and this corrected $\langle \Delta\sigma^2 \rangle$ is plotted along with the other two data sets. After correction, the combined simulation matches the pure signal simulation, demonstrating that the mixed correction is effective. The bottom plot shows the signal minus the corrected values. The solid error bars represent the statistical error propagated through the subtraction while the transparent error bars represent the estimated error from NLO contributions from the subtraction.	112

4.14	The Gaussian correlation model with ($A = -0.01, \sigma = 0.1$) is run with the 62 GeV STAR ϕ dependent efficiency to test the mixed event correction. A slight over-correction is observed as a consequence of the Gaussian correlation model coupling to the efficiency simulation.	113
4.15	The Gaussian correlation model with ($A = -0.01, \sigma = 1.0$) is run with the 62 GeV STAR ϕ dependent efficiency to test the mixed event correction. A noticeable over-correction is observed as a consequence of the Gaussian correlation model coupling to the efficiency simulation.	114
4.16	The simple clustering model with ($A = 0.165, \sigma = 0.8$) (Signal) is run with the 62 GeV STAR ϕ dependent efficiency (Background #1) as well as elliptic flow with $v_2 = 0.07$ (Background #2). The efficiency correction is performed on the combined simulation followed by the flow correction. The corrected $\langle \Delta\sigma^2 \rangle$ is compared to the pure simple clustering signal and found to agree well, with the variations shown in the bottom panel rather well explained by the NLO uncertainty estimates.	116
5.1	Example of an azimuthal partition proton multiplicity distribution in the most central AMPT data at 39 GeV. 72 partitions of width $w = 120^\circ$ are sampled randomly in each event and the number of proton tracks within each are counted. A histogram is constructed over all events with $N = 20$ total protons on the full azimuth. AMPT data is plotted as blue bars while the expected binomial distribution of N trials and $p = w/2\pi = 1/3$ probability of success is plotted in red dots.	118

5.2	Left: Variance of azimuthal partition multiplicity distributions as a function of N for the most central AMPT data at 39 GeV. Single event data is shown in blue circles, mixed event data in green circles and the red line shows the expected binomial variance. Right: $\Delta\sigma^2$ is calculated for the same distributions and plotted against N	119
5.3	$\Delta\sigma^2$ is plotted as a function of the total number of protons on the full azimuth, N , for the most central events in STAR 39 GeV data. The single event data is shown in blue circles, mixed event data in green squares, and the mixed corrected data in red triangles. Statistical uncertainties are shown as solid lines while systematic uncertainties are represented as transparent bars.	120
5.4	$\Delta\sigma^2$ is plotted for the 30-40% centrality class at 39 GeV. The uncorrected $\Delta\sigma^2$ is shown in blue dots, the mixed corrected data is shown in red triangles, and $\Delta\sigma^2$ after both mixed and v_2 correction is shown as black stars.	121
5.5	$\Delta\sigma^2$ after mixed event and elliptic flow corrections is shown for the most central events at 39 GeV. STAR data is shown as black circles, AMPT is represented as a red band, the default MUSIC+FIST model is shown as a blue band, and the MUSIC+FIST model with an excluded volume effect is shown as a purple band.	122
5.6	$\Delta\sigma^2$ is shown as a function of N . Each panel represents one of the six BES-I energies analyzed.	123
5.7	$\Delta\sigma^2$ averaged over the total number of protons per event, N , is shown as a function of center of mass energy. STAR data is shown as black circles while AMPT and the MUSIC+FIST default and excluded volume models are shown as red, blue and purple bands, respectively.	124

5.8	$\langle\Delta\sigma^2\rangle$ is plotted for each centrality class against the average reference multiplicity of that centrality class. The six AMPT energies are shown in the left panel as semi-transparent bands. STAR data is shown in the right panel as circles which are connected for each energy. A fit to AMPT data for all energies is displayed in both panels as a dashed line. Though not visible here, the fit clearly fails at high multiplicity where statistics are high and is shown only to guide the eye in comparison with STAR data.	125
5.9	The 62.4 GeV STAR and AMPT $\langle\Delta\sigma^2\rangle$ is fit as a function of reference multiplicity.	126
5.10	The fit of $\langle\Delta\sigma^2\rangle$ for 62.4 GeV is subtracted from all $\langle\Delta\sigma^2\rangle$ values to better visualize the energy dependence for AMPT data on the left and STAR data on the right.	127
5.11	The fit of $\langle\Delta\sigma^2\rangle$ for 62.4 GeV is subtracted from all $\langle\Delta\sigma^2\rangle$ values. These differences are averaged over centrality and plotted against beam energy.	128
5.12	$\langle\Delta\sigma^2\rangle$ is plotted against the azimuthal partition width and fit to a quadratic. . .	129
5.13	The baseline b from the quadratic fit of $\langle\Delta\sigma^2\rangle$ as a function of w is plotted against the reference multiplicity for AMPT on the left and STAR data on the right. . .	130
5.14	The curvature from the quadratic fit of $\langle\Delta\sigma^2\rangle$ as a function of w parameterized in z is plotted against the reference multiplicity for AMPT on the left and STAR data on the right.	131
5.15	Simulation $\Delta\sigma^2$ plotted as a function of the total number of protons per event. Six simulations are run in total, three repulsive and three attractive with $\sigma = 1$ and $A \in \pm[0.002, 0.006, 0.01]$	135
5.16	$\langle\Delta\sigma^2\rangle$ plotted on the y-axis against the amplitude (A) simulation parameter on the x-axis. Both attractive and repulsive simulations are run with $\sigma \in [0.8, 1.0, 1.2]$ and $A \in \pm[0.002, 0.004, 0.006, 0.008, 0.01]$. A linear fit is superposed for each of the six sets shown, illustrating the linearity between $\langle\Delta\sigma^2\rangle$ and A . .	136

5.17	$\langle \Delta\sigma^2 \rangle$ plotted against the width of the azimuthal partitions used. Simulation Amplitudes of $A \in \pm[0.002, 0.004, 0.006, 0.008, 0.01]$ are shown with $\sigma = 1$	137
5.18	$\langle \Delta\sigma^2 \rangle$ plotted against the width of the azimuthal partitions used. Simulation Amplitudes of $A \in -[0.006, 0.01]$ are shown with $\sigma \in [0.4, 1.2]$	138
5.19	The measured curvature of $\langle \Delta\sigma^2 \rangle$ as a function of azimuthal partition width, w is parameterized by z and plotted against the azimuthal correlation range from simulation, σ . Simulation Amplitudes of $A \in -[0.002, 0.004, 0.006, 0.008, 0.01]$ are shown with $\sigma \in [0.8, 1.0, 1.2]$	139
5.20	Parameters of the quadratic fits to $\langle \Delta\sigma^2 \rangle$ vs partition width w are plotted with the z parameter on the y-axis and the b parameter on the x-axis. The z parameter indicates the distance from 180° at which the fit intersects the x-axis and is inversely related to the curvature. The b parameter represents the maximum magnitude of the quadratic which, by definition, occurs in this parameterization at 180° . Simulation Amplitudes of $A \in \pm[0.002, 0.004, 0.006, 0.008, 0.01]$ are shown with $\sigma \in [0.8, 1.0, 1.2]$. Repulsive simulations have a negative b and are on the left of the plot while attractive simulations are on the right. The simulations are grouped by σ separately for attractive and repulsive. The mean is calculated for each group and represented as a horizontal line, with the standard error on the mean represented as a shaded horizontal band.	140
5.21	The baseline, b , of the quadratic fit of $\langle \Delta\sigma^2 \rangle$ vs w is plotted against the the simulation amplitude A for both attractive and repulsive simulations.	141
5.22	Number of events at each stage of event level cuts for 7 GeV. The left panel shows the cuts made before downloading the data locally. The right side shows the cuts made locally for the default analysis. The cuts on the right are varied for systematic checks.	143
5.23	Normalized primary vertex v_z (left) and reference multiplicity 3 (right) distributions.	144

5.24	Distributions of track variables for protons used in the analysis.	147
5.25	Azimuthal angle distributions of protons used in analysis. The TPC sector boundaries are clearly visible. The distributions in red correspond to the protons used in the default version of the analysis while the blue distributions correspond to a larger set of protons (including the red distribution) with much looser cuts, used for systematic variations. Most notably, the blue distribution includes protons with $ \eta \leq 2.1$ while the $ y < 0.5$ cut is performed to get the red distributions.	148
5.26	Rapidity distributions of a loose selection of protons in blue and the protons after all default cuts have been made in red.	148
5.27	STAR event plane resolution (left) and v_2 of identified protons (right) as a function of centrality for six BES-I energies.	151
5.28	AMPT event plane resolution (left) and v_2 of identified protons (right) as a function of centrality.	152
5.29	Event plane resolution (left) and v_2 of identified protons (right) for STAR, AMPT, and all three versions of the MUSIC+FIST model as a function of energy.	153
5.30	Dcaxy run average vs run index for all six BES-I energies analyzed. Only the most central 20% of events were averaged.	155
5.31	7.7 GeV Dcaxy event average vs event index for run 11126001. The purple rectangles enclose bad event ranges which are excluded from the analysis.	156
5.32	Left: Example of BTOFMatch vs Refmult correlation at 62.4 GeV. Right: The histogram on the left is rotated such that a linear fit is vertical. This should ensure that horizontal slices are Gaussian if the two detectors are linearly related.	158

5.33	Example horizontal slice of the rotated 62.4 GeV BTOFMatch vs refmult distribution. A Gaussian fit alone (left) is unable to model these distributions adequately. Fitting with a Gaussian plus two Woods-Saxon distributions multiplied by decaying exponential functions (right) seems to produce the most stable fitting results. From these fits, the $n\sigma$ x-axis locations with respect to the mean of the Gaussian distribution are calculated (yellow vertical lines).	159
5.34	The $n\sigma$ Gaussian cuts for each slice are rotated back to the original coordinates and fit with fourth order polynomials (extrapolated at the high end). These red lines define the boundary between the good events between then and the pile-up events above and below. The original green fitting points are difficult to see as they are covered by the red fit lines.	160
5.35	Visualization of the systematic checks made for the 0-5% centrality class in 7.7 GeV data using 120° partitions. On the bottom panel, the default $\langle\Delta\sigma^2\rangle$ value is shown on the left, with the solid error bars representing the statistical uncertainty and the transparent band representing the total systematic uncertainty. The $\langle\Delta\sigma^2\rangle$ value of each systematic variation is plotted. Each type of systematic check is grouped together and displayed in the same color. The top panel shows the contributions of each systematic source to the total systematic uncertainty, plotted on the left in black with the statistical uncertainty represented as a red line for comparison.	163
5.36	Visualization of the baseline values from the quadratic fit to $\langle\Delta\sigma^2\rangle$ vs partition width for the 0-5% centrality class at 7.7 GeV.	164

5.37	Table showing the contribution of each systematic source to the total systematic uncertainty of $\langle\Delta\sigma^2\rangle$ at 120° for all data sets. The top row (default) shows the size of the systematic uncertainty compared to the magnitude of $\langle\Delta\sigma^2\rangle$. Each of the succeeding rows represents a source of systematic uncertainty and its contribution to the total systematic uncertainty. Each column (excluding the default row) will sum to 100% within rounding error.	166
5.38	Table showing the contribution of each systematic source to the total systematic uncertainty of the $\langle\Delta\sigma^2\rangle$ vs partition width fit baseline for all data sets.	167
A.1	$\langle\Delta\sigma^2\rangle$ is calculated for events generated with M total particles (unfortunately labeled as N in the legend) and plotted against M . The data is fit to power law in M plus a constant. The fit is plotted and the fit parameters are displayed. . .	175

LIST OF TABLES

3.1	Average deviation of the bootstrap uncertainty estimates from the Delta Theorem estimates for bootstrap sets of size 25, 250, and 2500. The percent deviation is calculated by dividing by the average Delta Theorem uncertainty of 0.046. . . .	75
5.1	Minimum bias triggers and number of events used in analysis.	142
5.2	Track level cuts for selecting tracks counted in the reference multiplicity variables and for selecting proton tracks for analysis.	145
5.3	Dcaxy default and systematic trigger $n\sigma$ values.	157

ACKNOWLEDGMENTS

First and foremost, I thank my friends and family who have supported me through my life and academic career. I am extremely fortunate to have been raised by parents who have provided for my every need, allowing me to discover and pursue my passion.

I next thank the teachers who have nurtured my scientific curiosity: those such as Nancy Dwyer and Karen Smith who pushed me before I was ready to push myself, as well as Professors David Koltick, Norbert Neumeister, Zvi Bern and Terry Tomboulis who taught me physics. I am particularly grateful for Barbara Hagedorn who inspired and cultivated my nascent interest in science and helped me discover my talents and interests.

I thank Virgil Barnes whose free-spirited lectures convinced me to pursue physics. Vigil and Ephraim Fischbach provided an incredible introduction to research and were the best mentors a young scientist could hope to have. I also thank Purdue colleges Nathan Sahelijo, David Bernstein, Jonah Polly, and Ting-Wei Hsu.

I thank my UCLA mentors Prof Huan Huang, Gang Wang, Oleg Tsai, Steven Trentalange, and Roli Esha as well as colleagues Brian Chan, Zhiwan Xu, Xiatong Wu, Thomas Marshall, Aditya Dash, Yunshan Chen, Nihal Gozlukluoglu, Ryan Milton, and Mahsa Sergeeva. A special thanks to Zhongling Ji who helped tremendously with my analysis and thoroughly understanding various physics topics.

Finally, I thank STAR mentors Nu Xu, ShinIchi Esumi, Bedangadas Mohanty, Xiaofeng Lu, Toshihiro Nonaka, Hanna Zbroszczyk, Xu Sun, Saehanseul Oh, as well as colleagues Ashish Pandav, Yu Hu, Zach Sweger, Ben Kimmelman, and Matt Harasty.

VITA

2012 – 2016 Bachelor of Science in Physics and Mathematics,
Purdue University, West Lafayette

2016 – 2018 Teaching Assistant,
Department of Physics and Astronomy,
University of California, Los Angeles

2018 – 2023 Graduate Student Researcher,
Department of Physics and Astronomy,
University of California, Los Angeles

CHAPTER 1

Introduction

Physics had seen an explosion of growth in interest and understanding in the twentieth century. The discovery of quantum mechanics explained failures of the classical theories of Newton and Maxwell and gave an entirely new perspective on the fundamental laws governing matter and its interactions at the smallest distance scales. At the same time, General Relativity was developed to explain gravitational phenomena at astronomical distances, combining space and time into a fabric distorted by the presence of mass. By the latter half of the century, a new bedrock of theory was being codified in the Standard Model of particle physics. While this model has been enormously successful over the last half century it is, at best, an effective theory which must eventually give way to a theory which can describe gravity at high energy. Even so, the Standard Model will remain an indispensable component of any theory that succeeds it, and there is still much within the model, particularly within quantum chromodynamics, that has yet to be fully understood.

1.1 The Standard Model

The Standard Model of particle physics is currently our best model for understanding the matter our universe is made of. It describes the matter of the observable universe as a set of fundamental particles and their interactions. The model contains 12 fermions, their anti-particles and the 6 force carrying bosons. The fermions are comprised of the quarks and leptons which come in three generations, each no more than a heavier copy of the last. Within each generation, there are two flavors of quark and two leptons. The first-generation




Figure 1.1: Constituents of the Standard Model of particle physics.

quarks are the up (u) with electric charge $+2/3$ and the down (d) with charge $-1/3$. The up and down quarks combine to form protons (uud) and neutrons (udd) which are the basic building blocks of all atomic matter. The first-generation leptons are the electron (e^-) and the electron neutrino (ν_e). The next two fermion generations are identical copies of these four particles with larger mass (with the possible exception of the neutrinos, whose mass ordering is not yet known). The quarks exist only in bound states called hadrons: a quark and anti-quark pair is a meson and a combination of three quarks is a baryon. Baryon number and lepton number are conserved quantities in the Standard Model, along with electric charge.

Interactions between the fermions are mediated by the bosons. The photon (γ) is the electromagnetic force carrier and interacts with all fermions besides the neutrinos through their electric charge. Quantum Electrodynamics (QED) describes these interactions to an impressive level of precision. The W^\pm and Z bosons are the force carriers of the weak force, so called because these bosons are massive which suppresses their exchange. The weak and electromagnetic forces were unified into the Electro-Weak theory in the 1960s [1]. The massless gluon is the mediator of the strong interaction, which is described by the theory of Quantum Chromodynamics (QCD). The final boson in the Standard Model is the Higgs, discovered in 2012 [2][3] and responsible for giving the other particles of the theory their mass via the




Figure 1.2: An illustration of the electromagnetic, weak and strong coupling constants unifying at high energy under a supersymmetric theory.

Higgs Mechanism (W^\pm and Z bosons) and Yukawa Coupling (fermions).

1.2 Physics Beyond the Standard Model

The Standard Model provides a coherent description of the electromagnetic, weak and strong forces up to the current highest experimentally accessible energies of $E_{cm} = 14$ TeV. This energy frontier continues to march forward in hopes of providing evidence of new physics that must exist beyond the Standard Model.

1.2.1 Unification of the Quantum Theories

With the unification of first the electric and magnetic forces and later the weak force, it became an enticing goal of theoretical particle physics to unite all of the fundamental forces under a single theory. The development of QED and gauge theories lead to the understanding that the coupling constants which dictate the interaction strength of the strong, weak and electromagnetic forces were not constant, but depended upon the energy scale. It was plausible that, at very high energies, these coupling constants would converge to a common value. This convergence would be suggestive of a Grand Unified Theory (GUT) that splits

into the known effective theories of the Standard Model at energies below the GUT scale. As experimental data of the coupling constants as a function of energy has been collected, extrapolations to high energy are making this convergence less likely without some additional physics at mid-energy to divert the trajectories.

Supersymmetric (SUSY) theories, which envision heavy partners of all Standard Model particles, were invented to provide the needed adjustments to the running of the coupling constants. While these types of theories are a very attractive and elegant successors to the Standard Model, they necessitate the existence of these heavy supersymmetric particles along with a mechanism for proton decay. The lower bound for proton lifetime [4] and the lack of light SUSY particles at the LHC have ruled out many of the simplest theories. The search for new physics on the energy frontier continues, though the excitement around supersymmetric GUT theories has begun to fade.

1.2.2 Gravity

Even if light SUSY particles are observed in experiment, these supersymmetric theories only explain three of the four fundamental forces. The perennial issue of constructing a quantum theory of the gravitational force has plagued the theoretical community since the birth of quantum mechanics. While QCD and QED describe the strong and electroweak forces to astonishing accuracy, gravity has been best understood as the curvature of space-time as laid out in Einstein's General Theory of Relativity (GR). Gravity is dominant at astronomical distances and GR has been immensely successful in describing phenomena such as orbital trajectories and gravitational lensing. More recently, the 2017 Nobel Prize was awarded for the observation of gravitational waves, predicted by Einstein in 1916, which originated from the merging of two black holes [5]. However, the quantum field theories that describe the other three forces are typically constructed on a flat space-time background, neglecting gravitational interactions. Attempts have and are being made to construct consistent field theories in curved space-time, but none have yet been successfully unified with the Standard

Model to describe all of physics within a single Theory of Everything.

1.3 Quantum Chromodynamics (QCD)

The Standard Model is clearly not the Theory of Everything that physicists have sought for the last century. Within its realm of validity, however, its predictions are accurate to an astonishing level of precision. Despite its great successes, there are still many areas within the framework of the Standard Model that have yet to be fully explored and understood, chief among these being non-perturbative Quantum Chromodynamics (QCD).

1.3.1 The QCD Lagrangian

The theory of QCD is relatively straightforward, the Lagrangian easy to write down [6]:

$$\mathcal{L}_{\text{QCD}} = \sum_q \psi_{q,a} (i\gamma^\mu \partial_\mu \delta_{ab} - g_s \gamma^\mu t_{ab}^C \mathcal{A}_\mu^C - m_q \delta_{ab}) \psi_{q,b} - \frac{1}{4} F_{\mu\nu}^A F^{A\mu\nu} \quad (1.1)$$

Here the a index on the quark-field spinors $\psi_{q,a}$ runs from 1 to 3 and represents the three QCD colors charges. The q index runs from 1 to 6 and denotes the quark flavor with corresponding masses m_q . \mathcal{A}_μ^C represents the gluon field with C indexing the 8 gluons. t_{ab}^C are the eight generators of the SU(3) group and g_s is the strong coupling constant. $F_{\mu\nu}^A$ is the gluon field tensor. The four terms then correspond to the free quark kinetic term, the quark-gluon interaction term, the quark Yukawa mass term, and the free gluon term.

While this Lagrangian is easy to write down, it leads to enormously complex physics. For one, the gluon field strength term produces three and four point gluon-gluon interactions. While a photon must interact at a charged vertex with some fermion, a gluon carries color charge itself and a single gluon can split into two via the three point vertex. As opposed to the charge screening effect in QED via pair-produced dipoles, the three and four point gluon vertices produce an anti-screening effect in QCD, producing an effectively stronger

field at larger distance from the source. This anti-screening is encoded in the negative beta function of QCD which characterizes the running of the coupling strength as a function of energy. QED has a positive beta function, leading to a weaker coupling at larger energies. QCD's negative beta function [7][8], on the other hand, implies larger coupling as energy is decreased.

1.3.2 Perturbative QCD

Much of the success of quantum field theories has been in the perturbative regime. In terms of Feynman diagrams, fermion fields interact with each other via gauge boson propagators. At the simplest level, one fermion field is connected to another via a single propagator — a tree level interaction. It's perfectly reasonable that two fermions could be connected via a propagator that undergoes some other interaction between the two external fermion vertices, producing loops within a more complicated diagram. These loop diagrams can be arbitrarily complicated and, in order to calculate the cross section of a specific interaction, all possible diagrams must be considered and summed over. If there are an infinite number of arbitrarily complicated diagrams that produce the same final state fermions, this problem may seem intractable.

The saving grace of this type of calculation is that, when calculating the amplitude of a single diagram, the amplitude is proportional to the coupling constant to the power of the number of vertices. Each additional loop adds at least two vertices to the diagram, and the more complicated the diagram the larger the total number of vertices. If the coupling constant is small (much less than 1), this means that the more vertices there are in a diagram, the smaller its amplitude. This small parameter allows for the expansion of an interaction in powers of the number of vertices or loops, where the tree level diagrams have the fewest loops (zero) and therefore the largest amplitudes. More complicated diagrams necessarily have more vertices and therefore have successively smaller amplitudes as the complexity increases. This perturbative expansion allows for the calculation of interaction




Figure 1.3: The running of the strong coupling constant as a function of the momentum transfer [6].

cross sections to an arbitrary level of precision — higher order diagrams provide smaller and smaller corrections to the tree level calculation.

1.3.3 Limitations of Perturbative QCD

Perturbative QCD has been enormously successful in describing hard scattering in high energy collision experiments. As the interaction energy between partons decreases (a general term for the quarks and gluons which interact via the strong force), the strong coupling constant increases. As it approaches values on the order of unity, the suppression of higher order diagrams weakens. All arbitrarily complicated diagrams begin to become comparably important and can no longer be treated as corrections to the tree level calculation. In this strong coupling regime, the perturbative scheme breaks down and new methods must be pursued to make progress on understanding the relevant physics. While difficult for calculations, this infrared regime is very important for understanding the structure of hadrons

along with other novel systems such as neutron stars and the early universe.

1.3.4 Asymptotic Freedom and Color Confinement

The breakdown of perturbative QCD occurs at small interaction energies or, equivalently, large distances. We can imagine a quark and anti-quark interacting via the strong force as the distance between them increases. At very short distances the strong coupling is small as the interaction energy is large. The strength of the strong interactions between them is weak at these small distances and the quarks are therefore said to be asymptotically free. As the distance between these quarks increases, the corresponding interaction energy decreases. The gluons mediating this interaction carry color and interact amongst themselves. As the quarks separate, the mediating gluons are attracted to each other and form a color flux tube between the quarks. Increasing the distance between the pair increases the length scale and the tension in the flux tube. This leads to a large coupling constant and an increased importance of higher order diagrams. These diagrams correspond to more loops, some involving pair production of further quark anti-quark pairs. As the energy in the flux tube continues to increase and pair production becomes more prevalent, it eventually becomes energetically favorable to pair produce on shell quarks and pair them with the original two quarks. This produces two independent pairs of quarks with interactions only between the pair at short distances — two stable mesons.

This general picture makes it impossible to isolate a free quark, a phenomenon called color confinement. It states that only color singlet states, hadrons, are observable in nature. This has held up experimentally as no indication of free quarks has yet to be observed. It is important to note that, while this phenomenon is understood as the consequence of the negative beta function of QCD along with the fact that gluons carry color, there has been no analytical proof of color confinement from first principles and the QCD Lagrangian.

1.3.5 Lattice QCD

One method for understanding the non-perturbative physics of QCD is to give up entirely on an analytical solution and resort to numerical methods. Expectation values for observables can be calculated from the partition function, defined in the path integral formulation with the Euclidean QCD Lagrangian:

$$\langle \mathcal{O} \rangle = Z^{-1} \mathcal{O} Z \quad (1.2)$$

$$Z_{\text{QCD}} = \int D\psi D\bar{\psi} D\mathcal{A}_\mu \exp [-S_{\text{QCD}}(\psi, \bar{\psi}, \mathcal{A}_\mu)] \quad (1.3)$$

$$S_{\text{QCD}} = \int d^4x \mathcal{L}_{\text{QCD}}(x) \quad (1.4)$$

This formulation is only able to calculate time independent observables since the Lagrangian is Euclidean. The QCD Lagrangian and therefore the action can be decomposed into the free gluon term S_G which is independent of the fermion field and all of the terms that depend on the fermion fields S_F :

$$S_{\text{QCD}} = S_G(\mathcal{A}_\mu) + S_F(\psi, \bar{\psi}, \mathcal{A}_\mu) \quad (1.5)$$

We can then calculate the action of a given field configuration of quarks and gluons on a discrete grid of points with three spatial dimensions and one (imaginary) time dimension. Each lattice site (n) contains a vector of 12 values, 3 colors and 4 spins, corresponding to values of quark fields $\psi(n)$, $\bar{\psi}(n)$. Each link between lattice sites, indexed by lattice site n and direction μ , contains a 3x3 matrix representation of an $SU(3)$ element related to the gluon fields $U_\mu(n) = \exp[iag\mathcal{A}_\mu]$ where a is the lattice spacing.

The fermion fields are Grassmann valued and can actually be computed analytically [9].

$$Z_{\text{QCD}} = \int DU_\mu \prod_{q=1}^{N_f} \det M(m_q, U_\mu) \exp[-S_G(U_\mu)] \quad (1.6)$$

where $M(m_q, U_\mu)$ is a large matrix for each quark flavor. When an operator expectation value is calculated, the fermion integration of $\mathcal{O}Z$ must be carried out by hand. From there, the gluon action $S_G(U_\mu)$ can be calculated for a specific gluon configuration on the lattice U_μ by summing the action calculated at each link. The path integral over the fermion fields is then computed from $\det M(m_q, U_\mu)$, which are enormous matrices and the bottleneck for lattice calculations. The only thing left is to integrate over gluon configurations. This is done by generating many gluon field configurations and averaging over the resulting operator expectation value for each.

Exploring the full space of gluon field configurations on the lattice is prohibitively expensive. Instead, Monte Carlo sampling is performed on configurations near equilibrium. The gluon action acts as a weight for the importance of the gluon configuration, $\exp[-S_G(U_\mu)]$, and Markov chain algorithms are used to update the gluon field configuration and bring the system closer to its minimum action. These algorithms are stochastic, such that the system is brought into equilibrium but still samples fluctuations [10].

Lattice QCD is an extremely powerful tool and has been used to successfully predict the behavior of QCD configurations with equal amounts of matter and anti-matter. When a matter asymmetry is introduced, doping the QCD matter with more quarks than anti-quarks and producing a finite baryon chemical potential (μ_B), lattice QCD runs into a problem. The QCD action at zero baryon chemical potential is strictly real after Wick rotation, meaning that e^{-S} can be interpreted as a probability of a particular field configuration which allows importance sampling to work. The chemical potential term in the action, however, comes with an imaginary factor. This produces an e^{-S} with an imaginary part which ruins its role as a configuration weight. Without importance sampling, the convergence of the numerical approximation is much slower and no longer practically feasible. This is the notorious sign

problem of lattice QCD which restricts numerical calculations on the lattice to small baryon chemical potential.

Chemical Potential The magnitude of matter doping is encoded in the chemical potential, which is the energy required to add a new particle to the system. The chemical potential for bosons is zero, since multiple bosons can occupy the same state and new particles can then always be added to the ground state. For fermions, the Pauli Exclusion Principle prohibits two particles from occupying the same state. As new particles are added to a system with minimum energy, they cannot be placed in an already occupied ground state. Instead, they must be placed in the lowest energy state which is not occupied — the surface of the Fermi sea. The chemical potential is the energy of this state which must be provided to add a particle to the system.

Anti-particles can be viewed as the absence of a particle or a hole in the Fermi sea. If particles occupy states up to some energy E , this energy can be extracted from the system by adding an anti-particle (equivalent to removing a particle). Adding anti-particles will reduce the chemical potential while adding particles will increase it.

1.3.6 Quark Gluon Plasma (QGP)

Quarks and gluons are typically trapped within hadrons due to the color confinement of QCD. It is possible to bypass this confinement by either increasing the quark density of a system of hadrons or increasing its temperature. Increasing quark density mitigates gluon anti-screening by providing quarks multiple available partners within the Debye length of the system, replacing the usual string breaking response to anti-screening. By increasing temperature, the average kinetic energy and therefore interaction energy between partons in a hadron gas increases and therefore decreases the strong coupling strength, producing a larger Debye length. In either case, a system of color confined hadrons can transition to a system of semi-free quarks and gluons: Quark Gluon Plasma (QGP).




Figure 1.4: The phase diagram of QCD with temperature on the y-axis and baryon chemical potential on the x-axis [14]. The true structure of the phase diagram is not known, this is a simplified depiction of how it may look.

The asymptotic freedom of QCD implies that at very high temperatures (very large average interaction energy), QGP is a very weakly interacting gas of quarks and gluons. At lower temperatures, however, this plasma need not be a weakly interacting gas. In fact, experiment has found QGP at these temperatures to be a strongly interacting liquid [11] with the lowest shear viscosity to entropy density ratio of any known fluid [12]. This has also been understood theoretically by mapping the fluctuations on the event horizon of a black hole in GR to the thermodynamic characteristics of an infinitely strongly coupled system. This AdS-CFT correspondence has allowed the calculation of the shear viscosity to entropy density ratio in an infinitely strongly coupled fluid. The calculated value of $\frac{1}{4\pi}$ is not much lower than what is measured for QGP by experiment [13].

1.3.7 The Phase Diagram of QCD

The two relevant parameters for transitioning from hadron gas to QGP are temperature and quark density. Treating a system of quarks and gluons as a thermodynamic system, we can then explore the phase diagram of QCD with temperature on the y-axis and baryon chemical potential (μ_B , dual to baryon number in the QCD partition function) on the x-axis.

At low temperature and μ_B , a system of quarks and gluons will form a gas of color singlet hadrons. Due to the matter anti-matter asymmetry of our universe the diagram in Figure 1.4 only shows positive μ_B values, though the negative side of this diagram also exists for anti-matter dominated systems. Normal nuclear matter is found close to the x-axis (low temperature) around $\mu_B = 940$ MeV. Keeping the temperature low and increasing the baryon chemical potential may transition QCD matter into a color superconducting state. This would be similar to electric superconducting but the color charge, instead of electric charge, would produce quark Cooper pairs. The core of neutron stars is the only place in the universe hypothesized to contain QCD matter in this color superconducting state. Systems of QCD matter at high temperature, such as those formed in heavy ion collisions and the early universe, exist as QGP.

1.3.7.1 The Ising Model

Increasing the temperature of a hadron gas will eventually lead to a phase transition to QGP. Of particular physical interest is the order of this phase transition. The behavior of the baryon density as QCD matter transitions from one phase to another depends on the order of this phase transition. One of the most intuitive systems with which to understand phase transitions is the Ising model. This model consists of an N-dimensional lattice of spins which can be oriented either up or down. The spins interact with each of their nearest neighbors and it is energetically favorable for these neighbors to be aligned. We can introduce an external magnetic field (h) along which it is energetically favorable for each spin to align. In addition,




Figure 1.5: Phase diagram of the Ising model [14].

we define a temperature as a source of random fluctuations in the spin configurations, the larger the temperature the larger the expected deviation of a system from its ground state.

$$H = - \sum_{(i,j)} J s_i s_j - \sum_i^N h s_i \quad (1.7)$$

Taking the two dimensional case of the Ising model for simplicity (Hamiltonian in Equation 1.7), we can first look at the case of zero magnetic field and temperature. Here there is a degeneracy of the ground state: all spins pointing up or all pointing down. In both cases all spins are aligned and the system has minimum energy. In realistic models analogous to the Ising model (a ferromagnet) this symmetry is spontaneously broken as the system settles into one of the two degenerate states. Keeping the temperature zero, we can now introduce a finite magnetic field. All spins will align with the magnetic field, breaking the degenerate ground state explicitly. We can next imagine starting with a magnetic field pointing in one direction, slowly decreasing its magnitude to zero, and then slowly increasing it in the opposite direction. The spins of the model at zero temperature will all always align with the external field and will flip when the external field changes direction. As the external field is continuously varied there is a discontinuous variation of the spin directions.

$$\langle M \rangle = \frac{1}{N} \sum_i^N s_i \quad (1.8)$$

To better define and quantify the effects described above, we can define the magnetization as the sum of all spins on the lattice. The average magnetization, Equation 1.8, then quantifies the average direction of the spins in the system. In the ground state, this average magnetization (to be referred to as simply magnetization) has degenerate values of ± 1 and is discontinuous as the external magnetic field continuously changes sign. The magnetization gives information about the state of the system and observables such as these are called order parameters. The external magnetic field acts as a lever with which we can tune our system and these type of parameters are generally referred to as control parameters.

The zero temperature Ising model is said to undergo a phase transition as the control parameter flips the system from one degenerate state to the other. This transition occurs at $h = 0$ which is considered the phase boundary between the two phases. Since the order parameter is discontinuous as the control parameter is continuously varied to bring the system through the phase transition, this transition is called a first-order phase transition [15].

Now that the system and jargon is understood, we can introduce a finite temperature into the Ising model. At a very low but non-zero temperature, we expect only very slight thermal deviations from the ground state. In this case, we expect the behavior of the system to be almost identical to the zero temperature case: most spins will align with the external magnetic field and there will be a discontinuity of the magnetization at $h = 0$. As the temperature is increased, however, the thermal fluctuations will become more important, bringing the system further from the ground state on average. At very high temperatures (thermal energy very large compared to spin-spin and magnetic energy terms in the Hamiltonian), these thermal fluctuations dominate and the system's spins become random. At some large but finite temperature, the system is mostly random but is still sensitive to the external magnetic field. At zero magnetic field the spin alignment over a large lattice is random and the average magnetization is zero. As the magnetic field is




Figure 1.6: Mapping of the Ising phase diagram onto that of QCD.

increased, the magnetization will also increase slightly as spins are somewhat more likely to align with the magnetic field. In this region, as the external magnetic field is continuously varied from one direction to the other, the magnetization will follow but will encounter no discontinuities. The phase transition in this region of the phase diagram is called a crossover, as the order parameter and all of its derivatives are continuous as the control parameter is varied across the phase boundary.

From the arguments above we can see that the phase diagram of the Ising model has a phase boundary at $h = 0$. At low temperature this phase transition is first-order, while at large temperatures it is a smooth cross-over. These two transition regions must somewhere connect. This boundary between a first-order and cross-over phase transition is called the critical point.

1.3.7.2 From Ising to QCD

While the microscopic dynamics of the Ising model have nothing to do with those of QCD, the two systems actually share macroscopic similarities. QCD is expected to be in the same universality class as the 3D Ising model and all other systems which undergo a liquid-gas

phase transition [16][17]. Being in the same universality class implies that the two phase diagrams should both contain a critical point and exhibit similar behavior in this critical region. This would then imply that the QCD phase diagram contains both a cross-over and first-order phase transition, separated by a critical point. While the phase boundary of the Ising model exists on the $h = 0$ axis, the QCD phase transition from hadron gas to QGP depends on both baryon chemical potential and temperature, producing a phase transition curve in the QCD phase diagram.

A major goal of non-perturbative QCD, pursued by the BEST collaboration, is to uncover the equation of state of QCD matter [18][19]. This may be accomplished through a mapping of the 3D Ising phase diagram onto the QCD diagram in the region of the critical point. The mapping amounts to determining six parameters: two for the location of the critical point in T and μ_B , two for the orientations of the two Ising axes at this point, and two to account for the scaling of the axes. While two of these can be fixed by lattice QCD, the other four must be constrained through a combination of experimental data and equation of state modeling.

QCD Control Parameters The order parameters of QCD are the conserved charge densities: baryon, electric, and strangeness. A phase diagram can be constructed for each of these with the corresponding chemical potential on the x-axis. We will, however, focus exclusively on the baryon density and the phase diagram with baryon chemical potential as the control parameter on the x-axis.

1.3.7.3 The QGP to Hadron Gas Phase Transition

The QGP to hadron gas phase transition is the subject of intense theoretical and experimental investigation. On the theoretical side, the most accessible part of the diagram are the edges where the theory can be simplified as different limits are exploited. The $\mu_B = 0$ axis of the phase diagram is particularly approachable. Here, lattice QCD is not hindered by the sign problem and Monte Carlo techniques converge. The equation of state of the sys-




Figure 1.7: Lattice QCD calculations for thermodynamic quantities demonstrating that the phase transition from hadron gas to QGP is a cross-over at low μ_B .

tem is extracted from lattice QCD by computing thermodynamic quantities of the simulated configurations. Plotting the pressure, energy density and entropy density as a function of temperature, a rapid increase is found for each of these between 140-170 MeV. This comes from the increased number of degrees of freedom of the system as the quarks and gluons are liberated from hadrons to form QGP. Importantly, however, there is no discontinuity found in these quantities or their derivatives, indicating that the phase transition is a smooth cross-over at $\mu_B = 0$. While the sign problem confines the validity of lattice to $\mu_B = 0$, approximations have allowed for somewhat reliable probes off of this axis. The range of these approximations are limited but, with current techniques, no sign of anything but a cross-over transition has been observed in lattice calculations.

Investigating this phase transition at finite μ_B is a much more difficult task. Lattice QCD is unable to reliably probe large baryon chemical potentials and other theoretical approaches tend to be model dependent. The task is then left to experiment to explore the QGP to hadron gas phase transition at finite baryon chemical potential to determine if the cross-over




Figure 1.8: Visualization of the stages of a heavy ion collision.

transition at small μ_B ever yields to a thermodynamically discontinuous transition. It is expected that the QCD phase diagram shares the critical characteristics of the 3D Ising model, which would imply that the cross-over transition eventually becomes first-order at a critical point. But this prediction needs to be confirmed by experiment.

1.4 Relativistic Heavy Ion Collisions

Relativistic collisions of heavy ions at colliders are the only experiments capable of studying deconfined QCD matter. Heavy nuclei moving near the speed of light experience a Lorentz contraction in the lab frame along the direction of motion. The two ions then appear as colliding circular disks. On collision, a large system of high energy density is created. These systems form droplets of Quark Gluon Plasma which subsequently hadronize into hadron gas as the droplet expands and cools. The challenge of heavy ion experiments is to infer properties of the ephemeral QGP matter from the final state hadrons which reach the detectors.

1.4.1 Stages of a Heavy Ion Collision

Lacking a computationally simple solution, heavy ion collisions are best understood as a set of stages with different physics — the so-called hybrid model. The initial state refers to the configuration of the nuclei before collision. The locations of the nucleons inside of each nucleus can leave an imprint on the collision that is translated through to the final state hadrons. As the two nuclei begin to make contact, hard scattering processes between

high momentum partons take place. These perturbative processes eject high energy particles from the collision which directly enter detectors or form jets in the final state. As the two disks move through and past each other, many softer interactions take place. This produces a volume of high temperature partons in the overlap region which have been slowed relative to the passing nuclei. If the energy density is high enough, QGP forms in this volume and quickly thermalizes. From there, the QGP undergoes rapid hydrodynamic expansion which pushes the quarks and gluons of the plasma out from the collision site. The QGP cools as it expands until it eventually transitions back into hadron gas. The resulting hadrons will continue to interact and can change their hadronic identity until the system has spread out far enough to reach chemical freeze-out. Even then, these final state hadrons can still engage in elastic collisions until kinetic freeze-out, at which point they continue free streaming into the detector.

Rapidity and pseudorapidity The velocity of a relativistic particle is typically not a very good quantity to work with, as velocity differences are not Lorentz invariant. Instead, we can define the rapidity along the beam axis:

$$y = \frac{1}{2} \log \left(\frac{E + p_z c}{E - p_z c} \right) = \tanh^{-1} \left(\frac{p_z c}{E} \right) \quad (1.9)$$

Unlike velocity, the difference between the rapidities of two particles, $y_1 - y_2$, is invariant with respect to boosts along the beam direction. We also define the related pseudorapidity:

$$\eta = -\log \left[\tan \left(\frac{\theta}{2} \right) \right] \quad (1.10)$$

Here, θ is the polar angle with respect to the beam direction. The pseudorapidity is therefore simply a remapping of this polar angle. The pseudorapidity is equivalent to the rapidity for a massless particle and the two are closely related for relativistic particles in which the energy is much larger than the mass.

Center of Mass Energy to μ_B Changing the center of mass energy of a heavy ion collision will change the baryon chemical potential of the system formed at mid-rapidity — the lower the collision energy, the larger the μ_B . This can be understood as a consequence of the fact that the colliding systems are composed of matter rather than anti-matter. When a collision takes place, some of the baryon number is stopped and transported to mid-rapidity. At high beam energy, the colliding partons have large rapidities and are harder to stop and the mid-rapidity baryon density is lower. In addition, the energy density of the QGP is larger, which is converted to quark/anti-quark pairs which dilute the initial baryon excess. These two factors produce systems of QGP with very small μ_B in high energy collisions. As the collision energy is decreased, baryons are more easily stopped and fewer quark/anti-quark pairs are produced, so the relative asymmetry of baryon number at mid-rapidity increases. We utilize this dependence of μ_B on collision energy to map the QCD phase diagram with beam energy scans.

1.5 Experimental Signatures of a Critical Point

Experimentally determining the existence and location of a critical point in the QCD phase diagram has been a goal of paramount importance for the last two decades. The critical point is simply the boundary between a cross-over and first-order phase transition, so the discovery of a first-order transition in addition to the known cross-over transition at $\mu_B = 0$ would be sufficient to infer the existence of a critical point separating the two. While a first-order phase transition should produce final state hadrons in a significantly different manner than a cross-over transition, these differences may not produce signatures strong enough to be measured in our detectors. Searches are ongoing for both the first-order phase transition as well as the critical point itself.

1.5.1 Critical Point Search

The critical point itself may be easier to find than evidence of a first-order phase transition. On the first-order transition line, both states exist with distinct baryon densities but equally minimal free energies. Moving along the transition line to lower μ_B values, these distinct densities for the two states must approach each other since they are equivalent in the cross-over region. As these densities approach, the free energy barrier between them gradually decreases, allowing for easier fluctuation from one state to the other. At the critical point, the boundary between the cross-over and first-order phase transitions, the densities of the states have just become equal. The minimum of the free energy here is flat and becomes sharper as we move further into the cross-over region. A flat free energy corresponds to an ill-defined minimum density and therefore large fluctuations of the density of the system. Though the magnitude of these density fluctuations at the critical point are damped by finite size and finite time effects [20], it is expected that these fluctuations will still produce a signal that may be observed in experiment.

Baryon density fluctuations in heavy ion collisions are expected to manifest in the multiplicity of baryons produced at mid-rapidity. Baryon number is exactly conserved in each event and will therefore be identical in each event if all baryons are measured. If only a sub-volume of the event is considered, variations in the number of baryons in this sub-volume of phase space will be directly related to the baryon density fluctuations in the QGP system. The number of baryons produced at mid-rapidity in each event is measured in experiment and the resulting distribution over a large number of events is analyzed. For a stochastic production of baryons about an expected value, a Poisson distribution is expected which becomes Gaussian for large mean values. Fluctuations in the baryon density on an event-by-event basis will shift this expected value in each event and produce a non-Gaussian multiplicity distribution. When fluctuations are large, as is expected near the critical point, this baryon multiplicity distribution should deviate significantly from Gaussian.

The degree of deviation from Gaussian for these multiplicity distributions is measured via the moments or cumulants of the distributions. Cumulants are a measure of the shape of a distribution. The first cumulant is the mean, the second is the variance which measures the width of the distribution, the third is related to the asymmetric skew of the distribution, and the fourth is related to the peakedness of the distribution.

$$\begin{aligned}
C_1 &= \langle N \rangle \\
C_2 &= \langle (\delta N)^2 \rangle \\
C_3 &= \langle (\delta N)^3 \rangle \\
C_4 &= \langle (\delta N)^4 \rangle - 3\langle (\delta N)^2 \rangle^2 \\
\delta N &= N - \langle N \rangle
\end{aligned}
\tag{1.11}$$

The moments of a distribution can be written in terms of the cumulants:

$$\begin{aligned}
\mu &= C_1 \\
\sigma^2 &= C_2 \\
S &= \frac{C_3}{C_2^{3/2}} \\
\kappa &= \frac{C_4}{C_2^2}
\end{aligned}
\tag{1.12}$$

Cumulants themselves are extensive quantities which scale with the volume of the system. Taking ratios of cumulants cancels this volume dependence. Cumulant ratios are related to the corresponding ratios of thermodynamic susceptibilities which provides a convenient bridge between theory and experiment [21]. While the order parameter of interest is the conserved net-baryon number, experiment is typically unable to measure the neutral neutrons which make up half of the baryon number in a collision. Instead, proton or net-proton multiplicity distributions are used as a proxy for net-baryon number.

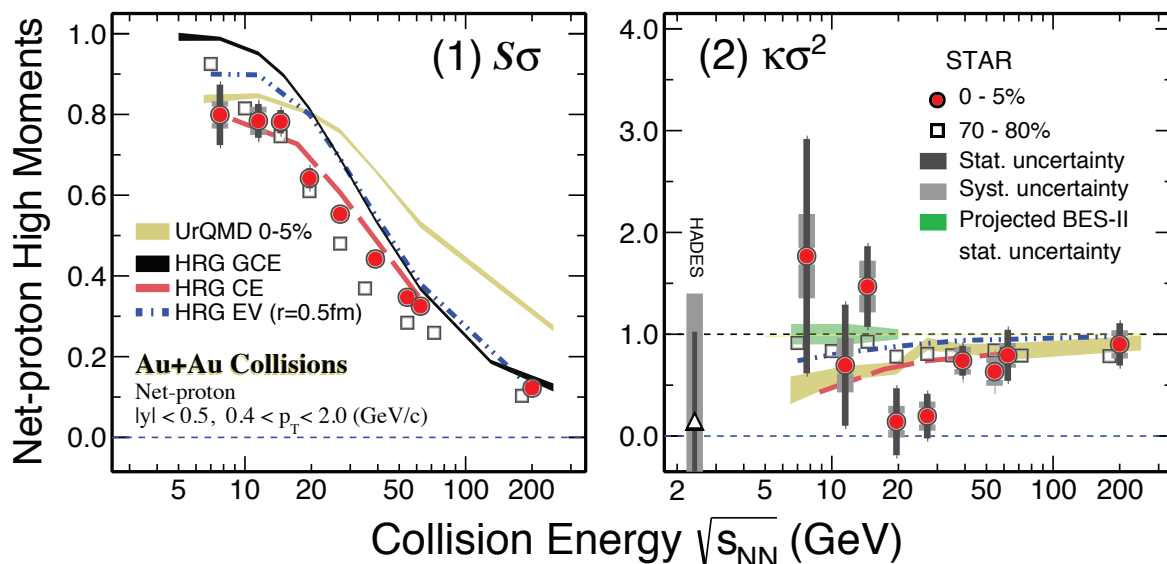


Figure 1.9: The normalized skewness and kurtosis of net-proton multiplicity distributions at mid-rapidity are plotted as a function of center of mass collision energy. The non-critical UrQMD and HRG models are shown along with data from peripheral collisions as baselines.

The cumulants of proton and net-proton multiplicity distributions at mid-rapidity have been key observables in the search for a critical point. In particular, the fourth order cumulant is predicted to exhibit non-monotonic energy dependence as the critical region is passed in a beam energy scan [22]. The kurtosis times variance squared was measured by the STAR experiment and a possible non-monotonic trend was observed, shown in Figure 1.9, where $\kappa\sigma^2 = C_4/C_2$ and $S\sigma = C_3/C_2$.

1.5.2 Signals of a First-Order Phase Transition

While direct observation of enhanced fluctuations near the critical point has been the observable of most interest to the heavy ion community, experimental observation of a first-order phase transition would equally well imply the critical point's existence. In a hypothetical first-order transition region, both states of QCD matter would have distinct and different baryon densities. Above the transition line, the QGP phase is favored and the free energy is minimum at the baryon density of the QGP state. The opposite is true on the hadron gas

side of the first-order line, where the hadron gas density is minimal. Exactly at the phase transition line, the free energy has two equal minima corresponding to each state's baryon density. This degeneracy results in the system being equally likely to exist in either state. Whenever there are two local minima, it is possible for the system to equilibrate into either phase producing the phenomenon of phase coexistence.

Phase coexistence is a general feature of first-order phase transitions and can lead to both of these two states existing in different regions of a given system. These regions of different density will experience a surface tension between them, producing bubbles or clusters. This is most familiar in the liquid to gas transition of water in which bubbles of gas form instead of gas being homogeneously produced throughout the liquid.

A similar effect may occur in a first-order phase transition of QGP to hadron gas. In this case, regions of cooling QGP may transition before others and produce “bubbles” of hadrons which may produce a detectable signal. One obstacle in searching for this signature is that the conjectured bubbles would be produced in coordinate space, which is not directly accessible via experiment. Instead, these bubbles would flow outward with the expanding QGP and be translated into some correlation in the momentum space of the final state hadrons. Models have explored this translation of a possible clustering effect from coordinate space to momentum space and have found that much of the clustering signal may be washed out in the process [23]. Clustering of final state hadrons, protons specifically, will nonetheless be the target of the experimental search presented in Chapter 5.

CHAPTER 2

RHIC and the STAR Detector

The STAR detector at Brookhaven National Laboratory (BNL) is designed to detect the remnants of heavy ion collisions provided by the RHIC complex. Both the collider complex and the STAR detector are immensely complex machines which require many subsystems working together to produce snapshots of the exotic QCD matter produced in relativistic heavy ion collisions.

2.1 RHIC

The Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL) accelerates and collides atomic nuclei at relativistic energies. Ion production and acceleration is accomplished via a chain of separate machines comprising the RHIC complex. Ions are produced in the Vacuum Arc Ion Source and sent to the Electron Beam Ion Source where about half of the electrons are stripped. The Booster and AGS synchrotrons accelerate the ions in stages before they are finally injected into the RHIC rings.

2.1.1 Vacuum Arc Ion Source

Metal ions bound for the RHIC rings typically begin as a cylindrical cathode in a vacuum arc ion source [24]. The cathode is placed in a vacuum and a voltage is applied to produce a spark between the cathode and an annular anode. Since this spark occurs in a vacuum, the charge carriers along the spark are a plasma of electrons and vaporized metal ions from




Figure 2.1: An aerial view of the RHIC complex.

the cathode. A fraction of this plasma travels to the anode to sustain the arc while another fraction passes through the anode's annulus. The metal ions are formed and accelerated into a beam.

2.1.2 Electron Beam Ion Source (EBIS)

The beam of $1+$ ions produced in the vacuum arc ion source is then sent to the Electron Beam Ion Source (EBIS) to be further ionized [25] [26]. The EBIS contains a series of drift tubes and a strong electron beam which is compressed in a magnetic field. The metal ions from the vacuum arc ion source are seeded into the EBIS and trapped axially by the drift tubes and radially by the space charge of the electron beam. In the trap, electrons are stripped from the metal ions through kinetic interactions with the electron beam. While the positive metal ions are trapped, the freed electrons are ejected from the EBIS. The longer ions are held in the trap, the more electrons are stripped. This allows the EBIS to produce any charge state up to the limit at which the positive ion charges neutralize the electron beam space charge, at which point the radial trap is compromised. Once the metal ions have been cooked in the trap long enough to produce the desired charge state, one side of the axial voltage is dropped (downstream of the electron beam) and the positive charges are accelerated out of the trap. As the ions must be trapped for some time to achieve the

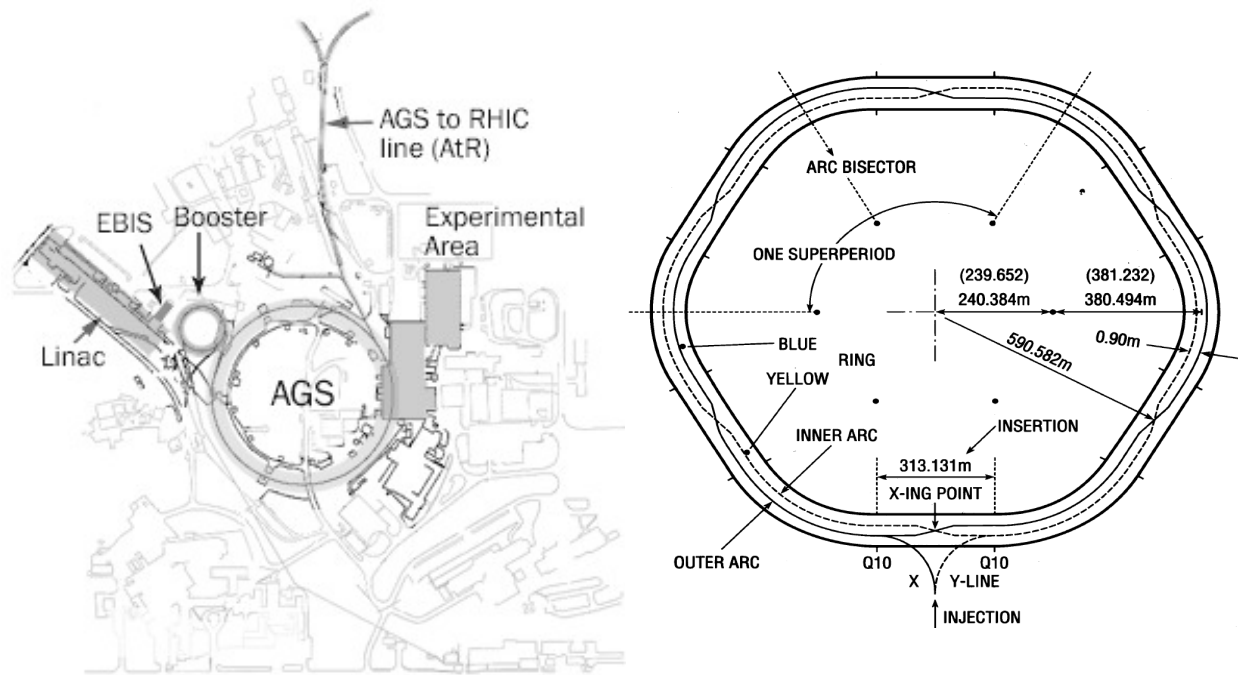


Figure 2.2: Left: Diagram of the accelerator complex leading to the RHIC ring. Right: Diagram of the RHIC rings.

desired charge state, they are ejected in pulses or bunches. The ions are then accelerated by a radio frequency quadrupole (RFQ) and a short Linac along with three stages of focusing, removing some undesired charge states in the process. For Au-197, the EBIS produces a substantial number of ions with charge states between +30 and +34 and with energy of 16.2 keV/u when leaving the trap. By the end of the subsequent acceleration, only the desired +32 charge state remains with $3e9$ ions per bunch at 2 MeV/u.

2.1.3 Accelerating Synchrotrons

From the EBIS, the ions next traverse a series of three synchrotrons which accelerate them to full energy. These accelerators all operate on the same principle: synchrotron acceleration. Charged particles are injected into a close-to-circular ring comprised of a series of modular sections. Most of the synchrotron circumference is made up of Focusing-Drift-Defocusing-Drift (FODO) cells [27]. These cells contain one or more dipole magnets which bend the

trajectories of the charged particles, keeping the beam moving in the close-to-circular orbit around the ring. As the beam is bent, it is defocused. To refocus the beam, a minimum of two quadrupole magnets are required, as a single quadrupole will focus the beam in one direction but defocus it in the transverse direction [28]. Higher moment magnets: sextupoles, octupoles, and some magnets even more complicated are also employed in the FODO cells to maintain a focused beam.

While these FODO cells along the circumference of the ring keep the beam on orbit and focused, often a single radio frequency (RF) cavity is used in a straight section of the ring to accelerate the charged particles. The RF cavity is designed to provide a resonance chamber for a specific frequency of radiowave, with a wavelength tailored to the size and speed of the beam's bunches. The RF phase is matched to the bunch orbit such that as the bunch passes, it is accelerated in the electric field. The radiowave's electric field magnitude oscillates as a sinusoid and the RF cavities also focus the bunch in the longitudinal direction by aligning the center of the bunch just after the electric field maximum. This gives late particles an extra kick and early particles less of a kick, compressing the bunch along the beam direction.

The bunches in a synchrotron circulate close to the speed of light and make hundreds of thousands of revolutions per second, passing through the RF cavity on each revolution (though the RF cavities typically operate in pulsed mode). As the RF cavities accelerate the beam's particles, the dipole magnets responsible for maintaining their orbit must increase their magnetic field to support the same trajectory at higher momentum. The focusing magnets must also adjust their fields to keep the beams focused. A synchrotron is limited in its maximum acceleration of heavy ions by the strength of the magnetic field the bending dipoles are able to produce. For this reason, these magnets are superconducting and account for most of the cost of these accelerators.

2.1.3.1 Booster

The first of the three synchrotrons in the RHIC complex is the Booster [29]. Focusing on the acceleration of gold ions, the partially stripped bunches enter the Booster from the EBIS at 2 MeV/u. Six bunches at a time circulate about the 200 m circumference of the ring and are accelerated to 100 MeV/u, or about one third the speed of light. Once fully accelerated, the ions are ejected from the ring by energizing a kicker magnet, one bunch at a time. The bunch is directed toward a copper foil roughly 1 mm thick. As the beam passes through the copper, most of the remaining electrons are stripped and the ions lose about 5 MeV/u. While the copper thickness is designed to maximize ions of the desired charge state, +77 for gold, a distribution of charge states will result. To select only the desired charge state, a series of two dipole magnets bend the stripped ions. The dipoles bend unwanted charged states out of the beam pipe and into shielding designed to absorb the energy of these discarded ions. The remaining ions of the desired charge state are refocused through a series of quadrupole magnets.

2.1.3.2 Alternating Gradient Synchrotron (AGS)

The next synchrotron in the acceleration chain is the Alternating Gradient Synchrotron (AGS) [30]. This machine was constructed in the 1960s after the development of the strong focusing principle [28] and was the highest energy accelerator in the world for most of the decade. It is now used as a stepping stone in the RHIC complex, accelerating protons and ions to injection energy for the RHIC rings.

The AGS is a half mile in circumference and gold ions of charge state of +77 are injected from the Booster at 95 MeV/u. 24 bunches are transferred from the Booster to the 12 AGS buckets in four cycles (2 bunches per bucket), after which the beam is debunched and rebunched into 4 bunches. The ions are then accelerated to a final energy of up to 10.5 GeV/u before they are ejected from the AGS, one bunch at a time. The bunches are sent to

the AGS to RHIC transfer line (ATR) where they pass through another stripping foil that removes the last two electrons from the gold ions, bringing the charge state from +77 to a fully stripped +79.

2.1.3.3 Relativistic Heavy Ion Collider (RHIC)

The fully stripped ions continue down the ATR where they encounter a switching magnet that directs the bunch vertically to either the “X” or “Y” beam lines which proceed in opposite directions to tangentially intersect the RHIC rings. The “Y” line injects bunches into the “Yellow” ring which circulates counter-clockwise while the “X” line injects into the clockwise circulating “Blue” ring. Each RHIC ring nominally has 360 buckets and 111 of these are typically filled with bunches from the AGS. This corresponds to a spacing of one bunch per three buckets, omitting 9 of these buckets to form an abort gap. When the beam needs to be dumped, an abort kicker magnet at 10 o’clock on the RHIC ring is turned on while these empty abort gap buckets are passing. The beam is then ejected from the ring and safely dumped into an absorber.

There are six interaction points on the RHIC ring which lie at the center of straight (insertion) sections of the beam pipe [31] [32]. Connecting each of these six insertion sections there are six arc sections comprised of FODO modules which bend the beam around the rings. RHIC’s RF cavities, acceleration cavities operating around 28.15 MHz and storage cavities at 197 MHz, are located at the 4 o’clock straight section. Gold ions can be accelerated to energies between the max AGS injection energy of 9.8 GeV/u up to 100 GeV/u, producing collisions with center of mass energy from $\sqrt{s_{NN}} = 19.6$ GeV per colliding nucleon pair to $\sqrt{s_{NN}} = 200$ GeV. Lower energy injections from the AGS can also be stored in the RHIC rings, pushing the lowest feasible center of mass energy down to $\sqrt{s_{NN}} = 7.7$ GeV.

In each of the six insertion sections around the RHIC ring, the yellow and blue beams can be steered such that they intersect at the interaction point at the center of the section. The collision rate R is determined by the geometric properties of the beams, encoded in the

luminosity \mathcal{L} , as well as the probability of individual ion pairs interacting (cross-section) σ : $R = \mathcal{L}\sigma$ [33]. The luminosity is proportional to the product of the number of ions in each bunch and inversely proportional to the cross sectional area of each bunch at the interaction point $\mathcal{L} \propto \frac{N_1 N_2}{A}$. As the beam circulates, the positively charged ions expand and need to be refocused by RHIC magnets to keep the beam area small. In the rest frame of the beam ions at high energy, little time passes between refocusing magnets and there is therefore less time for the bunch to expand before being refocused. At lower energies, the time between refocusing increases as the beams travel slightly slower and the time dilation effect is reduced. This makes it more difficult to keep the beam tight at lower energies which increases the cross sectional area of the bunches at beam crossing and reduces the luminosity.

2.1.4 Beam Energy Scan

The Beam Energy Scan (BES) program at RHIC was conducted to map out the QCD phase diagram. This energy scan took Au+Au data at various beam energies to probe the properties of QGP and, specifically, its phase transition back to hadron gas at various temperatures and baryon densities. While high center of mass energies were already available, lower energies would be able to produce higher baryon densities at mid-rapidity, pushing the produced systems into previously unexplored areas of the phase diagram. During the 2010 and 2011 runs, collisions were recorded for energies between 7.7 and 62.4 GeV. A second scan, BES-II, took place between 2019 and 2021 to improve the statistics at the lower energies.

2.2 The STAR Detector

The Solenoidal Tracker at RHIC (STAR) is the longest running experiment on the RHIC ring [34]. Commissioned in 2000, STAR has full azimuthal coverage and is designed to study QCD at high energy density. It has recorded Au+Au collisions at the maximum center of mass energy of $\sqrt{s_{NN}} = 200$ GeV along with two beam energy scans at lower




Figure 2.3: A rendering of the STAR detector as it was configured during the 2010 and 2011 Beam Energy Scan I.

energy. In addition, STAR has recorded polarized proton+proton and proton+Au collisions, deuteron+Au, Ruthenium+Zirconium isobar collisions, as well as collisions between several other heavy ion species. We will focus on Au+Au collisions, specifically the first beam energy scan which took place in 2010 and 2011.

The STAR detector is comprised of a suite of subsystem detectors. RHIC steers the yellow and blue beams such that they cross at around 1 mrad in the beam pipe at the center of STAR. On collision, final state particles from the colliding system emanate from the collision vertex. Those emitted at mid-rapidity first pass through the beam pipe and enter the Time Projection Chamber (TPC) which can reconstruct the three-dimensional particle tracks. From there, the tracks enter the Time Of Flight (TOF) which helps to identify the particle from its measured velocity. After the TOF, the particles then traverse the Barrel Electromagnetic Calorimeter (BEMC) which is intended to stop photons and electrons and measure their energy deposition. In addition, there are several detectors designed to detect

particles near beam rapidity, such as the Vertex Plane Detector (VPD), the Beam Beam Counter (BBC), and the Zero Degree Calorimeter (ZDC), among others. STAR is enclosed within a large room temperature solenoidal magnet which produces a uniform 0.5 T magnetic field that bends the trajectories of particles traveling through the detector, allowing their momentum to be measured by the TPC. Taken together, these subsystems, along with some others not mentioned, are able to reconstruct the trajectories and identify the particles produced in Au+Au collisions.

2.2.1 Time Projection Chamber (TPC)

The Time Projection Chamber (TPC) is a gaseous tracking detector and often referred to as the “heart” of STAR [35]. 4.2 meters long and extending radially from half a meter to two meters from the central axis, the TPC covers mid-rapidity from $\eta = -1.8$ to $\eta = +1.8$. As the first detector particles encounter after passing through the beam pipe, the design of the TPC as a gaseous tracking detector keeps track energy deposition low, allowing for the non-destructive measurement of particle trajectories.

The active volume of the TPC is filled with P10 gas (90% Argon and 10% methane by volume) held at a pressure of 2 mbar above atmosphere. The detector is split into West and East sub-volumes, separated at the center by a thin central membrane fixed at 28 kV that acts as a cathode. A relatively uniform electric field is produced between this cathode and the grounded endcaps of the TPC which act as anodes. As charged particles traverse the volume, they interact with and ionize gas molecules. Positive ions drift slowly toward the central membrane while the freed electrons drift to the endcap anodes. The drift velocity of the electrons is relatively fast, roughly 5.5 cm/ μ s, and is measured with lasers every few hours to maintain calibration. The drift electrons avalanche at the Multi-Wire Proportional Chambers (MWPC) on the endcaps and this signal is measured via an induced image charge on readout pads. This produces an electrical signal of amplitude proportional to the initial number of drift electrons.




Figure 2.4: Left: Diagram of STAR's TPC. Right: Pad row layout for one TPC sector.

The TPC is made up of twelve sectors which span the azimuth. Each of these sectors contains inner and outer radial subsections. The inner subsection contains 13 widely spaced rows of small pads. Since track density is higher radially closer to the interaction point, the inner subsector is optimized for track position resolution at the expense of poor energy deposition resolution. The outer subsector has 32 rows of larger, densely packed pads. Tracks are more spread out as they travel further from the interaction point and the outer subsector allows for tracking as well as good energy deposition resolution.

When reconstructing an event, TPC pads with readings surpassing a threshold value are considered hits and are assumed to correspond to a portion of a particle track, though some of these can also be background hits. For a single track, this produces a map of the hit points which can be connected radially between pad rows to determine trajectory of the particle in the radial and azimuthal planes. As the electron drift velocity is known from laser calibration, the longitudinal position of hits can be calculated via the arrival time relative to the start of the event multiplied by this drift velocity. This produces a full three-dimensional reconstruction of the track as it traverses the TPC volume.

There can be hundreds to thousands of tracks in each event, and a pattern recognition algorithm is necessary to connect these pad hits radially and reconstruct individual tracks. A particle traversing the TPC volume will not necessarily produce a signal on each of the




Figure 2.5: Energy loss per centimeter measured in the TPC as a function of the rigidity (momentum divided by charge) for 39 GeV BES-I data. Theoretical Bichsel curves are drawn as black lines for each particle species. Identified pions, protons, and kaons are plotted in red, green, and blue respectively.

pad rows due to gaps and statistical fluctuations, and the tracking algorithm must attempt to reconstruct these tracks even when only a fraction of the track’s hits are available. In analysis, we typically require that a good track has at least 15 pad row hits out of 45 as well as requiring that the fraction of possible hits be between 0.51 and 1.02, as the possible number of hits depends on the curvature and path of the track ($46/45 = 1.0222$ so 1.02 is used to deal with float comparison — including 45 hit tracks and excluding those with 46).

Once a group of hits are identified as a track, they are fit to a helical trajectory, characteristic of particles traveling in a uniform magnetic field, with corrections to account for energy loss as the particle traverses the gas volume. The radius of the helix is inversely proportional to the magnitude of the magnetic field and proportional to the momentum of the particle. Since the STAR’s magnetic field is known and uniform, the particle’s momentum is extracted from this helix’s radius (the magnetic field is mapped and corrections are made for deviations from uniformity). In addition, each pad hit has an amplitude proportional to

the number of drift electrons ionized, which provides a measurement of the energy loss of the particle per unit path length.

With the energy loss and momentum measured, the particle which left the track can be identified using the Beta Bloch equation for the energy loss of a particle traversing a medium. This equation has been extended to account for experimental effects by Bichsel [36]. A plot of energy loss (dE/dx) vs particle momentum divided by its charge (rigidity) is shown in Figure 2.5 along with the predicted Bichsel curves for pions, kaons and protons along with their anti-particles. Based on the distribution of track density in these plots about the predicted curves, each track is assigned an $n\sigma_{\text{particle}}$ value which characterizes the probability of each track being identified as a particular particle. The three bands in Figure 2.5 are well separated at low momentum and the particle identification is reliable in this region. As momentum increases the bands begin to merge and the TPC becomes unable to accurately differentiate the particle species at momenta larger than around 800 MeV.

2.3 Time Of Flight (TOF)

The Time Of Flight (TOF) detector is a gas based detector that surrounds the TPC [37]. This detector is designed to provide additional particle identification capability, especially at large momentum where the various particle bands of the TPC dE/dx vs rigidity plot merge. The detector is made from trays of Multi-gap Resistive Plate Chambers (MRPCs). STAR's Barrel TOF (BTOF) contains 120 MRPC trays, each of which are 2.1 meters long and extend along the direction of the beam from the center of the detector. There are 60 trays on the West side and 60 on the East, each of which cover 6° and together provide full azimuthal coverage from $\eta = -1$ to $\eta = +1$. In each tray, there are 32 MRPC modules, each angled perpendicular to the center of the detector.

Each MRPC module is a series of six resistive glass plates stacked between two outer electrodes with a large potential difference of 7 kV between them. A thin gas gap, 95%




Figure 2.6: Side view of a single TOF MRPC module.

Freon and 5% isobutane, separates each of the glass plates. Particles travel perpendicular to the stack of plates and ionize gas in the gap, producing avalanches. As the gas gaps are very thin, the avalanches produced in each gap are quenched at the next plate almost simultaneously, producing a signal with good timing resolution. The total avalanche signal read out from the outer electrodes will be the sum of the individual gas gap avalanches. This design allows for the detection of particles with timing resolution of around 60 ps.

The TOF cannot function in isolation but must match any hit it registers to a corresponding track from the TPC. The TPC tracks provide helical fits which can be extrapolated out of the TPC volume and into the enclosing TOF. If an extrapolated helix matches the position of the hit on the TOF, the hit is associated with the TPC track. The TOF measures the time at which the particle produces a hit. This time minus the start time of the event, t_0 as measured by the Vertex Plane Detector (VPD), is the flight time of the particle. The distance that the particle has traveled is the path length of the helical fit of the corresponding TPC track — from the collision vertex to the TOF cell.

Knowing the flight time and the distance traveled, the velocity of the particle can be calculated $\beta = \frac{dx}{dt}/c$. More massive particles will have larger momentum for the same velocity. This produces the well separated bands for each particle when plotting $1/\beta$ as a function of the momentum in Figure 2.7. In practice, the mass of each particle is calculated from its velocity and momentum $m^2 = p^2(1/\beta^2 - 1)$ and cuts are made on the mass for particle identification.




Figure 2.7: The inverse of measured track velocity measured in the TOF as a function of the rigidity measured by the TPC. The theoretical curves for pion, kaon and proton masses are plotted as black lines.

The TOF's mass measurement provides additional information about the matched TPC track that can be utilized to help identify the particle. Demanding that an identified TPC track have a measured TOF mass within some range of its expected value allows for a higher purity of identified particles. However, not all TPC tracks are matched to a TOF hit and if a TOF mass measurement is required, unmatched TPC tracks must be sacrificed. This TOF matching efficiency depends on the kinematic variables of the track and the collision system but is typically around 70-80%.

Vertex Plane Detector (VPD) The start time of an event is typically measured by the Vertex Plane Detector (VPD). This detector is comprised of 38 scintillator-PMT tubes, 19 on the East side and 19 on the West. Each side has two concentric rings of tubes, 5.6 meters from the center of the detector. The inner rings cover a pseudorapidity of $4.6 < \eta < 5.1$ while the outer rings cover $4.2 < \eta < 4.6$. The timing of a coincident signal between the

East and West VPD provides the start time of the event, which the TOF needs to calculate the flight time. The difference in timing between the East and the West also provides a measurement of the longitudinal position of collision vertex, which can be used to check the vertex position reconstructed from TPC tracks.

2.4 Trigger and Data Acquisition (DAQ)

An important aspect of collider experiments is determining which data to record. Bunch crossings occur at a rate of around 9 MHz and in each there is a small probability of two gold ions colliding. When a collision does occur, the trigger subsystem is responsible for identifying the collision and recording the information from each detector as an event. All of the detector subsystems are continuously read out, but data is only recorded when a Au+Au collision occurs. The TPC is a slow detector, unable to resolve tracks on the time scale of a bunch crossing, and is therefore not used in the trigger. The low-level readouts from other fast detectors are used in various combinations to form triggers. These triggers perform quick logical checks and fire if it is determined that a collision has occurred.

Once a trigger fires, the detectors are read out and this information is written to disk. The trigger operates on several levels. Level 0 corresponds to the fastest and simplest calculations and comparisons on low level detector information. Levels 1 and 2 combine detector information and perform rudimentary reconstruction. Level 3 is a full reconstruct the event. After being triggered at Level 0, an event can be rejected at any of the higher-level triggers if it fails to meet specific criteria. Events surviving the level two trigger are sent to a computer farm in the control room which performs full reconstruction of the event. The reconstructed event is then saved for further calibration, quality assurance and, eventually, physics analysis [38].

Zero Degree Calorimeter (ZDC) and the Minimum Bias Trigger The most important trigger for many STAR analyses is the minimum bias trigger, whose goal is to record Au+Au collisions of any impact parameter. The minimum bias data set allows the centrality of each event to be determined based on its multiplicity and utilizes the Zero Degree Calorimeter (ZDC) in tandem with the TOF for its level 0 trigger. The ZDC detectors are alternating layers of tungsten and scintillating fibers which are fed into PMTs and are situated downstream of the first bending magnetic on both sides of the interaction point. The magnet redirects any charged particles and the ZDCs therefore measure only the neutral particles from the event, mostly spectator neutrons. The minimum bias trigger fires on the coincident signal of at least one neutron in both Zero Degree Calorimeter (ZDC) detectors along with a handful of hits in the TOF.

2.5 Collision Vertex Reconstruction

When reconstructing a triggered event, a critical piece of information to be determined is the three-dimensional location within the beam pipe at which the collision occurred [39]. This collision vertex is reconstructed from the trajectories of TPC tracks. Extrapolating the trajectories of these tracks back to the beam pipe, the collision vertex is identified by clustering algorithms. This primary vertex location then provides an extra fit point for the helical fits of the TPC tracks. The Distance of Closest Approach (DCA) of the helical fit to the collision vertex for each track is also calculated. The DCA can be used in analyses to distinguish particles originating from the collision itself, called primary particles, from those originating from the decay of other particles emanating from the event, called secondary particles.

In addition, the coincidence timing in the VPD is used to identify the time at which the collision occurred, providing a t_0 for the primary vertex. The difference in timing between the East and West VPD can also be used to determine the longitudinal position of the vertex at

energies of 39GeV and above. In analysis, it is required that the longitudinal position of the collision vertex determined by the VPD match that of the TPC vertex within 3 centimeters. If the two don't agree within this criterion, the event is removed from the analysis.

2.6 Centrality Determination

Gold nuclei in their rest frame are spherical but are Lorentz contracted along the direction of motion when boosted to relativistic speeds. This means that in the STAR lab frame, the gold nuclei appear as flat discs. When these discs collide, the nucleons in the overlapping region interact and produce a system of large energy density. The nucleons outside of this overlapping region interact much less with the opposing nucleus and tend to pass relatively unimpeded down the beam pipe. The nucleons in the overlap region are referred to as participants while those outside of the overlap region which don't interact are referred to as spectators.

Neglecting the initial configuration of the nucleons, we can approximate each colliding nucleus as azimuthally symmetric about the center of the Lorentz contracted disk. The collision between two nuclei can therefore be described entirely in terms of the interaction region as parameterized by the impact parameter and reaction plane angle. The impact parameter is the distance between the center of the two colliding discs. The line between the two disk centers forms an angle with respect to the azimuth called the reaction plane angle. The overlapping region and corresponding number of participants is independent of the reaction plane angle, which only defines the azimuthal orientation of each collision. It is the impact parameter that dictates what is referred to as the centrality of each event.

Events with smaller impact parameters have larger overlap regions, producing systems of larger energy and volume at mid-rapidity. It may therefore be beneficial to order and group events by centrality to study observables as a function of system volume. It is then necessary to measure the centrality of each event. It is not possible to measure either

the impact parameter nor the number of participants directly, though models have shown that more central collisions will produce a larger final state particle multiplicity. We can therefore measure the multiplicity of each event, which should be directly related to the event's centrality.

This relationship can be complex: it depends on the center of mass energy and may be complicated by detector effects that influence the measured multiplicity. In addition, as events become more peripheral, the event multiplicity decreases. At some point, the collision is so peripheral and the multiplicity is so low that it becomes unlikely that the minimum bias trigger will fire, in which case the event will not be recorded. This is referred to as the detector's trigger inefficiency which results in a decreasing probability of recording an event as the impact parameter increases.

If all events with impact parameter less than some large and fixed value were recorded consistently between data sets, we could simply divide the data set into groups based on the measured event multiplicity distribution, called the reference multiplicity in STAR. Events with the top 5% largest reference multiplicity value, for instance, could be considered the most central events and could be compared with similarly grouped events in other data sets. The trigger inefficiency is data set dependent, however, and this will affect the shape of the reference multiplicity distribution, ultimately shifting the top 5% boundary. To deal with this issue, the reference multiplicity distribution is fit to a Glauber Model simulation [40]. This model simulates the number of tracks produced within a detector's acceptance based on the energy, species, and impact parameter of the colliding nuclei. It contains tunable parameters that allow the simulated reference multiplicity distribution to be adjusted to match the measured one at moderate to large values of the reference multiplicity where the trigger inefficiency is negligible. Once the distributions match, the Glauber distribution acts as a map between the measured reference multiplicity and the underlying impact parameter. This allows for centrality definitions which are consistent between data sets.

2.7 Pile-Up

The probability of two gold ions colliding in a bunch crossing is relatively low, on the order of one collision per 10,000 bunch crossings. This makes the probability of two collisions within the same bunch crossing very low, one per $10,000^2$ crossings. However, with a crossing rate of 9 MHz, this leads to the expectation of a single bunch crossing with two collisions occurring every 10 seconds. Two collisions being recorded within the same event is called a pile up event and can cause issues in analysis, as the two collisions may be interpreted by analyzers as one large event.

The problem of pile up is not quite as bad as the approximated p^2 . In addition to time information, the three-dimensional collision vertex is utilized to identify events with multiple vertices. In the case of two rather central collisions, if the collision vertices are separated in the beamline either longitudinally or radially by a few centimeters, the vertex reconstruction algorithm may have enough resolution to determine each vertex very precisely. If this precision exceeds the distance between the two vertices, the vertices can be successfully differentiated. The event can then be flagged as pile up and removed from the analysis. The situation becomes more difficult for the vertex algorithm when one collision is much more central than the other, in which case the less central collision is not identified as a separate vertex and its tracks are appended to the event along with those from the more central collision. Even in this case, however, analysis cuts on the Distance of Closest Approach (DCA) will be able to reject most of the less central collision's tracks if its vertex is spatially separated from the more central one by more than a few centimeters. If two collisions occur close in both time and in space, there is almost no hope of differentiating them. However, the probability of this type of extreme pile-up is very small.

While some pile-up events can be rejected based on their vertex positions during event reconstruction, more can be identified by their timing after the fact. We can compare the total event multiplicities recorded by the slow TPC with multiplicities recorded by fast

detectors, such as the TOF. In a normal event, the fast and slow detector multiplicities should be proportional within statistical fluctuations. However, when two collisions take place close in time, the fast detectors may be able to differentiate these two collisions while the slow detector integrates the tracks from both collisions. In this case, we would expect to see a larger multiplicity in the slow detector compared to what would be expected from the fast detector's multiplicity. We can then plot the correlation between the fast and slow detector multiplicities and define a cut on this plot to reject events in which the slow detector records a much larger multiplicity than the fast detector. This method is effective for removing pile-up from collisions occurring in separate bunch crossings, called out-of-time pile-up. If two collisions occur within the same bunch crossing, which was estimated to occur every 10 seconds, even the fast detectors will not be able to resolve the time difference and this in-time pile up cannot be rejected by its timing information.

CHAPTER 3

Fluctuations in Azimuthal Partitions

Previous STAR analyses have measured event-by-event net-proton multiplicity distributions about the full azimuth to search for an enhancement of fluctuations expected in the vicinity of a critical point. If these fluctuations exist across the pseudorapidity boundary, they should also exist in azimuthal space. We investigate the feasibility of measuring proton multiplicity within azimuthal subvolumes of each event. Correlation between proton tracks can be inferred by measuring azimuthal partition multiplicity distributions and comparing their moments to the uncorrelated binomial expectation. This chapter will focus on defining our methodology for constructing azimuthal partitions and demonstrating that our method works as expected when random tracks are generated. We will then show that we are free to take many overlapping partition measurements within the same event to improve statistics. We end the chapter with an explanation of our statistical uncertainty estimations and a demonstration of their efficacy.

3.1 Azimuthal Partitions

An azimuthal partition of width w is randomly placed on each event. Proton tracks are identified within $y < |0.5|$ and $0.4 < p_T < 2.0$ GeV. For each of these protons, the azimuthal angle of the initial trajectory is calculated. The number of protons whose azimuthal angle lies within the partition are counted. A diagrammatic example of this procedure is shown for a mock event on the left side of Figure 3.1. This procedure is repeated over many events containing the same number of total number about the entire azimuth (N) and an azimuthal

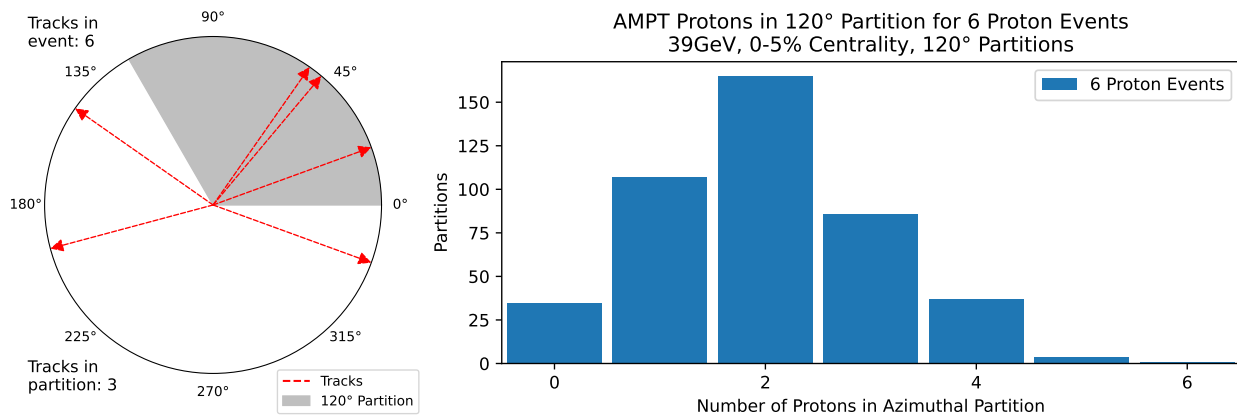


Figure 3.1: Left: Example event with 6 total proton tracks including 3 within the 120° azimuthal partition. The direction of the initial momenta of the tracks are represented as red arrows, and the gray area represents the azimuthal partition. Right: Example azimuthal partition multiplicity distribution for 500 events with 6 total protons.

partition multiplicity distribution is constructed from the number of protons within the azimuthal partition, an example of which is shown on the right of Figure 3.1.

To get a better intuition for the azimuthal multiplicity distributions, we can imagine placing each proton track in an event on the azimuth one by one. If the proton tracks were entirely uncorrelated with one another and placed at random, the probability of a track falling within an azimuthal partition of width w is equal to the fraction of the azimuth covered by this partition. Each track placed on the event would then correspond to a Bernoulli trial with probability of success $p = \frac{w}{2\pi}$. Placing N total tracks on an event would be a series of Bernoulli trials and would produce a binomial distribution. We would then expect that, for events which contain N total proton tracks on the full azimuth, the azimuthal multiplicity distribution produced would be a binomial distribution of N trials and probably of success $p = \frac{w}{2\pi}$ in the case of uncorrelated proton tracks. If the protons are correlated, the distribution should systematically deviate from this binomial expectation. This correlation can be quantified by measuring the moments of the multiplicity distributions and calculating their deviation from the known binomial moments.

3.2 Demonstrating Convergence to Binomial

We will first demonstrate that the moments of azimuthal multiplicity distributions produced by random tracks do in fact converge to the expected binomial moments. We will run a simulation in which N random azimuthal angles are generated in each event and then measure the azimuthal multiplicity distributions in these events. The true pdfs for these multiplicity distributions are only resolved as the number of entries goes to infinity. We therefore examine moments of the multiplicity distributions as a function of the number of entries (or number of events) and confirm that they do indeed converge to the expected binomial moments as the statistics are increased.

For simplicity, we consider only events with exactly $N = 15$ tracks and use an azimuthal partition width of $w = 120^\circ$, though the conclusions will generalize to all N and w . Generating random events and taking a single partition multiplicity sample from each, Figure 3.2 shows the distribution at different points in the simulation. Here, 100,000 total events are generated and the distributions are snapshots as the events are accumulated. The three distributions are therefore correlated. There is visible deviation from binomial when the statistics are low but the distribution approaches binomial as the number of events increases.

To better quantify this convergence, one option is to calculate the moments of the azimuthal partition multiplicity distribution and compare to the known values of the expected binomial. Limiting ourselves to the variance for now, Figure 3.3 shows the variance of the simulated distributions (the same simulation shown in Figure 3.2) vs the number of events in the distribution. The error bars here are calculated via the Delta Theorem (Section 3.4.1) and decrease as the number of entries in the distribution increases. The variance of the simulated distribution shows some scatter about the true binomial value when statistics are low and appears to converge to the true value as more events are generated. It is important to reiterate that 100,000 total events are generated and that each point on the x-axis is a snapshot of the distribution as events are added. The points are therefore correlated (adding




Figure 3.2: 100,000 events are generated with 15 random tracks in each. One random azimuthal partition of 120° is taken in each event. The azimuthal partition multiplicity distribution is plotted along with the expected binomial distribution at 100, 1000, and 100,000 events to illustrate the convergence. Note that 100,000 total events are generated and these three plots are snapshots of the distribution as it fills.




Figure 3.3: 100,000 events are generated with 15 random tracks in each and one random azimuthal partition of 120° is taken in each event. The variance of the distribution is plotted in blue against the number of events in the distribution. The expected binomial variance is plotted as a red dashed line. The error bars are calculated via the Delta Theorem.

a new event will only shift the variance slightly due to the inertia of the distribution).

We can see that, at least for this 100,000 event simulation, the distribution appears to converge to the expected binomial as more events are generated. However, it is apparent that this particular distribution approaches the binomial variance from above and, at points along the x-axis, appears to significantly deviate. This systematic overshoot of the width of the distribution would be problematic if it were a general feature of the analysis methodology. In order to investigate, we can produce a similar plot but, instead of generating 100,000 events and plotting the variance of a single distribution as each new event is appended, we generate an entirely new set of events at each point along the x-axis. In this scenario, the x-axis values are no longer correlated and are expected to scatter randomly about the true value, equally likely to overshoot the width as undershoot. The variance values from this new simulation are shown in Figure 3.4.

Here we are relieved to find that the simulation's variance values do indeed scatter both above and below the expected value and therefore indicate no obvious bias in the analysis. Due to the scatter of points both above and below the true value, though, the convergence of

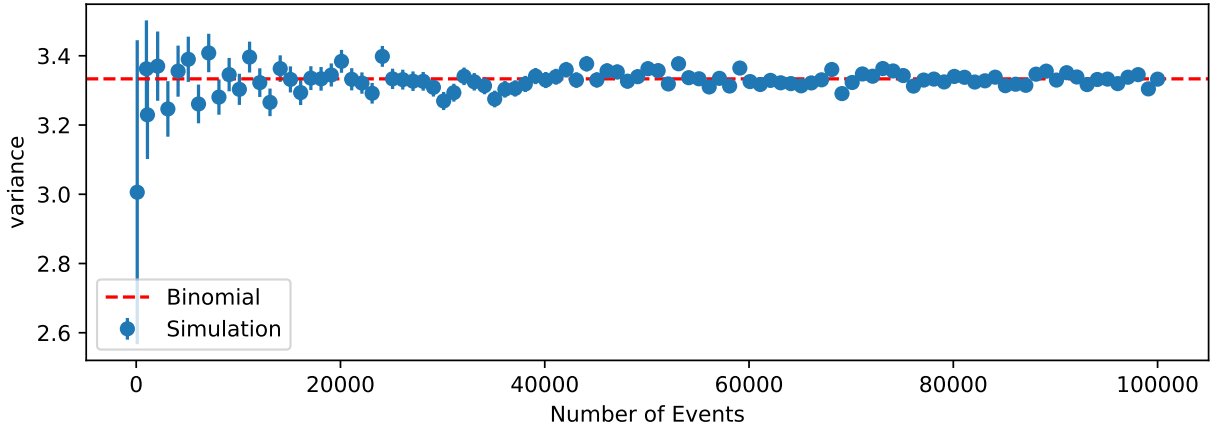


Figure 3.4: For each point, x events are generated with 15 random tracks in each. One random azimuthal partition of 120° is taken in each event. The variance of each independent distribution is plotted in blue against the number of events in the distribution. The error bars are calculated via the Delta Theorem and the expected binomial variance is plotted as a red dashed line. Each point is generated independently of all others.

the simulation variance to that of the expected binomial can be difficult to see in Figure 3.4. This can be quickly remedied by plotting the absolute value of the deviation of the variance from the binomial value $\Delta = |\sigma^2 - \sigma_{binom}^2|$ in Figure 3.5.

As each point is generated independently, there is considerable scatter in Figure 3.5, though it is somewhat apparent that the deviation from binomial is converging to zero as the number of events is increased. Looking at a single point on the x-axis of Figure 3.4 corresponding to a distribution with m events, we expect that if we were to repeat this measurement, we would obtain a spread of variances scattered about the true binomial value. We would expect this spread to increase for lower statistics and narrow for a larger number of events. Indeed, we can perform this simulation, running multiple “experiments” at a fixed m events each. The resulting distributions of variances are shown in Figure 3.6 for $m = (100, 1000, 100000)$. As expected, the widths of these distributions decrease as the number of events per experiment is increased. This test shows qualitatively that the measured values of the variance converge to that of the expected binomial and the width of this distribution of variances decreases as the number of events increases.




Figure 3.5: For each point, x events are generated with 15 random tracks and one random azimuthal partition of 120° is taken in each event. The absolute value of the variance of each independent distribution minus the binomial variance is plotted in blue vs the number of events in the distribution. The error bars are calculated via the Delta Theorem.

In order to better quantify this convergence, we can compute the absolute difference between the simulated distributions' variances and the variance of the binomial $\Delta = |\sigma^2 - \sigma_{\text{binomial}}^2|$. Figure 3.7 shows that the means of these Δ distributions converge to zero and their widths narrow as the number of events increases. We plot this convergence explicitly in Figure 3.8, which shows the mean of the Δ distribution at each point as a function of the number of events per distribution. The dark inner band represents the standard error on this mean while the light outer band represents the standard deviation of the Δ distribution.

The convergence of the variance with the number of events shown in Figure 3.8 tells us that, given enough statistics, the width of the azimuthally partitioned multiplicity distributions for the simulated data matches the width of the binomial distribution we expect to be emulating. This gives us confidence that, at least in this test case, our analysis methodology is working correctly when one azimuthal partition sample is taken per event. The variance is only one moment of a distribution, but we will demonstrate in Section 3.3.0.1 that the skewness and kurtosis in these simulations behave similarly.




Figure 3.6: Events are generated with 15 random tracks and one random azimuthal partition of 120° is taken. In each panel, m events are generated 300 times (“experiments”). The distribution of variances of the azimuthal partition multiplicity distribution for each experiment is displayed as a histogram. The mean, standard deviation and standard error on the mean ($\sigma/\sqrt{n_{\text{experiments}}}$) for these distributions are shown along with the expected binomial variance.




Figure 3.7: Events are generated with 15 random tracks and one random azimuthal partition of 120° is taken. In each panel, m events are generated 300 times (“experiments”). The distribution of the absolute difference between the variance for each experiment and the expected binomial variance is displayed as a histogram. The mean, standard deviation and standard error on the mean ($\sigma/\sqrt{n_{\text{experiments}}}$) for these distributions are shown.

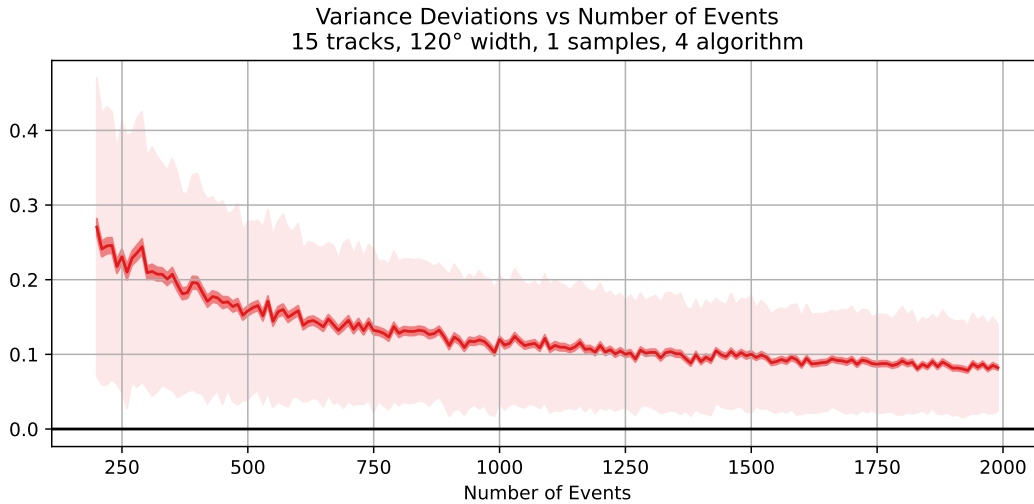


Figure 3.8: Events are generated with 15 random tracks and one random azimuthal partition of 120° is taken. At each point on the x-axis, 300 experiments of m events are generated and the deviations of the variances of the partitioned multiplicity distributions from the expected binomial variance (Δ) are calculated. The mean of the Δ distribution at each point on the x-axis is plotted, along with a dark band representing the standard error on this mean and a light band representing the standard deviation of the Δ distribution.

Standard Error on Mean vs Standard Deviation While the variance of the Δ distributions is informative, what we are really interested in is the convergence of the mean. For this reason, the standard error on the mean should be viewed as the uncertainty on the convergence of this trend. To expound upon this point, we can imagine conducting a single experiment at each point on the x-axis. This is the scenario presented in Figure 3.5. This gives an idea of the true convergence behavior we are trying to measure, but each point is only a single sample from the underlying Δ distributions shown in Figure 3.7. By running many experiments at each point on the x-axis, we are able to measure the true mean of this convergence behavior (mean of the Δ distribution) to arbitrary precision, as indicated by the $1/\sqrt{n_{\text{experiments}}}$ factor in the standard error on the mean definition. The ability to resolve this convergence at a precision limited only by computational power is the strength of running this type of simulation. This is opposed to a real world scenario in which any measurement will be limited by the statistics of the actual data taken.

3.3 Event Resampling

When the azimuth of an event is partitioned, it is necessary to specify not only the width but also where on the azimuth the partition will be placed. The STAR detector was designed to have full azimuthal coverage which would ideally be symmetric. Ignoring the realistic inhomogeneties for now (revisited in Section 4.5), the symmetry of the detector makes any placement of the partition an equally good choice. One option would be to arbitrarily choose a position on the detector and use that for each event. Another would be to randomly choose a position for each event. Either of these options should work, as the symmetry of the detector and of the collisions (on average) should render the choice inconsequential. If the partition placement in an event is arbitrary, it may be possible to take more than one of these samples per event. For instance, in the case of 120° partitions, the azimuth could be divided into exactly three partitions which could provide three separate multiplicity measurements per event. Taking additional partitions seems expedient, as more of the detector data is utilized which would naively improve statistics.

Our studies have found that algorithms taking multiple partitions per event do indeed improve statistics while also working as expected in simulations. In fact, the partitions taken in an event can even overlap as long as care is taken when estimating the statistical uncertainties. The algorithm developed for obtaining azimuthal partition multiplicity samples from an event starts by randomly generating s azimuthal angles between $[0, 2\pi)$. These random values are then sorted from least to greatest and represent the lower partition edge of each of s partitions. For each of these, the proton multiplicity between the lower edge (the random value from the list) and the higher edge (the partition width w plus the lower partition edge) is counted. In this way, s multiplicity measurements are made in each event. In practice, s is optimized such that the most information can be extracted from each event while keeping the computational cost reasonable. This optimal value used in the analysis was found to be $s = 72$.

3.3.0.1 Demonstrating Binomial Convergence with Event Resampling

We will demonstrate the viability of event resampling in a manner analogous to the case of a single sample per event described in Section 3.2. We will again generate events of $N = 15$ uncorrelated tracks and take azimuthal partitions of width $w = 120^\circ$. We will increase the number of experiments per data point to $n_{\text{exp}} = 1000$ to increase resolution. Figure 3.8 demonstrated that, as the number of events in the azimuthal partition distributions increased, the variance of the distributions monotonically approached the expected binomial value. In the top panel of Figure 3.9 we show this same convergence when taking 2 and 3 samples per event. We find that, in all three cases, the variance of the distributions converges to the binomial expectation. In addition, it appears that the larger the number of samples, the closer the distribution variance is to the expected binomial value for each number of events on the x-axis. This shows that, given a fixed number of events (as in experimental data), one can get a more precise measurement of the true variance of the distribution by taking multiple samples per event. The middle and bottom panels of Figure 3.9 show the same distributions with the skewness and non-excess kurtosis calculated for each distribution instead of the variance. As discussed at the end of Section 3.2, the standard deviation of the Δ distributions is not very useful, so from here on its representation as a light band in these plots will be omitted.

Non-Excess Kurtosis The kurtosis of a distribution is typically defined as the fourth central moment over the second central moment squared minus 3. This is typically called the “excess kurtosis” as in this formulation $\kappa = 0$ corresponds to the kurtosis of a Gaussian distribution. The excess kurtosis is used often enough that it is typically just referred to as the kurtosis. The range of the excess kurtosis is from -3 to ∞ , which can be slightly inconvenient. We will instead use the kurtosis without the -3 which we will call the non-excess kurtosis.

$$\begin{aligned} \kappa = \text{kurtosis} &= \frac{\mu_4}{\mu_2^2} - 3 \\ \text{non-excess kurtosis} &= \frac{\mu_4}{\mu_2^2} \end{aligned} \tag{3.1}$$

It is clear that the precision of measured moments increases as more samples are taken per event in the case of 1, 2, and 3 samples with a partition width of 120° . It is prudent to now check if this trend continues as the number of samples per event increases past 3. For this test, we'll pick a fixed number of events to work with and plot the mean of the Δ distributions vs the number of samples taken per event. Figure 3.10 shows that for 250 events with 120° partitions, the precision of the moments continue to increase as we take more samples per event. However, this trend appears to saturate rather quickly, with the variance accuracy flattening out around 17 samples, the skewness around 22, and the kurtosis around 25 (very rough and subjective estimates made by eye). It is interesting to note here that the higher moments seem to saturate later, implying that taking more samples per event may be more helpful in resolving higher moments.

3.3.0.2 Event Resampling Conclusions

Leveraging our understanding of events with totally uncorrelated tracks, we were able to show that, as the number of events simulated increases, the event re-sampling algorithm works as expected. We then demonstrated that increasing the number of samples taken per event increases the precision of the measured moments of the azimuthal partition multiplicity distributions. This increase in precision provides diminishing returns and, for all three moments checked, saturates at ten or twenty samples. Since the computational cost of the stochastic algorithm increases linearly with the number of samples, we have decided to use $s = 72$ samples per event as a standard for our analysis.

Stats Deviations vs Number of Events
 15 tracks, 120° width, 4 algorithm




Figure 3.9: Events are generated with 15 random tracks in each and [1,2,3] random azimuthal partitions of 120° are taken. At each point on the x-axis, 1000 experiments of m events are generated and the deviations of the variances (skewness for middle panel and kurtosis for bottom panel) of the partitioned multiplicity distributions from the expected binomial variance (Δ) are calculated. The mean of the Δ distribution at each point on the x-axis is plotted, along with a dark band representing the standard error on this mean and a light band representing the standard deviation of the Δ distribution.

Stats Deviations vs Number of Samples
 15 tracks, 120° width, 250 events, 4 algorithm




Figure 3.10: Events are generated with 15 random tracks with azimuthal partitions of 120°. At each point on the x-axis, 1000 experiments of 250 events are generated taking x samples per event and the deviations of the moments of the partitioned multiplicity distributions from the expected binomial moment (Δ) are calculated. The mean of the Δ distribution at each point on the x-axis is plotted, along with a dark band representing the standard error on this mean.

3.4 Statistical Error Estimation

When more than one azimuthal partition is taken per event, the multiplicity values from the same event are not independent. In this case, care must be taken when computing statistical uncertainties. Many methods for statistical approximation include an assumption that the data is independently sampled. When this is not the case, as in our analysis, numerical methods can be utilized instead.

3.4.1 Delta Theorem

For the case of a single multiplicity sample per event, each sample is independent of all others in the distribution because the events are independent of each other. In this case, the Delta Theorem can give an accurate approximation for the uncertainties on moments of the distribution. The equations for uncertainty on the moments along with other detail on the Delta Theorem can be found here [41].

3.4.2 Bootstrap

In order to deal with the case of multiple samples per event, the numerical bootstrap method is utilized. While traditional bootstrap techniques assume independent samples, the block bootstrap method was designed to deal with non-independent data.

Traditional Bootstrap The traditional bootstrap [42] takes a data set of \tilde{N} entries (to distinguish from N used as total number of protons per event) and constructs M “bootstrap sets” by randomly sampling \tilde{N} entries from the original set for each bootstrap set. This sampling from the original set is done at random, so it is expected that the same entry may be included more than once in a single bootstrap set while other entries are not included at all. In this way, M bootstrap data sets of size \tilde{N} are constructed from the original data. These bootstrap sets are each analyzed identically to the original set. When it is time to

estimate statistical uncertainties, the analysis of the original set is interpreted as the true measurement while the variation among the bootstrap sets is taken as a measure of the statistical uncertainty.

To give an example, in this analysis an azimuthal partition multiplicity distribution can be constructed by taking one multiplicity sample per event from many events. Given \tilde{N} events, this distribution is the “default”. M bootstrap distributions are constructed by randomly sampling \tilde{N} multiplicities (or, equivalently events, since only one sample is taken per event in this case) from the original set. If we are interested in a particular moment of the distribution, we can calculate its value in the default distribution as well as in each of the bootstrap distributions. The moment value from the default distribution will be the “measured” value. The moments calculated from each of the bootstrap sets produce a range of moment values from which a confidence interval can be constructed to serve as the uncertainty on the measured value.

Bootstrap Confidence Intervals Quite a few options exist for estimating the standard error on a measurement from a bootstrap set. A relatively simple method is to calculate the sample standard deviation of the distribution of bootstrap values and use this as the standard error on the measurement. This method is used frequently in physics analyses and has the benefit of being a quick and easily implemented calculation. In addition, the error is taken to be symmetric, which can make further analysis with this error easier. This assumption of symmetry can also be a detriment, however, if the bootstrap distribution is non-gaussian and the shift/skew is important information to retain.

When a more careful estimation of the standard error is required that can optionally accommodate non-symmetric error bars, a confidence interval can be constructed for the bootstrap distribution. Using a confidence interval of 68% provides a range equivalent to the standard error. Several methods for constructing confidence intervals on bootstrap sets have been developed since the birth of bootstrap [42]. The simplest method is the Percentile

Bootstrap Confidence Interval (PBCI) in which the boundaries of the symmetric 68th percentile provide the interval bounds. This method tends to have slight bias (the center of the interval is shifted from the ideal case) which has been addressed via more complex constructions such as the Bias-Corrected Bootstrap Confidence Interval (BCBCI). While these other constructions contain less bias, it is often found that this comes at the expense of worse coverage [43]. It is found that PBCI typically has very accurate coverage at the expense of its slight bias. This bias-coverage trade-off means that no single method is objectively better than others. Ideally, one would test each method in on a toy model for their analysis to justify the final choice.

In the case of our analysis, we use the standard deviation of the bootstrap set to estimate our standard errors.

Block Bootstrap For data sets in which the elements are not independent, the traditional bootstrap method fails. If the data set can be divided into subsets of independent “blocks”, however, these blocks can be sampled together in the bootstrap scheme [44]. This “block bootstrap” was developed historically to deal with the inherent correlations in time-series data where the blocking is only an approximate solution. In the case of this analysis, each event is independent of every other. Our blocks are therefore events and in constructing bootstrap sets entire events are sampled. When taking more than one partition per event, all multiplicity values from an event are appended to the bootstrap distribution when the event is selected. This block bootstrapping algorithm works identically to the traditional method except that whole independent blocks of values are randomly sampled instead of the values individually.

Poisson Bootstrap Sampling In constructing a bootstrap set, the traditional bootstrap method randomly samples from the original set until the bootstrap set is of the same size (has the same number of elements) as the original. This is typically accomplished by randomly

selecting indices on $[0, N - 1]$ for data with \tilde{N} values until \tilde{N} values have been selected. In practice, if each sampling operation is expensive, bootstrapping in this way can be computationally prohibitive when constructing a large number of bootstrap sets. Colloquially, this is the burden of having to run through a data set multiple times.

A computationally cheaper algorithm that approximates constructing a bootstrap set in the traditional way is called Poisson Bootstrapping [45]. Here, the original data set is run through only once. Each element (or block) is added to a specific bootstrap set r times where r is an integer sampled from a Poisson distribution with a mean of 1 — $p(r) = \frac{1}{e r!}$. This emulates randomly sampling from the data set while only requiring each element or block to be constructed once.

In this analysis, each event is an independent block for the bootstrap. Each event is processed and the multiplicities are counted in each of the azimuthal partitions. For each bootstrap set, an integer r is sampled from a Poisson distribution with mean 1 and all of the partition multiplicity values from the event are added to the bootstrap distribution r times. Using this Poisson Block Bootstrap technique allows for the computationally feasible estimation of uncertainties in data for which the individual elements are not independent.

3.4.3 Demonstrating Consistency of Statistical Error Estimation

Section 3.4 introduced the Delta Theorem and the Bootstrap method for estimating statistical uncertainties on the moments of multiplicity distributions. This section will utilize the binomial simulations described in Section 3.2 to show that the uncertainty estimates for this analysis are consistent and provide proper coverage with little bias or skew. We will first show that the Delta Theorem works in the case of taking a single sample per event but begins to under cover as the number of samples per event are increased. We will then show that the Poisson Block Bootstrap estimate produces the desired uncertainties in the case of multiple samples per event.




Figure 3.11: Example azimuthal partition multiplicity distribution of $m = 10000$ events with $N = 15$ randomly placed tracks and 120° partitions. The x-axis is the number of tracks within the azimuthal partition and the y-axis is the number of events with this multiplicity. The blue bars are a histogram of the simulated events and the red dashes show the expected binomial distribution.

3.4.3.1 Delta Theorem Demonstration

The Delta Theorem (Section 3.4.1) provides an analytical approximation for the uncertainties on the moments of a distribution when the entries in that distribution are independent. This is the case in our analysis when only one sample per event is taken. To demonstrate the validity of the Delta Theorem in estimating statistical uncertainty in this scenario, we will generate many “experiments” consisting of m generated events each. The tracks will be randomly distributed about the azimuth, in which case the azimuthal partition multiplicity distributions should sample from corresponding binomial distributions. An example azimuthal partition multiplicity distribution for a single experiment of $m = 10000$ events is shown in Figure 3.11. The moments of this distribution are then calculated along with their corresponding Delta Theorem uncertainties.

For a single distribution, it’s difficult to evaluate the efficacy of a statistical uncertainty

estimate. We can compare the value of the moment to the expected binomial value to get an idea of how accurate the simulation is, but the uncertainty can't be evaluated quite as easily. Instead, we need to consider what the uncertainty estimate represents. In an ensemble of identical experiments, the 1σ uncertainties estimated in these experiments are defined to cover the true value 68% of the time. The uncertainty estimates are typically interpreted as normally distributed about the estimated value, so twice the estimated uncertainty should correspond to 2σ and cover the true value 95% of the time. This generalizes to x times the estimate covering the area within $x\sigma$ of the mean under a standard normal PDF. The number of standard deviations from the mean is defined as the Z-value — $Z = \frac{x-\mu}{\sigma}$.

To verify the accuracy of an uncertainty estimate, we must generate an ensemble of experiments and ensure the estimates cover the true value the proper proportion of the time. To build intuition, we can consider an example of 100 identical experiments of random tracks. In this case, whatever moment of the distribution is measured should be distributed as a Gaussian about the true (binomial) value. The 1σ uncertainty bands for these experiments should cover the true value in about 68 of the experiments. Likewise, twice the uncertainty should cover in about 95 experiments. If the measured moment values indeed have a symmetric, Gaussian distribution about the true value and more experiments cover the true value than expected, this tells us that the uncertainties have been overestimated. Similarly, if fewer experiments cover than expected this suggests that the uncertainties have been underestimated.

The test above is carried out for our analysis, utilizing events with random tracks. Figure 3.12 shows the variance measurements and their associated statistical uncertainties calculated via the Delta Theorem for each experiment as a scatter plot. We can then calculate the Z-value for each of these experiments. A histogram of these Z-values should produce a standard normal distribution if the uncertainties are accurate. The top panel of Figure 3.13 shows a histogram of Z-values for these 100 experiments along with an overlaid standard normal curve. The statistics of only 100 experiments are limiting the comparison resolution

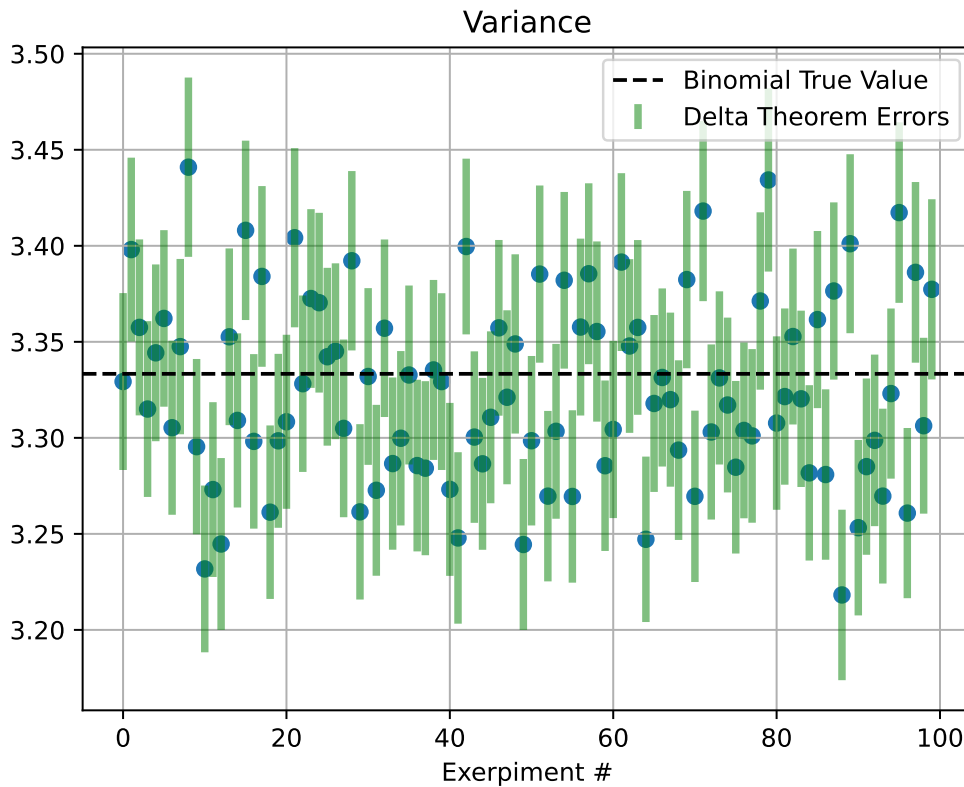


Figure 3.12: The variance of 100 simulated distributions plotted against the arbitrary experiment number index on the x-axis. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. The true binomial value is plotted as a dashed black line and the green error bars are calculated via the Delta Theorem.

but, at this level, the Z-value distribution appears relatively normal. The middle and bottom panels of Figure 3.13 show the Z-scores for the skewness and kurtosis. It is noted that all three moments see a significantly negative shift of about 2σ significance which we believe to be a statistical fluctuation in the small 100 experiment sample size. The kurtosis also appears to have a slightly longer negative tail than expected, which may be a significant though small asymmetry.

This test and the resulting Z-score distributions tells us that the Delta Theorem is producing relatively unbiased estimates of the uncertainty on the variance of our distributions and that it also has proper coverage, up to the resolution that 100 generated experiments allows.

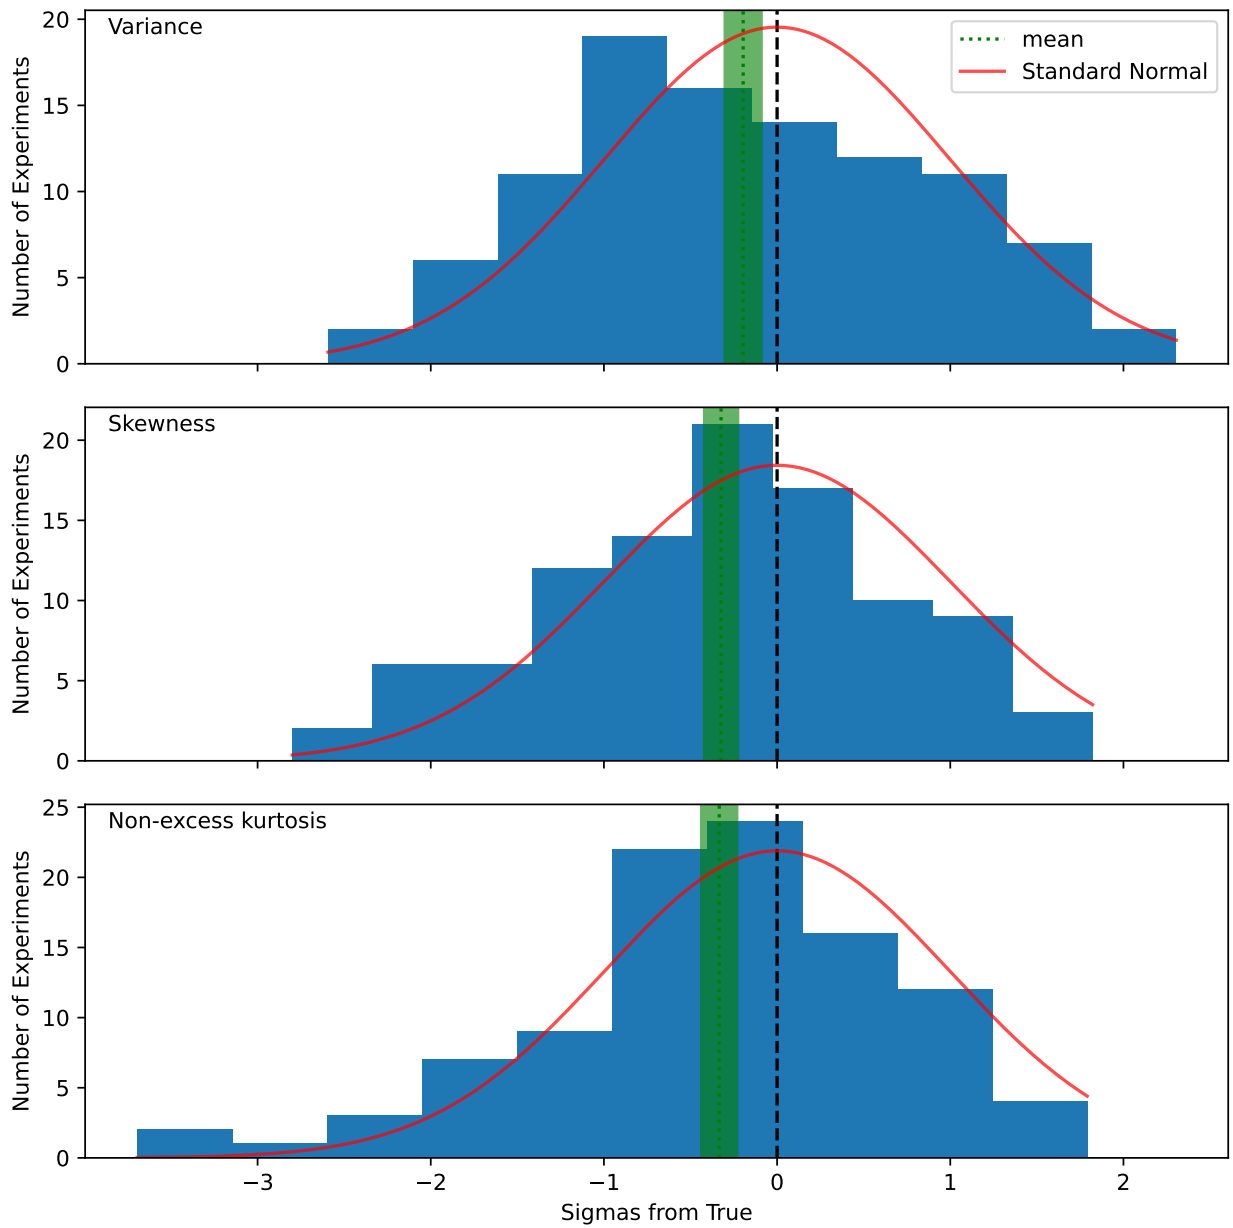


Figure 3.13: Z-value histograms of the moments of 100 simulated distributions. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. The mean of these distributions is plotted as a dotted green line and the green band about it is the standard error on this mean. A standard normal distribution scaled to the number of entries in the histogram is drawn in red.

3.4.3.2 Bootstrap Demonstration

Section 3.4.3.1 demonstrated that the Delta Theorem gives good estimates for the uncertainties on the moments of our distributions in the case of taking one sample per event. In the case of taking multiple samples per event, the entries in the distributions are no longer independent and the Delta Theorem is expected to break down. In this scenario, as discussed in Section 3.4, the Poisson Block Bootstrap technique is required to numerically estimate uncertainties on the distribution moments. We'll first demonstrate the bootstrap method in the case of one sample per event along side the Delta Theorem to show that they are comparable in most cases. Once faith in the bootstrap has been established in relation to the Delta Theorem, we will show how the Delta Theorem begins to under cover as the number of samples per event increases. We finally investigate the behavior of the bootstrap uncertainties for a large, realistic number of samples per event.

The Case of One Sample per Event While the bootstrap technique is employed in this analysis to deal with the non-independent samples when taking multiple samples per event, it should also work in the case of a single sample per event. In this scenario, we can calculate both the Delta Theorem and bootstrap uncertainties for a particular distribution and compare the estimates from the two methods. We will follow the discussion in Section 3.4.3.1 where confidence was established in the Delta Theorem estimates.

For a given distribution with \tilde{N} entries, the bootstrap uncertainty estimates are made by first constructing M sets through random sampling of \tilde{N} entries from the original distribution (the details of the Block and Poisson methods are described in Section 3.4.2). The moments of these M bootstrap distributions are then calculated. The bootstrap 1σ uncertainty can then be estimated as the standard deviation of these M moment values. $M = 250$ bootstrap sets are constructed by default unless otherwise stated.

For the distribution shown in Figure 3.11, the distributions of bootstrap moments for the variance, skewness, and kurtosis are shown in Figure 3.14. For all three moments,

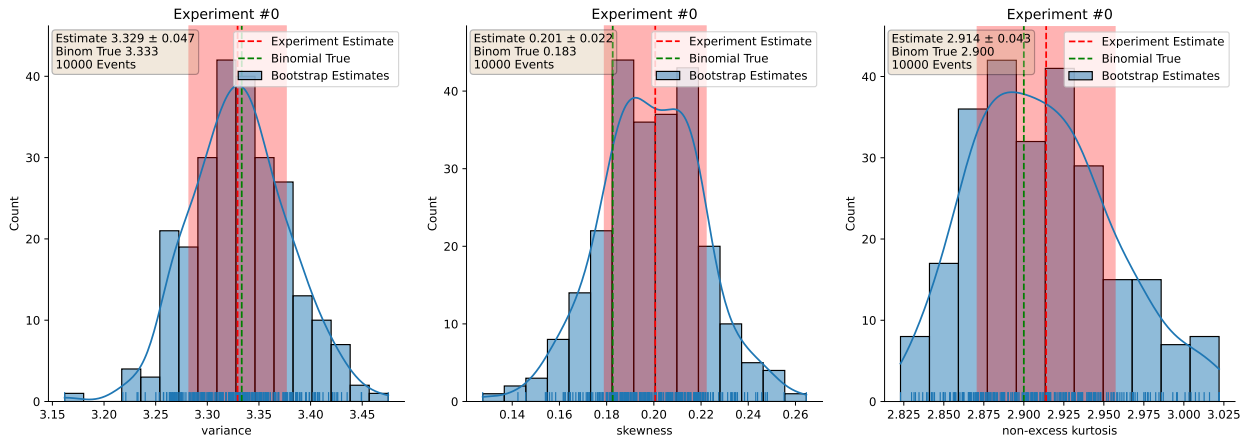


Figure 3.14: Distribution of bootstrap variance, skewness, and kurtosis values derived from the distribution in Figure 3.11. The dashed red line represents the mean of the distribution while the shaded red area corresponds to one standard deviation on each side of the mean. The green dashed line is the expected binomial mean. A rug plot of the individual bootstrap variance values is displayed along with a KDE to help visualize the true distribution.

the distribution of the 250 bootstrap moment values seem to form a reasonably normal distribution about the experiment estimate, which is relatively close to the binomial true value. To check if this is just statistical luck in Experiment #0, we need an ensemble of experiments.

Using the same 100 experiments generated for Section 3.4.3.1, we can now plot the moments for each experiment along with both the bootstrap and Delta Theorem uncertainty estimates. Figure 3.15 shows all three moments of each experiment's azimuthal multiplicity distribution plotted against the arbitrary experiment number index. The light green error bars again correspond to the Delta Theorem uncertainty estimates while the overlaid blue error bars now correspond to the estimates from bootstrap. The two types of estimate tend to agree very well in this example, to the point it is somewhat difficult to distinguish the two types of error bars. To help visualize the difference between the two uncertainty estimates, the left panel of Figure 3.16 shows the bootstrap uncertainties minus the Delta Theorem uncertainties for the variance as a function of the arbitrary experiment number. This plot shows a very balanced scatter about zero, suggesting no systematic differences (only




Figure 3.15: Moments of 100 simulated distributions plotted against the arbitrary experiment number index on the x-axis. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. The true binomial value is plotted as a dashed black line, the green error bars are calculated via the Delta Theorem and the blue error bars are calculated via bootstrap. The bootstrap and Delta Theorem uncertainties tend to agree well, so the difference in the two types of error bars can be difficult to see in this plot.

statistically random) between the two methods. The right side of Figure 3.16 expounds upon this conclusion by constructing a histogram of the differences which shows only a slight skew which is assumed to be due to low statistics.

How many Bootstrap Sets? We have used $M = 250$ bootstrap sets by default in this proof of concept and will continue to use this number in the data analysis. The more bootstrap sets constructed (the larger M) the more accurate the bootstrap uncertainty estimates will be. For an M too small, the uncertainties estimated by the bootstrap technique will




Figure 3.16: For the variances of the experiments shown in Figure 3.15, the bootstrap uncertainty estimate minus the Delta Theorem estimate is plotted as a function of the experiment number (left). A histogram is constructed from these differences (right) with a KDE and rug plot to help visualize the true distribution shape.

have larger statistical scatter about the true uncertainty. As Roger Barlow suggests [46], it's not productive to worry too much about the uncertainty on an uncertainty estimate or any recursion to that effect. However, it is prudent to understand how the accuracy of our uncertainty estimate is affected by the number of bootstrap sets we use. Running more bootstrap sets increases the computational expense on our analysis, so we will attempt to find a compromise between accurate uncertainty estimates and reasonable computational cost.

To quickly explore the dependence of our bootstrap uncertainty estimates on the number of bootstrap sets constructed, we compare the $M = 250$ case presented above to identical tests run with $M = 2500$ and $M = 25$. We first plot an analogue of Figure 3.15 for $M = 25$ bootstrap sets in Figure 3.17. From this plot, deviation between the bootstrap and Delta Theorem uncertainties is visible in some of the experiments, but is still very small. This is encouraging, as even as few as 25 bootstrap sets seems to produce rather good agreement with




Figure 3.17: Moments of 100 simulated distributions plotted against the arbitrary experiment number index on the x-axis. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. $M = 25$ bootstrap sets are constructed for each experiment. The true binomial value is plotted as a dashed black line, the green error bars are calculated via the Delta Theorem and the blue error bars are calculated via bootstrap. The bootstrap and Delta Theorem uncertainties tend to agree well, so the difference in the two types of error bars can be difficult to see in this plot.




Figure 3.18: For the variances of the experiments shown in Figure 3.15, a histogram is constructed from the bootstrap uncertainty estimates ($M = 25$ on left and $M = 2500$ on right) minus the Delta Theorem estimates. A KDE and rug plot are included to help visualize the true distribution shape.

the Delta Theorem. To quantify the agreement with the Delta Theorem, we can estimate the average deviation from the histogram in Figure 3.16 along with it's analogue for $M = 25$ and $M = 2500$ in Figure 3.18. Estimating from the Figures by eye and using the average Delta Theorem uncertainty for the variance of 0.046, we construct Table 3.1. We find that the spread on the bootstrap uncertainty estimates decreases roughly proportional to $1/\sqrt{M}$ and that the spread relative to the Delta Theorem at $M = 250$ is roughly 5%. A 5% variation on the variance uncertainty estimate is acceptable and we will continue to use $M = 250$ in future analysis to produce reasonably accurate uncertainty estimates while also keeping the computational cost under control.

M	Absolute Deviation	Percent Deviation
25	0.006	13.0%
250	0.002	4.3%
2500	0.0006	1.3%

Table 3.1: Average deviation of the bootstrap uncertainty estimates from the Delta Theorem estimates for bootstrap sets of size 25, 250, and 2500. The percent deviation is calculated by dividing by the average Delta Theorem uncertainty of 0.046.

Multiple Samples per Event Now that we have demonstrated the agreement between the bootstrap and Delta Theorem uncertainty estimates for the case of a single sample per event, we can investigate the behavior of these estimates when multiple samples are taken in each event. Event resampling is explained in Section 3.3 and its convergence to the expected value when simulating random tracks is demonstrated in Section 3.3.0.1. We will run the same test carried out in Section 3.4.3.2 but will now take $r = 72$ samples per event. The number of entries used to calculate the Delta Theorem uncertainties will be naively calculated in two ways:

1. Using the number of resamples r times the number of events m , which is the number of entries in each distribution.
2. Using the total number of events m .

We carry out these simulations with $r = 72$ and plot in Figure 3.19 the moments as a function of the arbitrary experiment number where the Delta Theorem uncertainties are estimated with $r \cdot m$ entries. The Delta Theorem assumes independent entries, which is not true in the case of event resampling. In effect, the “true number of independent entries” is smaller than $r \cdot m$, and the Delta Theorem uncertainties are therefore underestimated. This is clearly visible in the plot, where the Delta Theorem error bars are much smaller than the

bootstrap error bars and are clearly underestimating the uncertainties. When we ignore the resampling and assume only m entries for the Delta Theorem uncertainty estimates, we find in Figure 3.20 that the Delta Theorem uncertainties are now significantly overestimated. This makes sense, as the estimate now ignores the additional statistics we have gained by resampling the event. In fact, this plot provides a clear visualization of the statistical power gained by taking multiple samples per event as compared to a single sample.

We finally show in Figure 3.21 the Z-scores calculated using the bootstrap uncertainty estimates for the simulation run with $r = 72$. While running only 100 experiments leads to some fluctuations in these histograms, all three still appear to match relatively well to the overlaid normal distribution. This implies that the bootstrap uncertainty estimates continue to provide a decent estimation in the case of multiple samples taken per event.

3.4.4 Statistical Error Estimation Conclusion

We found that when taking a single sample per event, the Delta Theorem and the bootstrap method produce comparable and accurate uncertainty estimates. Using $M = 250$ bootstrap sets seems to be adequate to match the Delta Theorem estimates of the variance within 5%. When multiple samples are taken per event, the Delta Theorem uncertainty estimates break down but the bootstrap estimates remain accurate. The size of the bootstrap uncertainties relative to the Delta Theorem scaled uncertainties in Figure 3.20 allow us to visualize the additional statistical power gained when taking 72 samples per event instead of 1. These tests give us confidence in the accuracy of the bootstrap uncertainty estimates and illustrate the statistical benefit of taking multiple samples per event.




Figure 3.19: Moments of 100 simulated distributions plotted against the arbitrary experiment number index on the x-axis. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. $r = 72$ samples are taken in each event and the Delta Theorem uncertainties are calculated with $r \cdot m = 72 \cdot 10000$ entries. The true binomial value is plotted as a dashed black line, the green error bars are calculated via the Delta Theorem and the blue error bars are calculated via bootstrap.




Figure 3.20: Moments of 100 simulated distributions plotted against the arbitrary experiment number index on the x-axis. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. $r = 72$ samples are taken in each event and the Delta Theorem uncertainties are calculated with $m = 10000$ entries. The true binomial value is plotted as a dashed black line, the green error bars are calculated via the Delta Theorem and the blue error bars are calculated via bootstrap.

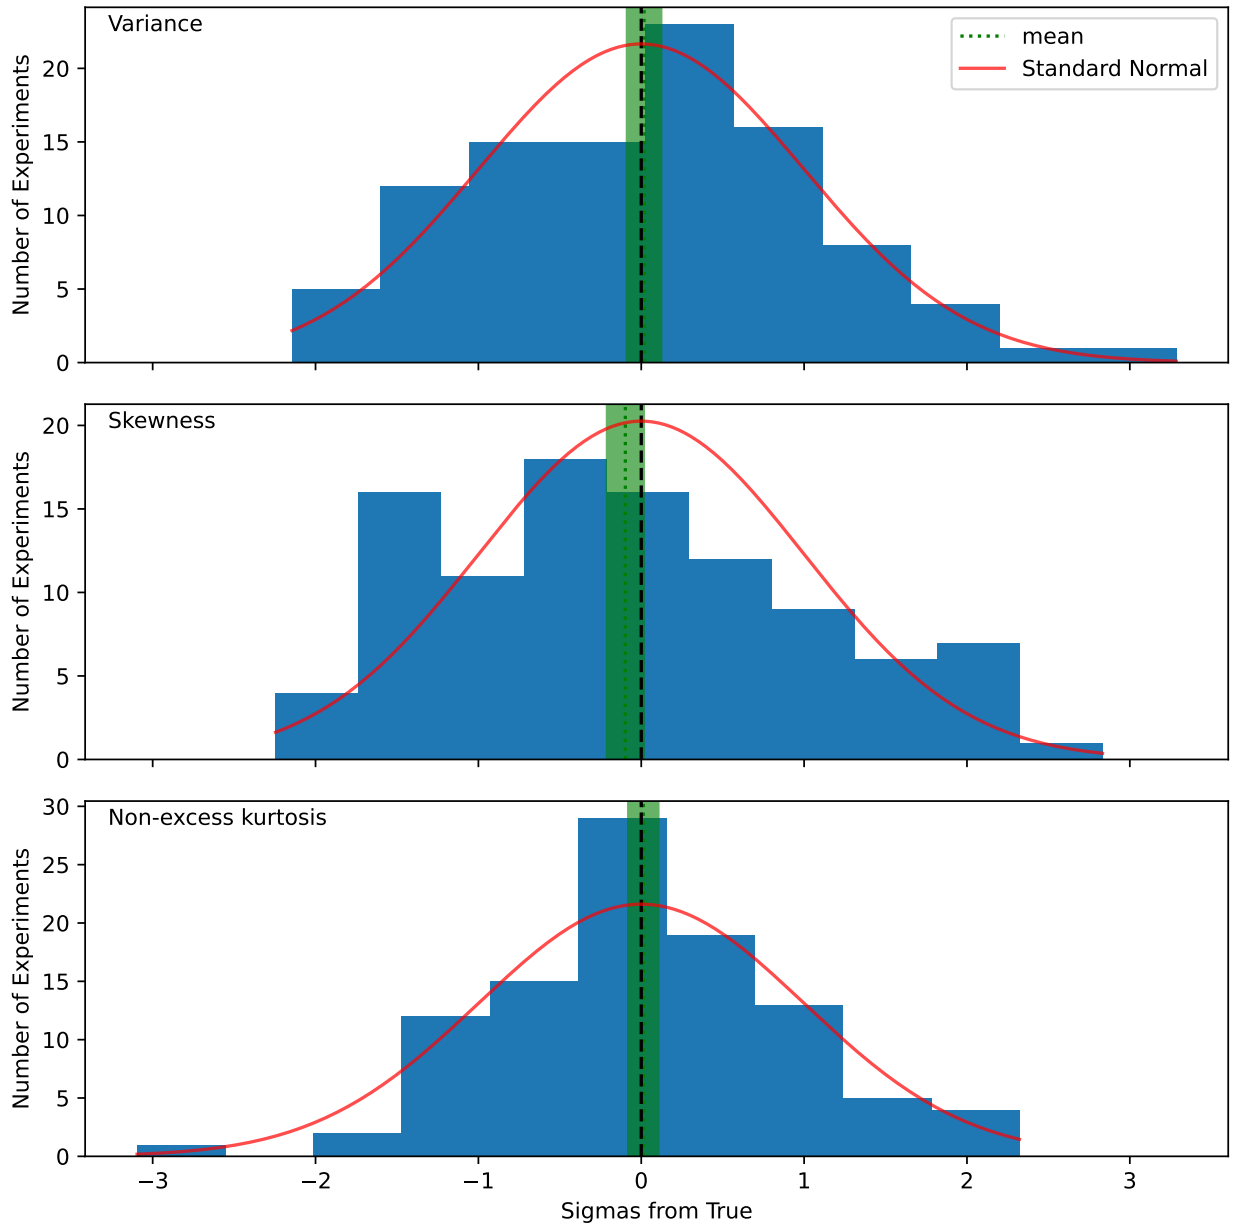


Figure 3.21: Z-values of the moments of 100 simulated distributions calculated from the bootstrap uncertainty estimates. In each experiment $m = 10000$ events are generated with $N = 15$ randomly placed tracks and 120° partitions. $r = 72$ samples are taken from each event. The mean of these distributions are plotted as a dotted green line and the green band about it is the standard error on this mean. A standard normal distribution scaled to the number of entries in the histogram is drawn in red.

3.5 Bias in Estimators

Naively measured moments of distributions are known to be biased. This bias can be eliminated or assuaged by correction factors, such as Bessel’s correction for the variance — $\frac{n}{n-1}$. In the case of Bessel’s correction, the bias in the variance is totally eliminated. However, when calculating the standard deviation as the square root of the variance, a bias returns. It has been shown that for certain statistics such as the standard deviation, an unbiased estimator does not exist. The bias can be reduced, however, and Bessel’s correction for the standard deviation accomplishes this. Reducing bias in an estimator tends to increase the MSE of the measured statistic. This is a manifestation of the bias-variance trade-off.

For the purposes of our analysis, we will quickly demonstrate the difference between biased and minimally biased estimators of the cumulants. The minimum bias estimators for the cumulants are called the Fisher k-statistics [47].

We generate events with 15 tracks randomly distributed about the azimuth and take one sample of a 120° partition. This will produce a binomial distribution whose true cumulant values are known. We calculate the cumulants and k-statistics for distributions of m events. If this measurement is made only once, the statistical fluctuations will be much larger than the systematic differences between the cumulant and k-statistic values. We therefore run 1000 “experiments” in which the random generation of m events and subsequent measurement of cumulants is performed. We then average these 1000 values and plot as a function of the number of events m in Figure 3.22, using the error on the mean as the uncertainty bands.

We find that the cumulants indeed systematically deviate from the true binomial value at all orders, reflecting their bias. This deviation is large for distributions with few events and converges to the true value as the number of events increases. The k-statistics are consistent with the true binomial value at all orders and independent of the number of events per distribution.

We will therefore use the k-statistics in our analysis to reduce bias. It is noted, how-




Figure 3.22: Cumulants in red and k-statistics in blue are plotted as a function of the number of events in the distribution. The second order cumulant (variance) is plotted in the top panel, the third order cumulant in the middle panel, and the fourth order cumulant on the bottom panel. The true value is plotted as a horizontal dotted line in each panel.

ever, that the cumulants converge rather quickly to the true value as the number of events increases. For a realistic distribution from data we expect at least a few hundred entries per distribution and therefore, in practice, the bias on the cumulants should be negligible compared to the statistical and systematic uncertainties.

3.5.1 Unbiased Estimator of Variance

The unbiased estimator for the variance is obtained via Bessel's correction (M represents the number of entries in the distribution):

$$k_2 = \frac{M}{M-1} \sigma_{\text{sample}}^2 = \frac{1}{M-1} \sum_i^M (x_i - \bar{x})^2 \quad (3.2)$$

When calculating the variance for the $\Delta\sigma^2$ observable, the second k-statistic, k_2 , as defined in Equation 3.2 will be used.

CHAPTER 4

The $\Delta\sigma^2$ Observable

Chapter 3 introduced the methodology for constructing azimuthal partition multiplicity distributions. In this chapter, we will investigate the variance of these distributions and define the $\Delta\sigma^2$ observable. We will discuss flow and azimuthal dependent efficiency as sources of background and then analyze the observable analytically. We will finally demonstrate the efficacy of background subtraction on $\Delta\sigma^2$.

4.1 Variance of Multiplicity in Azimuthal Partitions

The deviation of the measured azimuthal multiplicity distribution from binomial can be quantified via comparison of the variance between the two distributions. The variance of the binomial distribution is $\sigma_{\text{binomial}}^2 = Np(1-p)$ and the variance of the measured distribution can be calculated $\frac{1}{M-1} \sum_i (x_i - \bar{x})^2$. Uncorrelated protons will produce binomial multiplicity distributions within azimuthal partitions. In the case of an attractive interaction between protons, clusters of tracks will tend to form on the azimuth. When these clusters fall within an azimuthal partition, the multiplicity value will be relatively large. The presence of clustering necessitates the presence of voids on the azimuth, which will produce a relatively low multiplicity value when captured within a partition. Since the partitions are placed randomly on the azimuth, both clusters and voids will be captured which will produce a distribution wider than binomial with a correspondingly larger variance.

In the case of negative correlation between protons, the tracks will tend to spread out

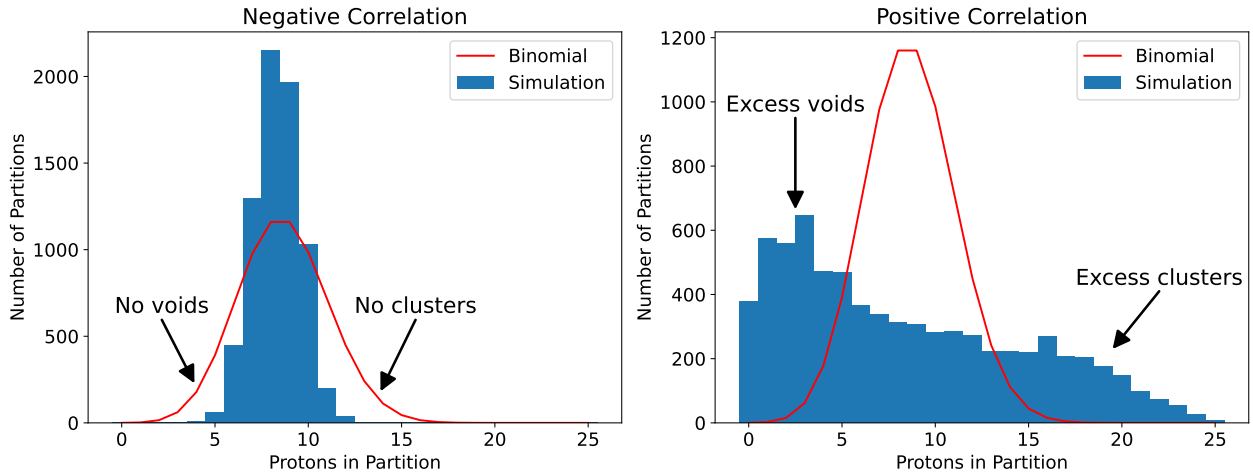


Figure 4.1: The azimuthal multiplicity distribution is compared to binomial in the case of extreme repulsion on the left and extreme attraction on the right. These distributions were simulated with the Gaussian Correlation Model from Section 5.2

within the available 2π of azimuthal space. In contrast to the attractive case, a repulsive interaction will produce fewer clusters and voids than would be expected randomly. As a result, there will be fewer extreme events resulting in a distribution narrower than binomial. This produces a variance smaller than binomial.

In summary, attractive interactions between proton tracks produce a variance greater than binomial while repulsive interactions produces a variance smaller than binomial. This is illustrated in Figure 4.1 where the azimuthal multiplicity distribution is compared with the binomial expectation for both extreme repulsion on the left and extreme attraction on the right.

4.2 Defining the Observable

We now compare the variance of the measured distributions (σ^2) to the binomial variance expected from uncorrelated tracks ($\sigma_{\text{binomial}}^2 = Np(1-p)$). We define the $\Delta\sigma^2$ observable as a function of the proton multiplicity over the full azimuth (N):

$$\Delta\sigma^2(N) = \frac{\sigma^2(N) - \sigma_{\text{binomial}}^2(N)}{N(N-1)} \quad (4.1)$$

A positive value of this observable indicates a positive correlation between tracks while a negative value indicates a negative correlation. A value of zero indicates that the measured variance equals the binomial variance and that the tracks are uncorrelated from the perspective of the distribution width.

4.3 Analytical Analysis of $\Delta\sigma^2$

An analytical expression for the variance of the number of proton tracks per partition can be obtained when each track is generated via a fixed probability as a function of the azimuthal angle. Under these conditions, placing a track on the event is a Bernoulli trial with probability of success equal to the integral of the azimuthal pdf between the partition edges, illustrated in Figure 4.2. For a general distribution $\tilde{p}(\phi)$ with azimuthal partition of width w and lower partition bound ψ , the probability of placing a proton within the azimuthal partition is given by Equation 4.2.

$$p(\psi) = \int_{\psi}^{\psi+w} \tilde{p}(\phi) d\phi \quad (4.2)$$

Fixing the lower bound ψ , the multiplicity distribution within azimuthal partitions will be binomial for N trials with the corresponding binomial mean and variance:

$$\mu_{\text{binom}}(\psi) = Np(\psi) \quad (4.3)$$

$$\sigma_{\text{binom}}^2(\psi) = Np(\psi)q(\psi) = Np(\psi)(1-p(\psi)) \quad (4.4)$$

We note here for future reference that, through a change of variables, ψ in Equation 4.2 can be expressed as a shift on ϕ :




Figure 4.2: Example of azimuthal probability distribution for elliptical flow with $v_2 = 0.1$ (Section 4.4). The gray area represents the integral of this probability distribution between partition edges.

$$p(\psi) = \int_0^w \tilde{p}(\phi + \psi) d\phi \quad (4.5)$$

While Equation 4.2 and Figure 4.2 suggest an interpretation of ψ as dictating the position of the azimuthal partition, Equation 4.5 demonstrates that it can equivalently be understood as a shift of the probability distribution relative to a fixed azimuthal partition. It will be convenient to switch between these interpretations of ψ for different calculations.

4.3.1 A Single PDF

In heavy ion collisions, even if we can approximate some type of track correlation as a fixed azimuthal probability distribution, this PDF will typically have an orientation relative to the reaction plane which rotates in each event. We will therefore investigate the effect of randomly rotating the azimuthal probability distribution for each event. Here, we imagine distributing N proton tracks per event and counting the number within the azimuthal partition. The azimuthal probability distribution is shifted randomly for each event. From Equations 4.2 and 4.5, shifting the azimuthal probability distribution is equivalent to shifting

the azimuthal partition by the opposite amount on a fixed pdf. We can therefore equivalently imagine that the lower partition bound, ψ , is randomly rotated between each event.

We now attempt to calculate the mean and variance of the multiplicity distributions within the azimuthal partitions over many events of randomly rotated ψ values. Over all M events, the mean is the usual average over the number of particles in each partition n_i :

$$\mu = \frac{1}{M} \sum_i^M n_i = \langle n \rangle \quad (4.6)$$

and the variance can be expressed as

$$\sigma^2 = \langle n^2 \rangle - \langle n \rangle^2 = \frac{1}{M} \sum_{i=1}^M n_i^2 - \left(\frac{1}{M} \sum_{i=1}^M n_i \right)^2 \quad (4.7)$$

We can next imagine a discrete set of possible ψ values and separate each event in the full ensemble into sets of events with the same ψ value (one group for $\psi = 1$, another for $\psi = 1.5$, ...). Each of these sets of events has a fixed ψ value and therefore samples from the same probability distribution (same phase). These individual sets then satisfy the previously discussed conditions to produce binomial multiplicity distributions within the azimuthal partition with corresponding binomial mean and variance (Equations 4.3 and 4.4). The probability of success will be equal to the integral of the fixed pdf between the partition edges, Equation 4.2, as before.

The sums in Equations 4.6 and 4.7 can be grouped by the sets of equal ψ , supposing there are R of these sets with m_i events each.

$$\mu = \frac{1}{M} \sum_i^M n_i = \frac{1}{\sum_j^R m_j} \sum_j^R \sum_i^N n_{ji} \quad (4.8)$$

and the variance can be expressed as

$$\sigma^2 = \langle n^2 \rangle - \langle n \rangle^2 = \frac{1}{\sum_j^R m_i} \sum_j^R \sum_i^N n_{ji}^2 - \left(\frac{1}{\sum_j^R m_i} \sum_j^R \sum_i^N n_{ji} \right)^2 \quad (4.9)$$

To simplify the calculations, we can assume that the number of samples in each ψ set is the same, $m_i = m$, in which case $\sum_j^R m_i = Rm$. This is a fair assumption if the rotation between each event is random. Making this replacement, we find sums over the average $\langle n \rangle$ and $\langle n^2 \rangle$ values for the individual ψ sets.

$$\mu = \frac{1}{R} \sum_{j=1}^R \sum_{i=1}^N \frac{n_{ji}}{m} = \frac{1}{R} \sum_{j=1}^R \langle n \rangle_j \quad (4.10)$$

$$\sigma^2 = \langle n^2 \rangle - \langle n \rangle^2 = \frac{1}{R} \sum_{j=1}^R \sum_{i=1}^N \frac{n_{ji}^2}{m} - \left(\frac{1}{R} \sum_{j=1}^R \sum_{i=1}^N \frac{n_{ji}}{m} \right)^2 = \frac{1}{R} \sum_{j=1}^R \langle n^2 \rangle_j - \left(\frac{1}{R} \sum_{j=1}^R \langle n \rangle_j \right)^2 \quad (4.11)$$

Here we have terms of the average multiplicity and multiplicity squared for each ψ set. Since each of these individual sets are binomial, their means and variances should be binomial with each j index mapping to a specific value of ψ :

$$\mu_{\text{binom}}(\psi) = Np(\psi) = \langle n \rangle_j \equiv \mu_j \quad (4.12)$$

$$\sigma_{\text{binom}}^2(\psi) = Np(\psi)(1 - p(\psi)) = \langle n^2 \rangle_j - (\langle n \rangle_j)^2 = \langle n^2 \rangle_j - (\mu_{\text{binom}}(\psi))^2 \equiv \sigma_j^2 \quad (4.13)$$

Plugging Equation 4.12 into 4.12, we find that the mean of the full distribution over all ψ sets is just the average of the individual ψ set means. However, plugging Equation 4.13 into 4.13, we find that the variance of the full distribution is the average of the variances of the individual ψ sets plus the average of their means squared minus the total set mean squared.

$$\mu = \frac{1}{R} \sum_{j=1}^R \sum_{i=1}^N \langle n \rangle_j = \frac{1}{R} \sum_{j=1}^R \mu_j \quad (4.14)$$

$$\sigma^2 = \frac{1}{R} \sum_{j=1}^R \langle n^2 \rangle_j - \left(\frac{1}{R} \sum_{j=1}^R \langle n \rangle_j \right)^2 = \frac{1}{R} \sum_{j=1}^R (\sigma_j^2 + \mu_j^2) - \left(\frac{1}{R} \sum_{j=1}^R \mu_j \right)^2 \quad (4.15)$$

At this point we can reevaluate the discretization of ψ values when constructing the R sets. In the continuum limit, this average over sets becomes an integral over ψ . This integral ranges from $(0, 2\pi)$ and is divided by the unit area, 2π . Taking this continuum limit and plugging the binomial mean and variance into Equations 4.14 and 4.15 gives

$$\mu = \frac{1}{2\pi} \int_0^{2\pi} \mu_{\text{binom}}(\psi) d\psi = \frac{1}{2\pi} \int_0^{2\pi} Np(\psi) d\psi = \frac{N}{2\pi} \int_0^{2\pi} \int_{\psi}^{\psi+w} \tilde{p}(\phi) d\phi d\psi = N \frac{w}{2\pi} = Np \quad (4.16)$$

$$\sigma^2 = \frac{1}{2\pi} \int_0^{2\pi} [\sigma_{\text{binom}}^2(\psi) + \mu_{\text{binom}}^2(\psi)] d\psi - \mu^2 = \frac{1}{2\pi} \int_0^{2\pi} [Np(\psi)(1-p(\psi)) + (Np(\psi))^2] d\psi - \mu^2 \quad (4.17)$$

Substituting $\phi = \phi' + \psi$ in the Equation 4.16 inner integral to switch the order of integration and using the normalization of the probability distribution ($\int_0^{2\pi} \tilde{p}(\phi) d\phi = 1$) we find that the mean of the full distribution is equal to $Nw/2\pi$. This is the probability we would expect for a flat probability distribution in ϕ . While the mean of a flat distribution in ϕ , what we'll denote as $\mu_{\text{binomial}} = Np$ (distinct from $\mu_{\text{binomial}}(\psi)$ which is a function of ψ), produces the same mean as any arbitrary ϕ pdf, the variance of a flat distribution, $\sigma_{\text{binomial}}^2 = Npq = Np(1-p)$, is different from Equation 4.17. In fact, we can rewrite Equation 4.17 in terms of $\sigma_{\text{binomial}}^2$ plus a modification:

$$\begin{aligned}
\sigma^2 &= \frac{N}{2\pi} \left[\int_0^{2\pi} p(\psi) d\psi - \int_0^{2\pi} p(\psi)^2 d\psi + \int_0^{2\pi} Np(\psi)^2 d\psi \right] - \mu^2 \\
&= Np + N(N-1) \int_0^{2\pi} \frac{p(\psi)^2}{2\pi} d\psi - N^2p^2 + (Np^2 - Np^2) \\
&= Np(1-p) + N(N-1) \int_0^{2\pi} \frac{p(\psi)^2}{2\pi} d\psi - N(N-1)p^2 \\
\sigma^2 &= \sigma_{\text{binomial}}^2 + N(N-1) \left[\int_0^{2\pi} \frac{p(\psi)^2}{2\pi} d\psi - p^2 \right]
\end{aligned} \tag{4.18}$$

The variance is then the binomial variance for a flat distribution plus a deviation from this binomial variance. Since $p = w/2\pi$, the only non-trivial term in Equation 4.18 is $\frac{1}{2\pi}N(N-1) \int_0^{2\pi} p(\psi)^2 d\psi$. This is the term we will attempt to isolate and that will carry information about the underlying processes responsible for distributing tracks about the azimuth. Rearranging, we define $\Delta\sigma^2$ as the normalized deviation of the distribution variance from binomial, as defined in Section 4.2.

$$\Delta\sigma^2 \equiv \frac{\sigma^2 - \sigma_{\text{binomial}}^2}{N(N-1)} = \int_0^{2\pi} \frac{p(\psi)^2}{2\pi} d\psi - p^2 \tag{4.19}$$

One additional step can be taken to simplify the right hand side of Equation 4.19. The original PDF, $\tilde{p}(\phi + \psi)$, can be written as a uniform distribution plus a perturbation:

$$\tilde{p}(\phi + \psi) \equiv \frac{1}{2\pi} + \delta\tilde{p}(\phi + \psi) \tag{4.20}$$

Plugging this in to the right hand side of Equation 4.19:

$$\begin{aligned}
\frac{1}{2\pi} \int_0^{2\pi} p(\psi)^2 d\psi - p^2 &= \frac{1}{2\pi} \int_0^{2\pi} \left[\int_0^w \tilde{p}(\phi + \psi) d\phi \right]^2 d\psi - p^2 \\
&= \frac{1}{2\pi} \int_0^{2\pi} \left[\int_0^w \left(\frac{1}{2\pi} + \delta\tilde{p}(\phi + \psi) \right) d\phi \right]^2 d\psi - p^2 \\
&= \frac{1}{2\pi} \int_0^{2\pi} \left[\frac{w}{2\pi} + \int_0^w \delta\tilde{p}(\phi + \psi) d\phi \right]^2 d\psi - p^2 \\
&= \frac{1}{2\pi} \int_0^{2\pi} \left[\left(\frac{w}{2\pi} \right)^2 + \frac{w}{\pi} \int_0^w \delta\tilde{p}(\phi + \psi) d\phi + \left(\int_0^w \delta\tilde{p}(\phi + \psi) d\phi \right)^2 \right] d\psi - p^2 \\
&= \left(\frac{w}{2\pi} \right)^2 + \frac{w}{2\pi^2} \int_0^{2\pi} \int_0^w \delta\tilde{p}(\phi + \psi) d\phi d\psi + \frac{1}{2\pi} \int_0^{2\pi} \left(\int_0^w \delta\tilde{p}(\phi + \psi) d\phi \right)^2 d\psi \\
&\quad - \left(\frac{w}{2\pi} \right)^2 \\
\Delta\sigma^2 &= \frac{1}{2\pi} \int_0^{2\pi} \left(\int_0^w \delta\tilde{p}(\phi + \psi) d\phi \right)^2 d\psi \equiv \frac{1}{2\pi} \int_0^{2\pi} \delta p(\psi)^2 d\psi
\end{aligned} \tag{4.21}$$

The constant term is canceled by p^2 . $\delta\tilde{p}$ is defined to be the deviation of \tilde{p} from its mean value, $1/2\pi$, and therefore has vanishing mean itself. Switching the order to perform the ψ integral first in the cross term, this term vanishes when integrated from 0 to 2π , as the mean is zero.

From Equation 4.21, it is immediately clear that the ψ integral is over a strictly positive quantity, due to the square. This then implies that $\Delta\sigma^2$ itself is a strictly positive quantity under this derivation. In addition, while $\Delta\sigma^2$ may be zero for specific values of w depending on the PDF considered, it is only zero for all w when the PDF itself is uniform. These important observations imply that when correlations are generated via static azimuthal probability distributions, as considered in this derivation, $\Delta\sigma^2$ will always be greater than or equal to zero and $\Delta\sigma^2 = 0$ implies that the tracks are randomly distributed.

The $\Delta\sigma^2$ definition from Equation 4.19 simplifies slightly to Equation 4.22 with the expansion about a uniform distribution:

$$\Delta\sigma^2 \equiv \frac{\sigma^2 - \sigma_{\text{binomial}}^2}{N(N-1)} = \frac{1}{2\pi} \int_0^{2\pi} \delta p(\psi)^2 d\psi \quad (4.22)$$

4.3.2 Two Independent PDFs

The discussion in Section 4.3.1 assumed that tracks were generated via a single azimuthal probability distribution. In reality, it is often convenient to envision the correlation between tracks as originating from some number of independent effects. This is particularly useful when trying to separate the effects of signal from those due to background. We will attempt to modify the arguments of Section 4.3.1 to accommodate a second independent pdf contributing to the final track distribution about the azimuth.

To start, we'll define an "independent" pdf in this case as a static distribution whose relative orientation on the azimuth is random. Under this definition, the two pdfs in question will have no correlation between them. The net effect of these two pdfs is the creation of a new pdf which is their renormalized product:

$$\tilde{p}_{12}(\phi; \psi_1, \psi_2) \equiv C(\psi_1, \psi_2) \tilde{p}_1(\phi + \psi_1) \tilde{p}_2(\phi + \psi_2) \quad (4.23)$$

$$C^{-1}(\psi_1, \psi_2) = \int_0^{2\pi} \tilde{p}_1(\phi + \psi_1) \tilde{p}_2(\phi + \psi_2) d\phi \quad (4.24)$$

Here, $C(\psi_1, \psi_2)$ is the normalization constant for the combined pdf, $\tilde{p}_{12}(\psi_1, \psi_2)$. This combined pdf will function exactly like the single pdf in Section 4.3.1, except now there are two ψ s, one to define the orientation of each original pdf. Expanding the distributions about a uniform distribution will again be useful, and we can therefore pick up at Equation 4.22, replacing $p(\psi)$ with $p(\psi_1, \psi_2)$:

$$\Delta\sigma_{12}^2 = \frac{1}{4\pi^2} \int_0^{2\pi} \int_0^{2\pi} \delta p(\psi_1, \psi_2)^2 d\psi_1 d\psi_2 \quad (4.25)$$

where

$$\begin{aligned}
\delta p_{12}(\psi_1, \psi_2) &\equiv \int_0^w \delta \tilde{p}_{12}(\phi; \psi_1, \psi_2) d\phi \\
&= \int_0^w \left(\tilde{p}_{12}(\phi; \psi_1, \psi_2) - \frac{1}{2\pi} \right) d\phi \\
&= \int_0^w \left(C(\psi_1, \psi_2) \tilde{p}_1(\phi + \psi_1) \tilde{p}_2(\phi + \psi_2) - \frac{1}{2\pi} \right) d\phi
\end{aligned} \tag{4.26}$$

An additional complication in this scenario is the extra factor, $C(\psi_1, \psi_2)$, required to maintain the normalization of the combined pdf. This normalization must be computed separately for each of the possible relative orientations of the two pdfs, and is therefore a function of ψ_1 and ψ_2 . We can simplify $C^{-1}(\psi_1, \psi_2)$ slightly by expanding the two pdfs:

$$\begin{aligned}
C^{-1}(\psi_1, \psi_2) &= \int_0^{2\pi} \left[\frac{1}{2\pi} + \delta \tilde{p}_1(\phi + \psi_1) \right] \left[\frac{1}{2\pi} + \delta \tilde{p}_2(\phi + \psi_2) \right] d\phi \\
&= \int_0^{2\pi} \left[\frac{1}{4\pi^2} + \frac{1}{2\pi} \delta \tilde{p}_1(\phi + \psi_1) + \frac{1}{2\pi} \delta \tilde{p}_2(\phi + \psi_2) + \delta \tilde{p}_1(\phi + \psi_1) \delta \tilde{p}_2(\phi + \psi_2) \right] d\phi \\
&= \frac{1}{2\pi} + \int_0^{2\pi} \delta \tilde{p}_1(\phi + \psi_1) \delta \tilde{p}_2(\phi + \psi_2) d\phi
\end{aligned} \tag{4.27}$$

The linear $\delta \tilde{p}_1$ and $\delta \tilde{p}_2$ vanish because \tilde{p}_1 and \tilde{p}_2 are themselves properly normalized pdfs, so the deviation from their mean vanishes when integrated over the full azimuth. This is a helpful simplification, but Equation 4.27 refers to the inverse of the normalization constant and the resulting term is in the denominator. We will now assume that the second term is small compared to the constant $1/2\pi$. This is reasonable, as the randomization of the ψ s will tend to produce destructive interference between the two deviation terms and integrating over all possible combinations of the ψ s will ensure this. A worst case scenario would occur when the pdfs are similar and the ψ s are aligned. Taking the example of two identical elliptic flow distributions, the term would yield v_2^2/π . The magnitude of this term with respect to $1/2\pi$ would still be small for reasonable $v_2 \lesssim 0.1$.

Under the assumption that the second term is small, $C(\psi_1, \psi_2)$ can be written as an expansion:

$$\begin{aligned}
C(\psi_1, \psi_2) = & 2\pi - 4\pi^2 \int_0^{2\pi} \delta\tilde{p}_1(\phi + \psi_1) \delta\tilde{p}_2(\phi + \psi_2) d\phi \\
& + 8\pi^3 \left(\int_0^{2\pi} \delta\tilde{p}_1(\phi + \psi_1) \delta\tilde{p}_2(\phi + \psi_2) d\phi \right)^2 + \mathcal{O}(\cdot^3)
\end{aligned} \tag{4.28}$$

We can now write Equation 4.25 in terms of the individual probability distributions expanded about uniform distributions and suppress the $(\phi + \psi)$ arguments of the $\delta\tilde{p}$.

$$\begin{aligned}
\Delta\sigma_{12}^2 = & \frac{1}{4\pi^2} \int_0^{2\pi} \int_0^{2\pi} \left[\int_0^w \left(\frac{[\frac{1}{2\pi} + \delta\tilde{p}_1(\phi + \psi_1)] [\frac{1}{2\pi} + \delta\tilde{p}_2(\phi + \psi_2)]}{\frac{1}{2\pi} + \int_0^{2\pi} \delta\tilde{p}_1(\phi + \psi_1) \delta\tilde{p}_2(\phi + \psi_2) d\phi} - \frac{1}{2\pi} \right) d\phi \right]^2 d\psi_1 d\psi_2 \\
= & \frac{1}{4\pi^2} \int_0^{2\pi} \int_0^{2\pi} \left[\int_0^w \frac{\frac{1}{4\pi^2} + \frac{1}{2\pi} \delta\tilde{p}_1 + \frac{1}{2\pi} \delta\tilde{p}_2 + \delta\tilde{p}_1 \delta\tilde{p}_2}{\frac{1}{2\pi} + \int_0^{2\pi} \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi} d\phi - \int_0^w \frac{1}{2\pi} d\phi \right]^2 d\psi_1 d\psi_2
\end{aligned} \tag{4.29}$$

We will now use the expansion of $C(\psi_1, \psi_2)$ from Equation 4.28 and keep only the leading 2π term.

$$\begin{aligned}
\Delta\sigma_{12}^2 &\approx \frac{1}{4\pi^2} \int_0^{2\pi} \int_0^{2\pi} \left[2\pi \int_0^w \left(\frac{1}{4\pi^2} + \frac{1}{2\pi} \delta\tilde{p}_1 + \frac{1}{2\pi} \delta\tilde{p}_2 + \delta\tilde{p}_1 \delta\tilde{p}_2 \right) d\phi - \frac{w}{2\pi} \right]^2 d\psi_1 d\psi_2 \\
&= \frac{1}{4\pi^2} \int_0^{2\pi} \int_0^{2\pi} \left[\frac{w}{2\pi} + \int_0^w \delta\tilde{p}_1 d\phi + \int_0^w \delta\tilde{p}_2 d\phi + 2\pi \int_0^w \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi - \frac{w}{2\pi} \right]^2 d\psi_1 d\psi_2 \\
&= \frac{1}{4\pi^2} \int_0^{2\pi} \int_0^{2\pi} \left[\left(\int_0^w \delta\tilde{p}_1 d\phi \right)^2 + \left(\int_0^w \delta\tilde{p}_2 d\phi \right)^2 + 2 \int_0^w \delta\tilde{p}_1 d\phi \int_0^w \delta\tilde{p}_2 d\phi \right] d\psi_1 d\psi_2 \\
&\quad + \frac{1}{\pi} \int_0^{2\pi} \int_0^{2\pi} \left(\int_0^w \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi \right) \left(\int_0^w \delta\tilde{p}_1 d\phi + \int_0^w \delta\tilde{p}_2 d\phi + \pi \int_0^w \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi \right) d\psi_1 d\psi_2 \\
&= \frac{1}{2\pi} \int_0^{2\pi} \left(\int_0^w \delta\tilde{p}_1 d\phi \right)^2 d\psi_1 + \frac{1}{2\pi} \int_0^{2\pi} \left(\int_0^w \delta\tilde{p}_2 d\phi \right)^2 d\psi_2 \\
&\quad + \frac{1}{2\pi^2} \left(\int_0^{2\pi} \int_0^w \delta\tilde{p}_1 d\phi d\psi_1 \right) \left(\int_0^{2\pi} \int_0^w \delta\tilde{p}_2 d\phi d\psi_2 \right) \\
&\quad + \frac{1}{\pi} \int_0^{2\pi} \int_0^{2\pi} \left(\int_0^w \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi \right) \left(\int_0^w \delta\tilde{p}_1 d\phi + \int_0^w \delta\tilde{p}_2 d\phi + \pi \int_0^w \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi \right) d\psi_1 d\psi_2 \\
&= \Delta\sigma_1^2 + \Delta\sigma_2^2 \\
&\quad + \frac{1}{\pi} \int_0^{2\pi} \int_0^{2\pi} \left(\int_0^w \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi \right) \left(\int_0^w \delta\tilde{p}_1 d\phi + \int_0^w \delta\tilde{p}_2 d\phi + \pi \int_0^w \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi \right) d\psi_1 d\psi_2
\end{aligned} \tag{4.30}$$

The cross term between $\delta\tilde{p}_1$ and $\delta\tilde{p}_2$ vanishes as the ψ integrals are performed first over the full azimuth. The leading quadratic terms in $\delta\tilde{p}_1$ and $\delta\tilde{p}_2$ are compared with the last line of Equation 4.21 and found to be the $\Delta\sigma^2$ contributions from the individual pdfs.

Equation 4.30 tells us that, to leading order, the effects of independent sources of correlation superimpose in the observable:

$$\Delta\sigma_{12}^2 \approx \Delta\sigma_1^2 + \Delta\sigma_2^2 \tag{4.31}$$

This allows us to confidently subtract background. We should be careful, however, as we have shown that the effects only superimpose at leading order. We will order the full expansion in powers of $\delta\tilde{p}_i$. The $\Delta\sigma_i^2$ terms themselves are second order in this expansion. The remaining terms in Equation 4.30 are of order three and four. We should be careful

here, as we took only the leading term in the expansion of the normalization, $C(\psi_1, \psi_2)$. We can look at the sub-leading term in $C(\psi_1, \psi_2)$ and retain only its order three terms, noting that they are negative:

$$-\frac{1}{4\pi^2} \int_0^{2\pi} \int_0^{2\pi} \left(\int_0^{2\pi} \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi \right) \left[2 \int_0^w \delta\tilde{p}_1 d\phi + 2 \int_0^w \delta\tilde{p}_2 d\phi \right] d\psi_1 d\psi_2 \quad (4.32)$$

We can now collect all terms up to order three in Equation 4.30 and 4.32 to write the expansion of $\Delta\sigma_{12}^2$ in $\delta\tilde{p}_i$ to third order:

$$\begin{aligned} \Delta\sigma_{12}^2 = & \Delta\sigma_1^2 + \Delta\sigma_2^2 \\ & + \int_0^{2\pi} \int_0^{2\pi} \left[\int_0^w \delta\tilde{p}_1 d\phi + \int_0^w \delta\tilde{p}_2 d\phi \right] \left[2\pi \int_0^w \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi - \int_0^{2\pi} \delta\tilde{p}_1 \delta\tilde{p}_2 d\phi \right] \frac{d\psi_1 d\psi_2}{2\pi^2} + \mathcal{O}(\delta\tilde{p}_i^4) \end{aligned} \quad (4.33)$$

Subtracting Background Reliably removing sources of background correlation will be important when measuring $\Delta\sigma^2$ in data or model. From Equation 4.31 we find that, at leading order, $\Delta\sigma^2$ of the combination of the two sources is simply the sum of the individual source $\Delta\sigma^2$ values. If we can estimate the value of the background $\Delta\sigma^2$ independently, we can simply subtract this from the combined $\Delta\sigma^2$. This will be tested numerically in the following Section 4.4.2 and more thoroughly with event by event simulations in Section 4.4.3.

We have noted that Equation 4.31 is only true to leading order and we will want to estimate the uncertainty associated with subtraction of background sources. Equation 4.33 shows that the expression to third order in $\delta\tilde{p}$ contains four potentially non-zero terms. All four of these terms are $\mathcal{O}(\delta\tilde{p}_i^3)$ by definition while the leading terms are $\mathcal{O}(\delta\tilde{p}_i^2)$. We will assume that we have measured the full $\Delta\sigma_{12}^2 \approx \Delta\sigma_1^2$, we know $\Delta\sigma_2^2$ and are trying to estimate the third order contribution $\delta\Delta\sigma_{12}^2$. A decent estimate would be the source term times the background term to the 3/4 power. We don't measure $\Delta\sigma_1^2$ directly, but we can use our $\Delta\sigma_{12}^2$ measurement minus the background $\Delta\sigma_2^2$ as a good approximation (this is the

background corrected $\Delta\sigma^2$ that we are trying to estimate in the first place). The third order uncertainty estimate on Equation 4.31 is given by Equation 4.34.

$$\delta\Delta\sigma_{12}^2 \approx (\Delta\sigma_1^2\Delta\sigma_2^2)^{3/4} \approx [(\Delta\sigma_{12}^2 - \Delta\sigma_2^2)\Delta\sigma_2^2]^{3/4} \quad (4.34)$$

So far in this analysis we have worked with a fixed a partition width w and $\Delta\sigma^2$ will depend upon the partition width used. $\Delta\sigma_2^2$ from a particular source may be zero for a particular value of w , in which case $\Delta\sigma_{12}^2 = \Delta\sigma_1^2$ at leading order. The higher order terms, however, won't necessarily vanish and this is not reflected in the uncertainty estimate of Equation 4.34. We therefore modify the uncertainty estimate such that the $\Delta\sigma_2^2$ contribution is averaged over all w while still using the corrected value for the specific w .

$$\delta\Delta\sigma_{12}^2(w) \approx [(\Delta\sigma_{12}^2(w) - \Delta\sigma_2^2(w))\langle\Delta\sigma_2^2\rangle_w]^{3/4} \quad (4.35)$$

This uncertainty estimate will be tested numerically in Sections 4.4.2 and 4.4.3 and is found to be appropriate in the cases considered.

4.4 Elliptic Flow Correction

Elliptic flow in heavy ion collisions is strong due to the geometry of the colliding systems and the hydrodynamic evolution of the QGP medium. The overlapping area in which two pancake shaped colliding nuclei interact produces an almond shape droplet of QGP. This almond shape produces a larger density gradient along the short axis than the long axis perpendicular to the reaction plane. This produces larger hydrodynamic flow along the short axis and the final state particles retain this asymmetric momentum distribution after freeze-out. This asymmetry is parameterized in the second harmonic of a Fourier expansion about the reaction plane of a collision, v_2 .

Elliptic flow will produce a signal that appears to $\Delta\sigma^2$ like an attractive interaction

between proton tracks. In order to eliminate this expected source of correlation from the $\Delta\sigma^2$ observable, we must measure v_2 in our data (Section 5.4) and subtract its contribution (Equation 4.39) from the measured $\Delta\sigma^2$ values. In Section 4.4.1 we will derive an analytical equation for $\Delta\sigma^2$ as a function of the azimuthal partition width for a given v_2 magnitude. We will combine flow with a simple clustering signal in Section 4.4.2 and numerically integrate to demonstrate the efficacy of flow subtraction. Finally, in Section 4.4.3 we will utilize an event-by-event Monte Carlo simulation of flow combined with other signals to show that the subtraction works in more realistic scenarios.

4.4.1 Exact Integration

Some models for distributing tracks about the azimuth can be expressed as a simple pdf as used in the derivation from Section 4.3. One important effect that can be written this way is hydrodynamic flow. Flow is modeled as a Fourier expansion of the azimuthal probability distribution of emitted particles [48]:

$$\frac{dN}{d\phi} = \frac{1}{2\pi} \left[1 + \sum_{n=1}^{\infty} 2v_n \cos[n(\phi - \psi_R)] \right] \quad (4.36)$$

This is a pdf about the azimuth, exactly what is treated in Section 4.3. We can therefore compute the integral on the right side of Equation 4.19 or 4.22. While Equation 4.36 is linear in the v_n harmonics, the square on $p(\psi)$ in Equation 4.19 will mix terms. This means that if more than one harmonic has an appreciable magnitude, the fully mixed integral will have to be computed directly. In the case of only one harmonic with appreciable magnitude, we are left with:

$$\frac{dN}{d\phi} = \frac{1}{2\pi} [1 + 2v_n \cos[n(\phi - \psi_R)]] \quad (4.37)$$

Replacing $p(\psi)$ with Equation 4.37 in Equation 4.19, we can solve for $\Delta\sigma^2$ at a given harmonic n :

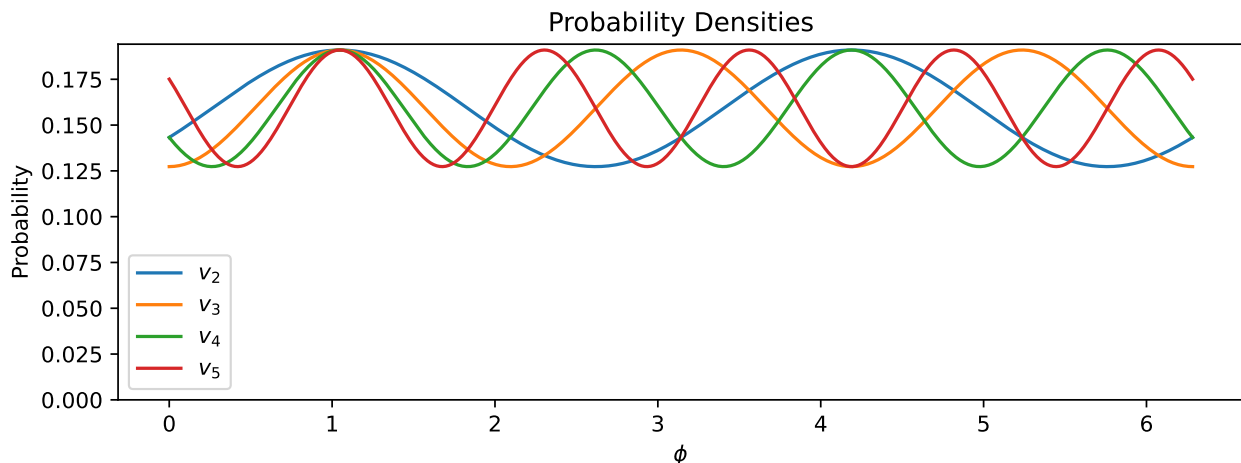


Figure 4.3: Probability distributions for v_n flow distributions in which only one harmonic (v_n) has a non-zero magnitude.

$$\Delta\sigma_n^2 = \frac{2v_n^2}{\pi^2 n^2} \sin\left(\frac{nw}{2}\right)^2 \quad (4.38)$$

The probability densities corresponding to Equation 4.37 for $n \in [2, 3, 4, 5]$, $v_n = 0.1$, and $\psi_R = \pi/3$ are shown in Figure 4.3. The corresponding $\Delta\sigma^2$ integrals are plotted in Figure 4.4 as a function of the azimuthal partition width, w . The presence of the n^2 term in the denominator of Equation 4.38 implies that the magnitude of $\Delta\sigma^2$ decreases quadratically with the order of the harmonic, which is reflected in Figure 4.4.

Direct flow, v_1 will be exceedingly small due to momentum conservation as our acceptance is longitudinally symmetric about the collision vertex. Equation 4.38 has a factor of n^2 in the denominator, which quickly reduces the contribution from higher order flow components, regardless of their v_n magnitudes. Elliptic flow is large in heavy ion collisions due to the geometric asymmetry of the overlap region between two ions. We will therefore focus on correcting for elliptic flow using Equation 4.39 and ignore the contributions from higher orders. In the future, it may be worth investigating the $\Delta\sigma^2$ integral (Equation 4.19) for the full Fourier series in Equation 4.36, though we expect the contribution from higher harmonics to be small due to their small observed magnitudes relative to the second harmonic.




Figure 4.4: Analytical integral of $\delta p(\psi)$ in Equation 4.22 where $\tilde{p}(\phi - \psi)$ is of form Equation 4.37, corresponding to the pdfs plotted in Figure 4.3.

$$\Delta\sigma^2(w; v_2) = \frac{v_2^2}{2\pi^2} \sin(w)^2 \quad (4.39)$$

4.4.2 Numerical Integration

When it is not possible to carry out the $\Delta\sigma^2$ integral of Equation 4.22 analytically, the integration can always be approached numerically. Even a simple Gaussian pdf produces an intractable analytic integration, so numerical methods are necessary in many cases. For well defined analytic pdfs such as flow (Equation 4.38), Gaussian distributions, or a combination of these, the `nquad` and `quad` functions from Python’s `scipy` package are used. This allows for the calculation of $\int_0^w \delta\tilde{p}(\phi + \psi) d\phi$ in Equation 4.22 given an azimuthal partition width w and pdf orientation ψ . The azimuth is periodic and the azimuthal partition may extend past this periodic boundary at 0 or 2π . We therefore wrap pdfs around this boundary, as explained in Section 5.2. A comparison between the analytic and numerical $\Delta\sigma^2$ for elliptic flow with $v_2 = 0.1$ is shown in Figure 4.5.

Numerical integration of azimuthal pdfs allows us to investigate the $\Delta\sigma^2$ directly, without generating any event-like data. If we can model a correlation as a static probability distri-




Figure 4.5: Comparison of the analytic vs numerical computation of the $\Delta\sigma^2$ integral for elliptic flow with $v_2 = 0.1$. The bottom panel shows that the difference between numerical and analytic is exceedingly small.

bution about the azimuth, we can then directly investigate the expected effect on $\Delta\sigma^2$. This allows us to test an important scenario in which two different types of correlation influence the final state proton azimuthal distributions. Such is the case when both signal and background correlations are present. In this case, we hope to somehow remove the background such that the final analysis can be implemented on signal alone. Numerical integration will allow us to construct these signal/background combinations and test the efficacy of correction techniques on the $\Delta\sigma^2$ observable without the complication of statistical fluctuations present in event-by-event simulations.

We will test background subtraction by generating an admixture of a simple Gaussian clustering signal with an elliptic flow background. The effect of elliptic flow on $\Delta\sigma^2$ is demonstrated in Figure 4.5. An important point is made in Section 4.4.1: if two different types of correlation are themselves correlated event-by-event, the full azimuthal pdf of their combination must be integrated in Equation 4.19. The most pertinent example is different




Figure 4.6: Azimuthal probability distributions for elliptic flow with $v_2 = 0.07$ and $\psi_{v_2} = \pi/3$ and the simple Gaussian clustering model with $A = 0.165$, $\sigma = 0.8$, and $\psi_{\text{gaus}} = \pi$. Also shown is the normalized product of the elliptic flow and Gaussian cluster pdfs in this particular ψ orientation.

flow harmonics, which are all oriented relative to the same reaction plane in each event. Even if these correlations are themselves linear, as in Equation 4.36, the $p(\psi)^2$ term will produce cross terms such that the contribution to $\Delta\sigma^2$ is not a simple linear combination of the contributions from the individual pdfs.

When two effects are not correlated in an event, however, the situation becomes workable as explained in Section 4.3.2. We can calculate $\Delta\sigma_{12}^2$ with Equation 4.25, being careful to keep $C(\psi_1, \psi_2)$ properly normalized for each relative orientation of ψ_1 and ψ_2 . In our demonstration, we will use a simple Gaussian clustering signal (Appendix A.5) whose orientation is independent from the reaction plane orientation of an elliptic flow background. Figure 4.6 shows the two pdfs along with their normalized product in the particular configuration of $\psi_{v_2} = \pi/3$ and $\psi_{\text{gaus}} = \pi$.

$\Delta\sigma^2$ is calculated for each of these three pdfs and plotted as a function of the azimuthal partition width w in Figure 4.7. The elliptic flow $\Delta\sigma^2$ is then subtracted from the combi-




Figure 4.7: $\Delta\sigma^2$ calculated via numeric integration for a simple Gaussian clustering signal with $A = 0.165$ and $\sigma = 0.8$, an elliptic flow background with $v_2 = 0.07$, and for their combination. The elliptic flow $\Delta\sigma^2$ is subtracted from the combination and this difference is found to adequately match the pure Gaussian clustering signal. The Gaussian clustering signal minus the correct value is plotted on the bottom panel along with a band corresponding to the estimated NLO contribution.

nation $\Delta\sigma^2$ and plotted as a red dashed line. After subtraction, the “corrected” $\Delta\sigma^2$ value matches the original Gaussian $\Delta\sigma^2$ value very well, indicating that the background subtraction works as expected. The bottom panel of Figure 4.7 shows the difference between the corrected value and the pure Gaussian cluster signal. The red band represents the estimated NLO term calculated via Equation 4.35. This numerical integration demonstrates that subtracting an independent background effect produces a $\Delta\sigma^2$ very close to the true signal value and that the uncertainty due to higher order terms is adequately estimated via Equation 4.35.

4.4.3 Event by Event Simulation

The flow background subtraction can be more directly tested via Monte Carlo event simulations. For these tests, we will produce simulations of events with both signal and flow background induced correlations. The signal correlations will be modeled via the simple Gaussian clustering model described in Section A.5 and the Gaussian correlation model described in Section 5.2. Flow is simulated as described in Appendix A.4. The produced events are run through the same analysis as STAR data, described in Section 5.1.

We begin by combining a simple clustering model ($A = 0.165$, $\sigma = 0.8$) with an elliptic flow simulation ($v_2 = 0.07$). We then use Equation 4.39 with $v_2 = 0.07$ to calculate the background due to the elliptic flow. This background is subtracted from the combined $\langle\Delta\sigma^2\rangle$ and compared with the pure signal in Figure 4.8. We find that the flow correction works well and the deviations in the bottom panel are covered by the estimated NLO uncertainties.

We next test the flow correction using the slightly more involved Gaussian correlation model, described in Section 5.2. This model can produce a negative $\Delta\sigma^2$ values, which correspond to a repulsive interaction between proton tracks. A repulsive correlation is particularly important to test because this is what we observe in both STAR and model data.

The tests are conducted in a similar manner to the previous one. A pure repulsive signal data set is simulated along with a combination data set of both repulsive signal and flow background. The flow correction is made on the combination data set and the result is compared to the pure signal data set.

Figure 4.9 shows the results for this test with Gaussian correlation model parameters of ($A = -0.01$, $\sigma = 0.1$). In this test the correction works pretty well, though the corrected $\langle\Delta\sigma^2\rangle$ values are slightly lower than the pure signal values, indicating an over-correction. When σ is changed from 0.1 to 1.0 (leaving $A = -0.01$), the repulsive signal becomes stronger and when the flow correction is made in Figure 4.10 there is an even larger over-correction. In both cases, the over-correction is larger than the estimated NLO contribution. This is




Figure 4.8: A simulation is run with pure clustering signal and another is run with clustering as well as elliptic flow. The contribution from the elliptic flow is calculated from Equation 4.39 with $v_2 = 0.07$ and subtracted from the combined $\langle \Delta\sigma^2 \rangle$. After correction, the combined simulation matches the pure signal simulation, demonstrating that the flow correction is effective. The bottom plot shows the signal minus the corrected values. The solid error bars represent the statistical error propagated through the subtraction while the transparent error bars represent the estimated error from NLO contributions from the subtraction.




Figure 4.9: The Gaussian correlation model with $(A = -0.01, \sigma = 0.1)$ is run with an elliptic flow of $v_2 = 0.07$ to test the flow correction. A slight over-correction is observed as a consequence of the Gaussian correlation model coupling to the flow simulation.

understood to be a consequence of the coupling of the Gaussian correlation model to the simulated flow background, as explained in Section 5.2. In both of these tests, however, the discrepancy between the pure signal and the corrected data set is small and understood as due to the coupling, and we therefore remain confident in the flow correction for removing v_2 background.

4.4.4 Elliptic Flow Closure Test Conclusion

We find that the elliptic flow subtraction works very well in the case of simple clustering, where the slight variation of the corrected value compared to the true value is explained by the NLO uncertainty. When flow is simulated alongside the Gaussian correlation model a slight over-correction beyond what can be explained by the NLO contribution is observed. This is due to the two models coupling which is expected to produce slight non-linearity in the combined effect. Overall the flow correction seems to work as expected within the




Figure 4.10: The Gaussian correlation model with ($A = -0.01, \sigma = 1$) is run with an elliptic flow of $v_2 = 0.07$ to test the flow correction. A slight over-correction is observed as a consequence of the Gaussian correlation model coupling to the flow simulation.

uncertainties estimated from NLO terms.

4.5 Mixed Event Correction

Mixed events are utilized in STAR analysis to correct for detector effects. A mixed event with N proton tracks is generated by sampling a single proton track from N single events. Mixed events will capture detector effects that are present in each event while washing out any intra-event signal such as correlation between tracks. In particular, azimuthal dependent inefficiencies will be captured by the mixed events and this contribution to the $\Delta\sigma^2$ signal can be subtracted.

Single events are sorted into classes based on the longitudinal position of their primary vertices and their centrality. We use 5 equal sized bins within the full longitudinal vertex range and nominal centrality class binning (0-5%, 5-10%, 10-20%, 20-30%, ...). As a data

set is being analyzed, each single event is sorted into one of these mixed classes. When a minimum of 150 single events are accumulated in a mixing class, mixed events are generated from this pool of events each time a new single event is added. We typically generate 5-10 mixed events for each single event, with the same number of total protons (N) as the single event being appended. Once 250 events are accumulated in a mixing class, an event is randomly removed when a new single event is added to the class. This maintains a maximum of 250 single events per mixing class which keeps the computer memory from overloading.

$\Delta\sigma^2$ is calculated for these mixed events in an identical manner to the single events. To correct the single event data, the mixed event value is subtracted from the single event value: $\Delta\sigma_{\text{Mix Corrected}}^2 = \Delta\sigma_{\text{Single}}^2 - \Delta\sigma_{\text{Mix}}^2$. We will test this correction's ability to account for ϕ dependent detector inefficiencies. This will be done first with numerical integration and then with event-by-event simulations.

4.5.1 Numerical Integration

In Section 4.4.2 we demonstrate that for correlations which can be modeled as static probability distributions about the azimuth, $\Delta\sigma^2$ can be calculated directly from the pdf. If two separate correlations of this type have independent orientations in each event, we expect that their combined $\Delta\sigma^2$ will be the sum of the individual $\Delta\sigma^2$ s. We will test this for the case of an azimuthal dependent detector inefficiency with a simple Gaussian clustering signal described in Section A.5.

The azimuthal efficiency pdf we will use is generated from STAR BES-I data at 62 GeV, as described in Appendix A.3. A histogram of the azimuthal angles of all protons used in the analysis is constructed and normalized, as shown in Figure 4.11 along with the Gaussian pdf. In the efficiency pdf, the twelve sectors of STAR's TPC are clearly visible, with a dramatic decrease in the number of protons detected with an azimuthal angle lying between sectors. In addition, each sector has its own characteristic efficiency as some readout pads may not be functioning. These ϕ dependent efficiencies will look like local clusters to the




Figure 4.11: Azimuthal probability distributions for a simple Gaussian clustering signal and a ϕ dependent efficiency background. The efficiency pdf is derived from the distribution of proton tracks in 62 GeV STAR data. The simple Gaussian clustering model has an amplitude of $A = 0.165$ and $\sigma = 0.8$. Also included is the normalized product of these pdfs in this specific relative orientation.

$\Delta\sigma^2$ observable, so must be accounted for.

The efficiency pdf is not an analytic function, but a series of values for each angle ϕ . This type of non-analytic pdf is not integrated well by Python's `scipy` package and a manual integration is employed instead. The algorithm calculates the azimuthal pdf at some number of equally spaced points, typically 1000, within the $[0, 2\pi)$ bounds. By a change of variables the ψ shift can also be written as the lower bound of the integral $\int_{\psi}^{\psi+w} \delta\tilde{p}(\phi) d\phi$. The azimuth is periodic and the azimuthal partition may extend past this periodic boundary at 2π . We deal with this numerically by duplicating the original pdf and attaching it in range $[2\pi, 4\pi)$. The azimuthal partition width w is at most 2π , so a total range of $[0, 4\pi)$ is sufficient. The periodicity about the azimuth must be respected by the pdfs and is enforced for Gaussians by wrapping, as described in Section 5.2. To compute the the $\frac{1}{2\pi} \int_0^{2\pi} \delta p(\psi)^2 d\psi$ term in Equation 4.22, ψ values are generated between $[0, 2\pi)$ and, for each ψ , the $\delta p(\psi)$ integral is




Figure 4.12: $\Delta\sigma^2$ calculated via numeric integration for the simple Gaussian clustering signal and STAR efficiency separately as well as for their combination. The efficiency $\Delta\sigma^2$ is subtracted from the combination and this difference is found to adequately match the pure Gaussian clustering signal. The difference between the pure signal and the corrected value is shown in the bottom panel along with a shaded band representing the estimated uncertainty due to NLO contributions.

calculated. These resulting integrals are squared and then summed to produce the desired integral for $\Delta\sigma^2$.

Figure 4.12 shows the numerically integrated $\Delta\sigma^2$ observable for the efficiency and Gaussian pdfs separately as well as their combination. The efficiency $\Delta\sigma^2$ values (representing background) are subtracted from the combination and the result matches the Gaussian clustering signal rather well. The bottom panel shows the difference between the corrected $\Delta\sigma^2$ and the pure Gaussian clustering signal along with an error band representing the estimated contribution of NLO terms (Section 4.3.2). The deviation of the corrected value from the true value is rather small and covered by the uncertainty estimated for the NLO contribution.

4.5.2 Event by Event Simulation

We can test the mixed correction more directly via Monte Carlo event simulations, analogous to Section 4.4.3. The signal correlations will be modeled via the Gaussian correlation model described in Section 5.2 and the simple Gaussian clustering model described in Section A.5. The ϕ dependent efficiency is modeled with the STAR data described in the previous section (Figure 4.11).

We first simulate STAR efficiency along with simple clustering with $A = 0.165$ and $\sigma = 0.8$. When distributing N tracks in an event, each track will first be sampled from a Gaussian pdf with probability p or from a uniform distribution with probability $1 - p$ (here $p = 0.05$, see Appendix A.5 for converting between A and p). The resulting azimuthal angle is then rejection sampled according to the efficiency distribution. This process is repeated until the desired N tracks are produced (Section A.3). The number of tracks per event is sampled from a uniform distribution between 2 and 80. These events are azimuthally partitioned, the variance of the azimuthal multiplicity distributions are calculated for each N , then $\Delta\sigma^2$ is calculated from these variances and averaged over N .

Figure 4.13 shows $\langle\Delta\sigma^2\rangle$ for single events vs the azimuthal partition width for a pure simple clustering simulation (Signal) and a simulation of this same simple clustering along with STAR ϕ dependent efficiency (Combination). Also shown is the clustering and efficiency simulation in which the mixed event $\Delta\sigma^2$ (Background) is subtracted (Corrected). These mixed subtracted points overlap those from the pure clustering signal, indicating that the efficiency effect was successfully corrected. The bottom panel shows the difference between the signal and the corrected points. We observe in this scenario that the NLO (Section 4.3.2) error estimation is overestimated.

We next test the mixed correction using the Gaussian correlation model (Section 5.2). Figure 4.14 shows the results for this test with Gaussian correlation model parameters of ($A = -0.01, \sigma = 0.1$) while Figure 4.15 shows ($A = -0.01, \sigma = 1.0$). As in the flow




Figure 4.13: A simulation is run with pure clustering signal and another is run with clustering as well as STAR ϕ dependent efficiency. In the data set with both signal and background, the mixed correction is implemented $\langle \Delta\sigma_{\text{Single}}^2 - \Delta\sigma_{\text{Mixed}}^2 \rangle$ and this corrected $\langle \Delta\sigma^2 \rangle$ is plotted along with the other two data sets. After correction, the combined simulation matches the pure signal simulation, demonstrating that the mixed correction is effective. The bottom plot shows the signal minus the corrected values. The solid error bars represent the statistical error propagated through the subtraction while the transparent error bars represent the estimated error from NLO contributions from the subtraction.




Figure 4.14: The Gaussian correlation model with $(A = -0.01, \sigma = 0.1)$ is run with the 62 GeV STAR ϕ dependent efficiency to test the mixed event correction. A slight over-correction is observed as a consequence of the Gaussian correlation model coupling to the efficiency simulation.

case, there is over-subtraction in both cases. We again attribute this to the coupling of the Gaussian correlation model to the efficiency simulation (Section 5.2). We note that the over-correction is small in both cases and, even though the NLO uncertainty estimates clearly don't cover the true value, they aren't too far off. As we understand the source of this small over-correction, we again conclude that the mixed correction is adequate.

4.5.3 Efficiency Correction Closure Test Conclusion

While a slight over-correction is observed when combining a ϕ dependent efficiency simulation with a repulsive Gaussian correlation model, this is expected to be a consequence of the repulsive model coupling to the ϕ dependent efficiency simulation. The correction is very




Figure 4.15: The Gaussian correlation model with ($A = -0.01, \sigma = 1.0$) is run with the 62 GeV STAR ϕ dependent efficiency to test the mixed event correction. A noticeable over-correction is observed as a consequence of the Gaussian correlation model coupling to the efficiency simulation.

close to the pure signal for the repulsive simulation and exactly matches the signal in the simple Gaussian clustering and flow test which allows us to confidently conclude that the correction by mixed events is effective in removing ϕ dependent efficiency contribution to $\Delta\sigma^2$.

4.6 Combined Efficiency and Flow Correction

We can finally consider contribution from two backgrounds: both elliptic flow and a phi dependent efficiency. Figure 4.16 shows the combination of elliptic flow where $v_2 = 0.07$, STAR 62 GeV phi dependent efficiency, and the simple clustering model with $A = 0.165$ and $\sigma = 0.8$. The efficiency correction is first applied to the combined simulation, followed by the flow correction. We find that in this case the corrected data still agrees with the pure clustering simulation and that the deviations are roughly covered by the NLO uncertainty estimation. We therefore conclude that the efficiency and flow corrections still work when both are applied sequentially.




Figure 4.16: The simple clustering model with ($A = 0.165, \sigma = 0.8$) (Signal) is run with the 62 GeV STAR ϕ dependent efficiency (Background #1) as well as elliptic flow with $v_2 = 0.07$ (Background #2). The efficiency correction is performed on the combined simulation followed by the flow correction. The corrected $\langle \Delta \sigma^2 \rangle$ is compared to the pure simple clustering signal and found to agree well, with the variations shown in the bottom panel rather well explained by the NLO uncertainty estimates.

CHAPTER 5

Proton Fluctuations in Azimuthal Partitions at STAR

Chapters 3 and 4 established the azimuthal partition methodology and the $\Delta\sigma^2$ observable. This chapter will focus on measuring $\Delta\sigma^2$ in STAR BES-I data and comparing to model and simulation predictions. In the presence of a first-order phase transition, we may expect bubbles or clusters of high baryon density in the transition from QGP to hadron gas [23]. These clumps in coordinate space should be translated into momentum space via the radial flow of the system. The clumps of baryon density would hadronize into final state baryons which could then be detected. While the size of this signal may be significantly reduced when translated from coordinate to momentum space, it is possible that a detectable signal may survive.

The $\Delta\sigma^2$ observable is used to measure the azimuthal correlation between proton tracks as a function of energy in the STAR BES-I data. We look to form a baseline of correlation at high energy where the transition is expected to be a cross-over. As energy is reduced, signs of increasing attractive correlation may be indicative of cluster formation associated with the onset of a first-order transition. The measurements are compared to models which do not contain this first-order behavior.

5.1 Analysis of $\Delta\sigma^2$ in STAR BES-I

Proton azimuthal partition multiplicity distributions in the STAR BES-I data set are analyzed for signs of clustering which may be indicative of a first-order phase transition. In each




Figure 5.1: Example of an azimuthal partition proton multiplicity distribution in the most central AMPT data at 39 GeV. 72 partitions of width $w = 120^\circ$ are sampled randomly in each event and the number of proton tracks within each are counted. A histogram is constructed over all events with $N = 20$ total protons on the full azimuth. AMPT data is plotted as blue bars while the expected binomial distribution of N trials and $p = w/2\pi = 1/3$ probability of success is plotted in red dots.

event, 72 azimuthal partitions of width w are randomly sampled and the number of identified proton tracks with $y < |0.5|$ and $0.4 < p_T < 2.0$ GeV within each partition are counted. An azimuthal partition multiplicity distribution is constructed separately over many events for each energy, centrality and azimuthal partition width. In addition, the total number of protons on the full azimuth, N , are counted and distributions are constructed separately for each N . An example of one of these distributions is shown in Figure 5.1 for AMPT data (Appendix B.1). The distribution is very similar to the overlaid binomial distribution for N trials and probability of success $p = w/2\pi$.

The variance of this distribution and similar distributions for all N are computed and plotted on the left side of Figure 5.2 as a function of N , along with the variances of the corresponding mixed event data (Section 4.5) and the expected binomial variance — $\sigma_{\text{binomial}}^2 = Np(1-p)$. It is clear from this plot that the widths of the distributions are very similar to binomial for all N , though the single event variances seem slightly lower. This difference can be magnified by calculating the $\Delta\sigma^2$ observable, as established in Chapter 4,




Figure 5.2: Left: Variance of azimuthal partition multiplicity distributions as a function of N for the most central AMPT data at 39 GeV. Single event data is shown in blue circles, mixed event data in green circles and the red line shows the expected binomial variance. Right: $\Delta\sigma^2$ is calculated for the same distributions and plotted against N .

which is the normalized deviation of the measured variance from the binomial expectation.

$$\Delta\sigma^2 = \frac{\sigma^2 - \sigma_{\text{binomial}}^2}{N(N-1)} \quad (5.1)$$

From the plot of $\Delta\sigma^2$ as a function of N on the right side of Figure 5.2, it is clear that while the mixed events from AMPT have the same width as the expected binomial distributions, the single event distributions are significantly narrower. This implies that the mixed event proton tracks are randomly distributed about the azimuth with no significant correlation while protons in the single AMPT events are negatively correlated, effectively repelling each other.

$\Delta\sigma^2$ is next calculated for STAR single event and mixed event data. The mixed event $\Delta\sigma_{\text{Mix}}^2$ is subtracted from the corresponding single event $\Delta\sigma_{\text{Single}}^2$ for each N , as described in Section 4.5. Figure 5.3 shows an example of the effect of this correction for the most central 39 GeV data. We see that $\Delta\sigma_{\text{Mix}}^2$ is consistently positive, indicating a positive correlation between tracks. This is due to the azimuthal efficiency inhomogeneities in the STAR detector, which the $\Delta\sigma^2$ observable interprets as clustering. The single event $\Delta\sigma^2$




Figure 5.3: $\Delta\sigma^2$ is plotted as a function of the total number of protons on the full azimuth, N , for the most central events in STAR 39 GeV data. The single event data is shown in blue circles, mixed event data in green squares, and the mixed corrected data in red triangles. Statistical uncertainties are shown as solid lines while systematic uncertainties are represented as transparent bars.

values are significantly negative and become slightly more negative after correcting for the mixed events.

Systematic error bars, shown on all STAR data plots, are estimated for each point at the time of plotting from the systematic variation data sets (see Section 5.7). Systematic uncertainties are never propagated via the error propagation formula. Statistical uncertainties for single and mixed event $\Delta\sigma^2$ are calculated from the standard deviation of the distribution of $\Delta\sigma^2$ values from 250 bootstrap sets (Section 3.4.2). These statistical uncertainties are then propagated algebraically. When subtracting the mixed data, the square of the statistical error on the difference is the sum of the squares of the single and mixed bootstrapped error. The $\Delta\sigma^2$ NLO uncertainty estimate from Equation 4.35 is then added to the statistical uncertainty in quadrature.

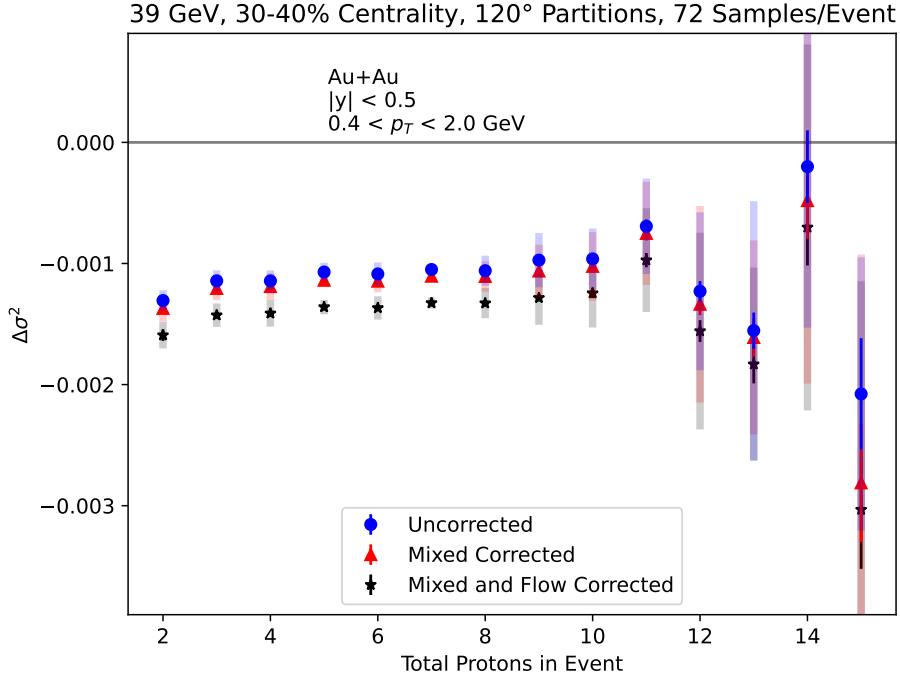


Figure 5.4: $\Delta\sigma^2$ is plotted for the 30-40% centrality class at 39 GeV. The uncorrected $\Delta\sigma^2$ is shown in blue dots, the mixed corrected data is shown in red triangles, and $\Delta\sigma^2$ after both mixed and v_2 correction is shown as black stars.

After correcting for mixed events, $\Delta\sigma^2$ is next corrected for elliptic flow. v_2 is measured for all protons within the centrality class and energy being analyzed, as described in Section 5.4. The $\Delta\sigma^2(w; v_2)$ contribution for this data set is calculated from Equation 4.39 using the measured v_2 and the partition width w of the data set being analyzed. In Figure 5.4 we see that the elliptic flow correction is small but significant in the 30-40% most central 39 GeV data. We note that the elliptic flow and therefore its correction is largest at this centrality. The proton multiplicity in peripheral events is lower, reflected in the contraction of the x-axis when compared to the most central events in Figure 5.3.

Figure 5.5 shows $\Delta\sigma^2$ after mixed event and elliptic flow corrections have been made in the most central events at 39 GeV. STAR data is compared to the AMPT (Appendix B.1) and MUSIC+FIST (Appendix B.2) models. We find that $\Delta\sigma^2$ in STAR data is significantly negative, indicating a negative correlation or repulsive interaction between proton tracks. A significantly negative $\Delta\sigma^2$ is also measured in AMPT data, though with a magnitude




Figure 5.5: $\Delta\sigma^2$ after mixed event and elliptic flow corrections is shown for the most central events at 39 GeV. STAR data is shown as black circles, AMPT is represented as a red band, the default MUSIC+FIST model is shown as a blue band, and the MUSIC+FIST model with an excluded volume effect is shown as a purple band.

about half of that measured in STAR. From the Gaussian correlation model described in Section 5.2, we find that the magnitude of $\Delta\sigma^2$ measures the strength of the correlation between tracks. This implies that, while both AMPT and STAR measure repulsion between proton tracks, the magnitude of repulsion is around twice as large in STAR data.

We use the MUSIC+FIST model to verify this interpretation. The default version of the model measures very little correlation and when an excluded volume effect is added a small but significant negative correlation is measured. The excluded volume effect is expected to produce an effective repulsion between baryons and it is encouraging that this is reflected in the measured $\Delta\sigma^2$. In addition, the MUSIC+FIST model conserves baryon number in each event, which seems to have little or no effect on the correlation between protons as measured by $\Delta\sigma^2$.

Figure 5.6 shows $\Delta\sigma^2$ as a function of N for all six BES-I energies analyzed. At all energies




Figure 5.6: $\Delta\sigma^2$ is shown as a function of N . Each panel represents one of the six BES-I energies analyzed.

and for both STAR and model data, we find that the measured $\Delta\sigma^2$ has no significant trend with N (11 GeV being a possible exception). This is an indication that the $N(N-1)$ normalization in Equation 5.1 as derived in Section 4.3.1 is describing the only significant dependence of the variance on the total number of protons per event. This justifies computing an average of $\Delta\sigma^2(N)$ over all N , weighted by the statistical uncertainties. These averages are plotted in Figure 5.6 as horizontal lines, with a transparent horizontal band representing the standard error on this weighted mean, $\sqrt{(\sum_i 1/\sigma_i^2)^{-1}}$. We will continue our analysis with these averages — $\langle\Delta\sigma^2\rangle$.

We note that detector inefficiency shifts the STAR data in Figure 5.6 to the left relative to the models. $\Delta\sigma^2$ measurements are made separately for each total number of protons in the event, N , which is also reduced by detector inefficiency. Since we observe no dependence of $\Delta\sigma^2$ on N , an azimuthally homogeneous detector inefficiency will simply shift the data in Figure 5.6 to the left, having no effect on the $\Delta\sigma^2$ measurement. This insensitivity to azimuthally homogeneous detector inefficiency is a convenient feature of this analysis.

Figure 5.7 shows the N averaged $\Delta\sigma^2$ as a function of center of mass energy for $w = 120^\circ$ partitions and the most central events. We find a large repulsive signal in STAR data which




Figure 5.7: $\Delta\sigma^2$ averaged over the total number of protons per event, N , is shown as a function of center of mass energy. STAR data is shown as black circles while AMPT and the MUSIC+FIST default and excluded volume models are shown as red, blue and purple bands, respectively.

displays no significant trend with energy. The MUSIC+FIST models are also relatively flat as a function of energy. The default MUSIC+FIST model $\langle \Delta\sigma^2 \rangle$ is slightly negative though very close to zero while the excluded volume version of the model is significantly negative, though with a magnitude much smaller than AMPT or STAR data. A repulsive interaction is also measured in AMPT data, with a magnitude that increases from about half that of STAR data at 62.4 GeV to almost comparable to STAR at 7.7 GeV. It is notable that, along with under-predicting the magnitude of repulsion found in STAR data, AMPT predicts a strong energy dependence which is not observed at STAR.

The initial motivation of this analysis was to search for proton clustering which may be indicative of clustering characteristic of a first-order phase transition. What we have found is a significant and consistent source of repulsion between proton tracks in both STAR data and the AMPT model. This is unexpected, and we will now try to determine the source of this repulsive signal.




Figure 5.8: $\langle \Delta\sigma^2 \rangle$ is plotted for each centrality class against the average reference multiplicity of that centrality class. The six AMPT energies are shown in the left panel as semi-transparent bands. STAR data is shown in the right panel as circles which are connected for each energy. A fit to AMPT data for all energies is displayed in both panels as a dashed line. Though not visible here, the fit clearly fails at high multiplicity where statistics are high and is shown only to guide the eye in comparison with STAR data.

5.1.1 Event Multiplicity Dependence

We investigate the centrality dependence of the observed repulsion in Figure 5.8. Here, $\langle \Delta\sigma^2 \rangle$ is plotted for each centrality class¹ as a function of the average reference multiplicity value for that class. We find that for both STAR and AMPT data there is a strong dependence on the reference multiplicity, which is a proxy for the total number of particles at mid-rapidity in each event. In the most central collisions with the largest event multiplicities, we find a significant but relatively small $\langle \Delta\sigma^2 \rangle$. As the event multiplicity decreases, $\langle \Delta\sigma^2 \rangle$ becomes more negative, indicating a stronger repulsive correlation between proton tracks. The magnitude of $\langle \Delta\sigma^2 \rangle$ becomes very large in the most peripheral collisions with the smallest event multiplicities, though uncertainties do increase as there are fewer protons in these events.

From Figure 5.8 we find that $\langle \Delta\sigma^2 \rangle$ is strongly dependent on the event multiplicity.

¹Eight centrality classes are shown, from 0-70%. The 70-80% class could be added but there are very few protons in these peripheral collisions which can lead to instability in these measurements. One issue encountered was an imaginary resolution ($\sqrt{\langle \cos(2(\psi_{\text{east}} - \psi_{\text{west}})) \rangle}$) in three of the systematic sets.




Figure 5.9: The 62.4 GeV STAR and AMPT $\langle \Delta\sigma^2 \rangle$ is fit as a function of reference multiplicity.

The AMPT energy dependence in Figure 5.7 can then be explained by the trend found in Figure 5.8 — the event multiplicity is smaller at lower energies and the repulsion therefore becomes stronger with decreasing energy. But, by this logic, STAR should see a similar energy dependence. Looking at the AMPT data in Figure 5.8, we find that $\langle \Delta\sigma^2 \rangle$ seems to follow a universal curve, independent of the energy. STAR data also seem to follow an energy independent curve at 62.4 and 39 GeV, but as energy decreases the energy curves become slightly less negative. The deviation of the STAR energy curves increases as energy decreases which counteracts the expected trend in Figure 5.7 due to lower energies having smaller event multiplicity.

This energy dependence of STAR data in Figure 5.8 could be consistent with a source of positive correlation that increases in magnitude as energy decreases in addition to a larger source of negative correlation which depends only on the event multiplicity. The positive correlation component in this scenario could be related to the clustering signal expected in a first-order phase transition. But to make any measurements of a positive component




Figure 5.10: The fit of $\langle \Delta\sigma^2 \rangle$ for 62.4 GeV is subtracted from all $\langle \Delta\sigma^2 \rangle$ values to better visualize the energy dependence for AMPT data on the left and STAR data on the right.

of the measured correlation, the much larger negative component must be understood and corrected.

This source of negative correlation is present in both STAR and AMPT data, rapidly increases in magnitude as event multiplicity decreases, and seems to approach zero correlation as event multiplicity increases. This rules out many detector effects which are expected to become worse and therefore cause stronger correlation as event multiplicity increases. One such effect is track merging, which is an expected background that has not yet been explored. It should, however, produce a negative correlation whose magnitude increases to some degree as the event multiplicity increases. We therefore conclude that this is not the observed background signal and is likely a rather small effect.

One possible source of this negative correlation is global momentum conservation. Intuitively, each particle coming from an event is partially responsible for ensuring that the net momentum of the event, specifically transverse momentum at mid-rapidity, remains zero. For an event with only two particles with finite transverse momentum, these two particles must have equal and opposite transverse momenta, corresponding to an exact anti-correlation. As more particles with finite transverse momentum are added to the event, there are more




Figure 5.11: The fit of $\langle \Delta\sigma^2 \rangle$ for 62.4 GeV is subtracted from all $\langle \Delta\sigma^2 \rangle$ values. These differences are averaged over centrality and plotted against beam energy.

degrees of freedom with which to satisfy conservation of transverse momentum. The negative correlation between particles therefore becomes smaller in magnitude as the number of particles increases. This matches the general trend observed in STAR and AMPT data.

We have constructed a simple simulation of pure momentum conservation, described in Appendix A.6. We find a trend with event multiplicity similar to what is observed in Figure 5.8 and that is highly consistent with an a/M fit. It seems that this behavior is not robust, however, as increasing the momentum of the simulated particles seems to produce an entirely different trend. We suspect that the observed reference multiplicity dependence of $\langle \Delta\sigma^2 \rangle$ is due in large part to momentum conservation, though we are yet unable to prove this or correct the effect.

In lieu of correcting the reference multiplicity dependence of $\langle \Delta\sigma^2 \rangle$, we attempt to subtract a proxy for the effect to highlight deviations. We note that, at least in STAR data, the curves in Figure 5.8 seem to converge as energy is increased. We therefore fit the 62.4 GeV data, the highest energy, with Equation 5.2 which seems to capture most of the trend.




Figure 5.12: $\langle \Delta \sigma^2 \rangle$ is plotted against the azimuthal partition width and fit to a quadratic.

$$\langle \Delta \sigma^2 \rangle = \frac{a}{M^c} + b \quad (5.2)$$

This fit, plotted in Figure 5.9, is subtracted from the other energies and the result is shown in Figure 5.10. The strong energy dependence in STAR data is more apparent and we note that the 7.7 GeV data in AMPT is consistently larger than the other energies. We can then perform an error weighted average over all centralities (using the statistical and systematic uncertainties added in quadrature) and plot the result as a function of center of mass energy in Figure 5.11. This shows the clearly increasing excess attraction with decreasing energy in STAR data.

5.1.2 Azimuthal Partition Width Dependence

The discussion up to this point has been limited to a partition width of $w = 120^\circ$. Figure 5.12 shows $\langle \Delta \sigma^2 \rangle$ as a function of the partition width for $w \in (60^\circ, 72^\circ, 90^\circ, 120^\circ, 180^\circ, 240^\circ, 270^\circ, 288^\circ, 300^\circ)$. We find a symmetry about 180° which we attribute to the symmetry of a binomial distribution about $p = 0.5$, though it is likely that a more precise explanation for this symmetry may exist.




Figure 5.13: The baseline b from the quadratic fit of $\langle \Delta\sigma^2 \rangle$ as a function of w is plotted against the reference multiplicity for AMPT on the left and STAR data on the right.

It is intuitive that changing the width of the partition may provide sensitivity to the azimuthal range of the correlation being measured. We investigate this with the Gaussian correlation model in Section 5.2 where we find that the curvature of $\langle \Delta\sigma^2 \rangle$ as a function of w is indeed monotonically related to the correlation range. We parameterize this curvature with the quadratic fit of Equation 5.3, whose z parameter is found to be most directly related to the Gaussian correlation range in the model.

$$\langle \Delta\sigma^2 \rangle = b \left[1 + \left(\frac{w - 180}{z} \right)^2 \right] \quad (5.3)$$

From Equation 5.3 we also extract the baseline parameter, b , corresponding to the maximum $\langle \Delta\sigma^2 \rangle$ magnitude which occurs at $w = 180^\circ$. The baseline is found to be directly related to the magnitude of correlation in the Gaussian correlation model. In Figure 5.13 we plot this baseline against the reference multiplicity to produce a plot extremely similar to Figure 5.8. The conclusions from this plot are the same as for the $w = 120^\circ$ case.

Finally we show the curvature (parameterized as z in Equation 5.3) from the $\langle \Delta\sigma^2 \rangle$ vs w fit as a function of the reference multiplicity in Figure 5.14. We find that the AMPT data show very little reference multiplicity dependence while the STAR data seem to follow




Figure 5.14: The curvature from the quadratic fit of $\langle \Delta\sigma^2 \rangle$ as a function of w parameterized in z is plotted against the reference multiplicity for AMPT on the left and STAR data on the right.

an energy independent curve of decreasing z as reference multiplicity decreases. This corresponds to a correlation range that increases as the reference multiplicity decreases. It is not immediately clear why the STAR data follow this trend. If the majority of the correlation being measured is a momentum conservation effect, it is possible that this increased in range is due to the increasing sparsity of tracks, causing correlation over longer distances. It is not clear why this explanation would not hold in AMPT.

5.1.3 STAR Analysis Conclusion

A strong repulsive signal has been observed in STAR BES-I data. The magnitude of this repulsion increases with decreasing event multiplicity, a trend also observed in the AMPT model and possibly attributed to a global momentum conservation effect. Energy dependence is observed in STAR data when plotting $\Delta\sigma^2$ as a function of event multiplicity while AMPT energies seem to follow a universal curve. It is possible that the energy dependence in STAR comes from a superimposed attractive correlation, though more theoretical development is needed to understand the repulsive background before any conclusion can be drawn about the nature of the dynamics responsible for this signal.

5.2 Gaussian Correlation Model

It was found in Section 4.3.1 that static azimuthal probability distributions can only produce positive $\Delta\sigma^2$ values, indicating attractive correlations. In data and models, however, we find consistently repulsive correlations. It is useful to model these types of negative correlations to test the $\Delta\sigma^2$ observable. To this end, a simple toy model was constructed, attempting to perturb independent Bernoulli trials to produce negative correlation between tracks.

Tracks in this model are sampled sequentially from a specific azimuthal probability distribution that depends upon the positions of all previously generated tracks. The PDF for the first track is flat — it's equally probable to be placed anywhere between 0 and 2π . The PDF of the second track is a flat distribution plus a Gaussian centered at the first track's location: $P(\phi_2) \propto 1 + \frac{A}{2\pi\sigma} e^{-\frac{1}{2}\left(\frac{\phi_2 - \phi_1}{\sigma}\right)^2}$ with the proper normalization. The second track is sampled from this PDF and the PDF for the third track is the product of the PDF from the first track multiplied by a similarly constructed PDF from the second track. This process continues and the PDF for the n th track is the product of contributions from all previous tracks:

$$P(\phi_n) \propto 1 \cdot \prod_{i=1}^{n-1} \left(1 + \frac{A}{2\pi\sigma} e^{-\frac{1}{2}\left(\frac{\phi_n - \phi_i}{\sigma}\right)^2} \right) \quad (5.4)$$

This model has two parameters: A which characterizes the amplitude or strength of correlation and σ which determines the correlation's angular range. The σ parameter is strictly greater than zero. Small σ values produce a sharp and thin peak, resulting in large correlations in the vicinity of the track but having little influence at further azimuthal distances. A large σ produces a weak but far reaching correlation. The A parameter can be positive to produce a positive correlation (clustering) or negative to produce negative correlation (repulsion). When the magnitude of A is small the strength of the correlation is small. Increasing the A parameter produces stronger correlations. The range of A , $[-1, +\infty)$, is asymmetric about zero. $A = -1$ makes it impossible for any track to be placed exactly

where a previous track exists while A must approach positive infinity in order to ensure a track is placed exactly where a previous track exists. This asymmetry means that positive and negative values of A should not be compared directly, though for small magnitudes there is rough equivalence for similar positive and negative magnitudes.

PDF Wrapping These probability distributions are over azimuthal space which is periodic from $[0, 2\pi)$. This is slightly problematic as the support for a normal distribution is $(-\infty, \infty)$. The PDFs must then be wrapped around this boundary such that a Gaussian centered on zero has a tail going in the positive direction from zero and its other half going in the negative direction from 2π . Practically, this is accomplished by calculating Equation 5.4 in $(-\infty, \infty)$ and then splitting this range into 2π chunks (eg ..., $[-4\pi, -2\pi)$, $[-2\pi, 0)$, $[0, 2\pi)$, $[2\pi, 4\pi)$, $[4\pi, 6\pi)$, ...). Each of these chunks are then simply superimposed on the $[0, 2\pi)$ range to produce the final azimuthal PDF.

To perfectly represent this function one would need to superimpose 2π chunks from the full $(-\infty, \infty)$ range. However, since the normal distribution approaches zero on both sides of this range, summing only a finite range and truncating the rest will serve as a good approximation. In practice, we determine this range based on the σ parameter of the model, and sum chunks out to n_σ away, where typically $n_\sigma = 8$. For example, for $\sigma = \pi$, the number of 2π chunks summed is $n_\sigma\sigma/2\pi = 8\cdot\pi/2\pi = 4$. So four chunks will be calculated on each side of the main $[0, 2\pi)$ range. This “wrapping” procedure allows for the approximation of the azimuthal PDF to an arbitrary precision specified by n_σ and dependent on the simulation’s σ parameter.

Limitations of Model While this simple model has been very useful in helping to understand the types of negative correlations found in data, it is important to note its limitations. Chief among these is the sequential placement of tracks. As each track is placed on the azimuth, it is influenced by all previously placed tracks. This is important because a negative

correlation cannot be generated by a static PDF — any perturbation of a flat distribution will only produce positive correlation between tracks. The sequential placement in this model leads to a dynamic PDF in which each track has information about the previously placed tracks in the event. Only with this information are tracks able to effectively spread out. This leads to an asymmetry, however, in that tracks placed first are influenced by fewer tracks than those placed later. While this seems to work for our purposes, the asymmetry of this procedure seems to lack elegance and may possibly miss some subtle correlations that may exist in nature.

An additional difficulty of this model is its coupling with other models when trying to combine different types of correlation. Take, for instance, a ϕ dependent efficiency which is flat apart from the range $(\pi/3, 2\pi/3)$ in which the probability of producing a track is zero. If the Gaussian correlation model is run along side a static efficiency PDF of this form, the model will “interpret” the hole in the efficiency distribution as a void in ϕ space and try to fill it if the correlation being modeled is repulsive ($A < 0$) or avoid it if attractive ($A > 0$). This necessarily produces a coupling between the efficiency pdf and the Gaussian correlation model, which we would not expect to be the case in an actual experiment (this would correspond to the collision knowing about and being influenced by the efficiency of the detector). A solution for this problem would be to generate the tracks with the Gaussian correlation model first and then reject based on the efficiency, though this produces difficulties in ensuring the proper number of tracks are produced in a given event. This coupling effect reduces the utility of this model as a stand in for physical correlation between tracks. These couplings, however, are typically small, so with some interpretation this model can still be useful when coupled to other sources of correlation.

5.2.1 $\Delta\sigma^2$ Analysis

This Gaussian correlation model can be used to test the $\Delta\sigma^2$ observable. Varying the A parameter allows us to test the observable’s response to different correlation magnitudes




Figure 5.15: Simulation $\Delta\sigma^2$ plotted as a function of the total number of protons per event. Six simulations are run in total, three repulsive and three attractive with $\sigma = 1$ and $A \in \pm[0.002, 0.006, 0.01]$.

while varying σ gives a handle on its response to correlation range. Events are generated in this model given a fixed number of tracks per event, between 2 and 80, in order to match the maximum range in AMPT and STAR data. $\Delta\sigma^2$ is then calculated for these generated events from the azimuthal partition multiplicity distributions. We simulate various combinations of the A and σ parameters.

For simulations with a handful of different A values, we plot $\Delta\sigma^2$ as a function of the total number of protons per event in Figure 5.15. In this model, there are no detector effects to capture so the mixed events will be statistically identical to binomial distributions and $\Delta\sigma_{\text{Mixed}}^2$ will therefore be zero. No flow is necessarily simulated, so no flow correction will be made unless otherwise specified. Six simulation sets are shown in Figure 5.15 with $\sigma = 1.0$ — three attractive and three repulsive. The attractive sets all exhibit a positive $\Delta\sigma^2$ while the repulsive sets are all negative. In addition, larger A values correspond to larger $\Delta\sigma^2$ magnitudes. This example suggest that the $\Delta\sigma^2$ observable may be correlated with the




Figure 5.16: $\langle \Delta\sigma^2 \rangle$ plotted on the y-axis against the amplitude (A) simulation parameter on the x-axis. Both attractive and repulsive simulations are run with $\sigma \in [0.8, 1.0, 1.2]$ and $A \in \pm[0.002, 0.004, 0.006, 0.008, 0.01]$. A linear fit is superposed for each of the six sets shown, illustrating the linearity between $\langle \Delta\sigma^2 \rangle$ and A .

simulation's correlation amplitude, A . As described in Section 5.1, since the $\Delta\sigma^2$ values are flat with respect to the total number of protons in the event, we average over this variable and work with this average $\langle \Delta\sigma^2 \rangle$ (or $\overline{\Delta\sigma^2}$) from here on.

To quantify the relationship between $\langle \Delta\sigma^2 \rangle$ and the A parameter, we plot the two in Figure 5.16. From the superposed fit lines, we can see that there is a linear relationship between $\langle \Delta\sigma^2 \rangle$ and the A parameter in both the attractive and repulsive simulations. This linear relation is dependent upon the value of σ used in the simulation. In addition, extrapolating the linear fits to $A = 0$ gives a $\langle \Delta\sigma^2 \rangle$ intercept of zero as expected since a simulation with $A = 0$ corresponds to totally uncorrelated tracks. This indicates that the measured $\langle \Delta\sigma^2 \rangle$ is proportional to the correlation amplitude as it is parameterized in this model.

Next we will try to extract the σ parameter from the simulation data as we did the A parameter. In the Gaussian correlation model, σ parameterizes the range of the correlation. We might then expect that information about this range may be accessible through variation




Figure 5.17: $\langle \Delta\sigma^2 \rangle$ plotted against the width of the azimuthal partitions used. Simulation Amplitudes of $A \in \pm[0.002, 0.004, 0.006, 0.008, 0.01]$ are shown with $\sigma = 1$.

of the width of the azimuthal partitions. We plot $\langle \Delta\sigma^2 \rangle$ as a function of the azimuthal partition width in Figure 5.17. In this plot, all simulations are run with the same correlation range of $\sigma = 1$. The magnitude of $\langle \Delta\sigma^2 \rangle$ is linearly proportional to the simulation amplitude A for all w . This was demonstrated for $w = 120^\circ$ in Figure 5.16 but is here found to be true for any width. The constant of proportionality is found to have strong dependence on the azimuthal partition width. In Figure 5.17 the maximum $|\langle \Delta\sigma^2 \rangle|$ for each simulation occurs at $w = 180^\circ$. The magnitude of the observable then decreases from this maximum symmetrically on either side. This general trend seems to hold for all 10 simulations shown — the curvature of the attractive simulations are negative and the curvature of the repulsive simulations are positive. These trends can be roughly described as quadratic and fits of form Equation 5.5 are superimposed. These fits consistently undershoot the magnitude of the $w = 180^\circ$ point and are therefore not perfect, though they are found to be useful.

$$\langle \Delta\sigma^2 \rangle = b \left[1 + \left(\frac{w - 180}{z} \right)^2 \right] \quad (5.5)$$




Figure 5.18: $\langle \Delta\sigma^2 \rangle$ plotted against the width of the azimuthal partitions used. Simulation Amplitudes of $A \in -[0.006, 0.01]$ are shown with $\sigma \in [0.4, 1.2]$.

The fit is parameterized such that the curvature (z) is expressed in units of w . The b parameter corresponds to the maximum magnitude of the fit, what we will call the baseline. The z parameter corresponds to the distance in angle between $w = 180^\circ$ and the fit's x-intercept. Visually, it is apparent that all simulation fits seem to intersect the x-axis around the same place in Figure 5.17 and therefore have similar z values. This is found to be true for simulations with other σ values, though the intersection point changes with σ . In addition, the intersection can be different between the attractive and repulsive simulations. Focusing on repulsive correlations, Figure 5.18 shows that simulations with the same value of σ intersect the x-axis at the same place and this intersection changes with σ .

In Figure 5.19 we plot the correlation between the $\langle \Delta\sigma^2 \rangle$ curvature, z , and the simulation azimuthal range, σ . The grouping of the points at each x-axis value demonstrates that the curvature is not heavily dependent on the simulation amplitude A parameter. As the simulation σ parameter increases, the z parameter decreases slightly. This trend is not nearly as significant as the one for correlation amplitude in Figure 5.16, though it still




Figure 5.19: The measured curvature of $\langle \Delta\sigma^2 \rangle$ as a function of azimuthal partition width, w is parameterized by z and plotted against the azimuthal correlation range from simulation, σ . Simulation Amplitudes of $A \in -[0.002, 0.004, 0.006, 0.008, 0.01]$ are shown with $\sigma \in [0.8, 1.0, 1.2]$.

demonstrates a monotonic decrease of z as σ increases. Larger curvatures of $\langle \Delta\sigma^2 \rangle$ vs w (smaller z) therefore imply a longer azimuthal correlation range.

The dependence of the quadratic fit curvature on the simulation range can be alternatively visualized in Figure 5.20 which plots the quadratic z parameter vs the baseline b parameter. Simulations are grouped by σ and their average is plotted as a horizontal line. Attractive simulations have a positive b parameter and are grouped separately from repulsive simulations which have a negative b .

We have demonstrated that $\langle \Delta\sigma^2 \rangle$ is linearly related to the simulation A parameter at $w = 120^\circ$ and stated that this is also true for all other partition widths. In this case, if $\langle \Delta\sigma^2 \rangle$ is well described by a fit, we would expect the fit function to also exhibit a linear relationship with A , at least at the fitting points (at which the fit and the values are ideally identical). From there, we may expect that this linearity holds everywhere on the fit curve, if the curve accurately interpolates and extrapolates the true trend in $\langle \Delta\sigma^2 \rangle$. In this case,




Figure 5.20: Parameters of the quadratic fits to $\langle \Delta\sigma^2 \rangle$ vs partition width w are plotted with the z parameter on the y-axis and the b parameter on the x-axis. The z parameter indicates the distance from 180° at which the fit intersects the x-axis and is inversely related to the curvature. The b parameter represents the maximum magnitude of the quadratic which, by definition, occurs in this parameterization at 180° . Simulation Amplitudes of $A \in \pm[0.002, 0.004, 0.006, 0.008, 0.01]$ are shown with $\sigma \in [0.8, 1.0, 1.2]$. Repulsive simulations have a negative b and are on the left of the plot while attractive simulations are on the right. The simulations are grouped by σ separately for attractive and repulsive. The mean is calculated for each group and represented as a horizontal line, with the standard error on the mean represented as a shaded horizontal band.




Figure 5.21: The baseline, b , of the quadratic fit of $\langle \Delta\sigma^2 \rangle$ vs w is plotted against the simulation amplitude A for both attractive and repulsive simulations.

we would expect that the maximum of the quadratic fit, parameterized by b which is defined as the value of the curve at $w = 180^\circ$, would share a linear relation with the simulation amplitude A . To confirm, we plot b vs A in Figure 5.21.

From Figures 5.20 and 5.21 we conclude that the quadratic fit to $\langle \Delta\sigma^2 \rangle$ vs the azimuthal partition width w can provide information about the azimuthal correlation between the data set's proton tracks. The baseline b parameter of the fit measures the amplitude of this correlation and its sign: positive for attractive correlations and negative for repulsive. The curvature of $\langle \Delta\sigma^2 \rangle$ as a function of the azimuthal partition width is encoded in the z parameter, which increases for decreasing azimuthal correlation range.

While this Gaussian correlation model is rather simple, we believe that these tests of the $\langle \Delta\sigma^2 \rangle$ observable motivates performing similar quadratic fits on STAR and model data and interpreting the b parameter of these fits as a proxy for correlation strength and sign and the z parameter as a proxy for the correlation range.

Energy	Triggers	# Events (M)
7.7	290001, 290004	2.1
11.5	310004, 310014	3.8
19.6	340001, 340011, 340021	10
27	360001	19
39	280001	57
62	270001, 270011, 270021	28

Table 5.1: Minimum bias triggers and number of events used in analysis.

5.3 STAR Data Set

Six collision energies from the STAR BES-I data set, collected between 2010 and 2011, are analyzed. Table 5.1 lists each energy analyzed along with the triggers used and the number of events used in the analysis. We require at least two good proton tracks within our acceptance to measure a correlation, so this event number will be slightly lower than other STAR analyses.

5.3.1 Event Level Cuts

Figure 5.22 shows the number of remaining events at 7 GeV after each of the various event cuts are made. The first significant cut is to remove bad runs based on run averaged values such as: reference multiplicity, DCA, transverse momentum, pseudo-rapidity, or azimuthal angle. If any of these average values for a run are beyond 3σ from the mean of the data set, the run is marked as bad and removed from the analysis.

The next significant cut is on the primary vertex longitudinal (v_z) and radial (v_r) position. The primary vertex is required to be radially within 2 cm of the center of the beam pipe for all energies. The longitudinal position of the primary vertex is required to be within 30 cm of the center of the TPC for all energies except 7 GeV where the requirement is relaxed




Figure 5.22: Number of events at each stage of event level cuts for 7 GeV. The left panel shows the cuts made before downloading the data locally. The right side shows the cuts made locally for the default analysis. The cuts on the right are varied for systematic checks.

to 50 cm to improve statistics. For 39 and 62.4 GeV the TPC v_z position must match the v_z position determined by the VPD within 3 cm. Events passing all of these cuts (the left panel of Figure 5.22) are stored in ROOT trees to expedite further analysis.

The event level cuts listed on the right panel of Figure 5.22 pick up where the left panel leaves off (“Original” on right is equivalent to “Good VPD Vz” on left). The DCA_xy (Section 5.5) and Pile Up (Section 5.6) cuts are varied for systematic checks, so are performed with the other cuts on the right when the azimuthal partition analysis is run on the generated trees.

5.3.2 Track Level Cuts

Table 5.2 lists the track level cuts for selecting the protons used in our analysis along with the criteria for counting the reference multiplicity variables. In all cases, primary tracks are first selected, defined as those global tracks which intersect the primary vertex.

“Refmult” is the default reference multiplicity definition used in STAR for defining event centrality. Refmult counts all primary tracks with a pseud-rapidity between ± 0.5 and a distance of closest approach (DCA) of less than or equal to 3 cm. In addition, the number of TPC hits used in the helical fit is required to be 10 or higher and the momentum is required




Figure 5.23: Normalized primary vertex v_z (left) and reference multiplicity 3 (right) distributions.

to be greater than 10^{-10} GeV/c to remove track artifacts. In this analysis we use refmult only for quantifying the event multiplicity in Figures 5.8, 5.13, and 5.14.

For defining the event centrality, we use the “refmult3” definition. The goal of refmult3 is to exclude any protons (or anti-protons) that we may use in the analysis to avoid auto-correlation between the analysis and the centrality. This exclusion is accomplished by counting tracks within a pseudo-rapidity of ± 0.5 and requiring that the $n\sigma_{\text{proton}}$ of TPC tracks is less than -3 which selects only particles lighter than the proton — mostly pions. In addition, for tracks which have TOF information, the m^2 as determined by the time of flight is required to be less than 0.4 GeV. The refmult3 multiplicity distribution (right side of Figure 5.23) is then compared to a similarly defined multiplicity in Glauber simulations (Section 2.6) to define the centrality class thresholds. In each event, refmult3 is compared to these thresholds to determine the centrality class of the event.

Similar to the reference multiplicity counting, analysis protons are selected as those tracks with a positive charge, a DCA less than or equal to 1 cm, the number of TPC hits fit greater

Track Variable	Refmult	Refmult3	Protons
Charge	-	± 1	+1
$ \eta $	< 0.5	< 1	-
DCA	≤ 3 cm	≤ 3 cm	≤ 1 cm
nHitsFit	≥ 10	≥ 10	> 20
nHitsFit / nHitsMax	-	-	≥ 0.52 and ≤ 1.05
$n\sigma_{\text{proton}}$	-	< -3	≥ -2 and ≤ 2
m_{TOF}^2	-	< 0.4	> 0.6 GeV and < 1.2 GeV
p	$\geq 10^{-10}$	$> 10^{-10}$	≤ 1 GeV without TOF ≤ 3 GeV with TOF
p_T	-	-	≥ 0.4 GeV ≤ 0.8 GeV without TOF ≤ 2.0 GeV with TOF
$ y $	-	-	< 0.5
nHitsDedx	-	-	> 5

Table 5.2: Track level cuts for selecting tracks counted in the reference multiplicity variables and for selecting proton tracks for analysis.

than 20, and the absolute value of the TPC $n\sigma_{\text{proton}}$ less than 2². The ratio of the number of TPC hits fit divided by the maximum number of hits possible is required to be between 0.52 and 1.05, this upper bound accounting for floating point comparison. The number of hits used to determine the track energy loss in the TPC, dEdx, is required to be greater than 5 to ensure a reliable measurement. Instead of the pseudo-rapidity, the proton *rapidity* is required to be between ± 0.5 . Proton tracks with transverse momentum between 0.4 GeV and 0.8 GeV and passing all other criteria are accepted as good protons. Protons with transverse momentum between 0.8 GeV and 2.0 GeV are required to be matched to a TOF hit with m^2 between 0.6 GeV and 1.2 GeV.

The distributions of these variables for protons passing all cuts are shown in Figure 5.24. A set of proton tracks with looser cuts than those listed in Table 5.2 are retained in order to perform systematic studies which may include a relaxation of specific cuts. This looser set of protons includes those with $|\eta| \leq 2.1$, $n\text{HitsFit} > 15$, $\text{DCA} < 3$ cm, $n\sigma_{\text{proton}} < 2.5$, and $0.3 \text{ GeV} \leq p_T \leq 2.2 \text{ GeV}$. Figures 5.25 and 5.26 show the distributions for the azimuthal angle and rapidity, respectively, for the set of protons with loose cuts and the set with the default analysis cuts.

²For 27 GeV the $n\sigma_{\text{proton}}$ is compressed by a factor of 2 (Figure 5.24), and all cuts involving this variable are similarly scaled for the 27 GeV data set.

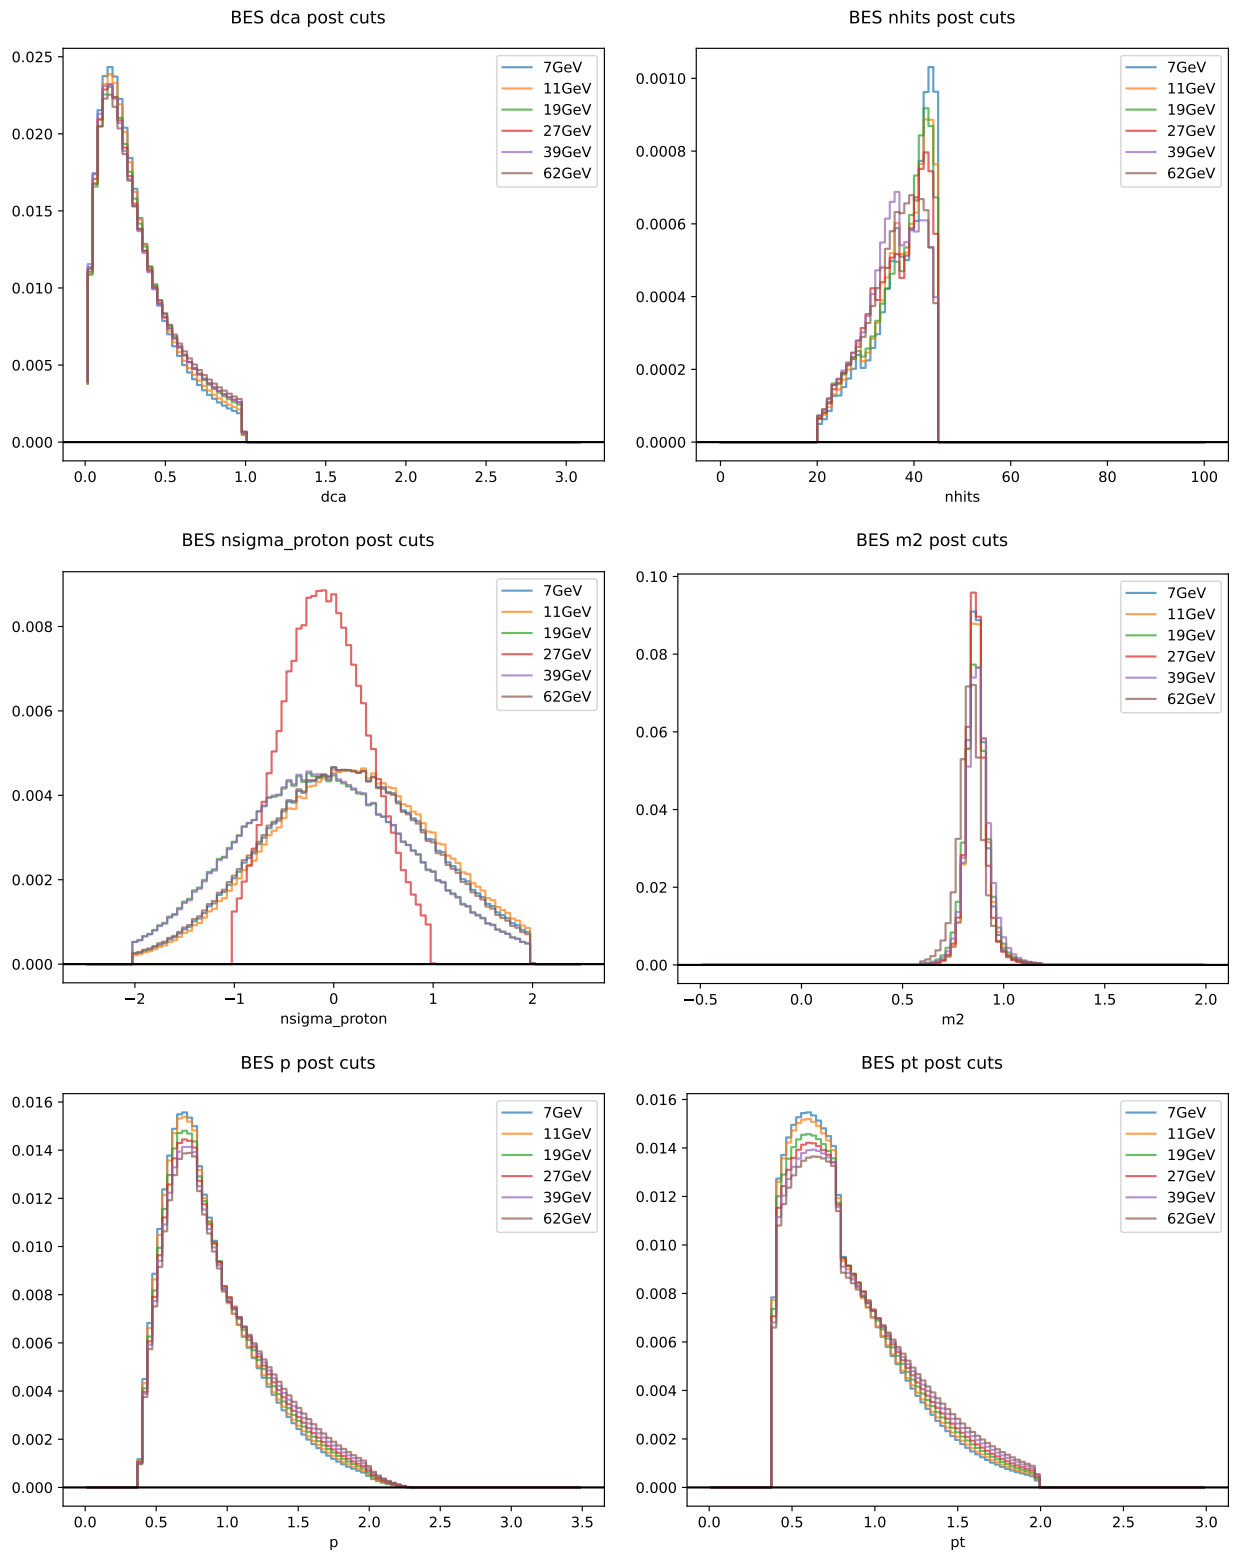


Figure 5.24: Distributions of track variables for protons used in the analysis.

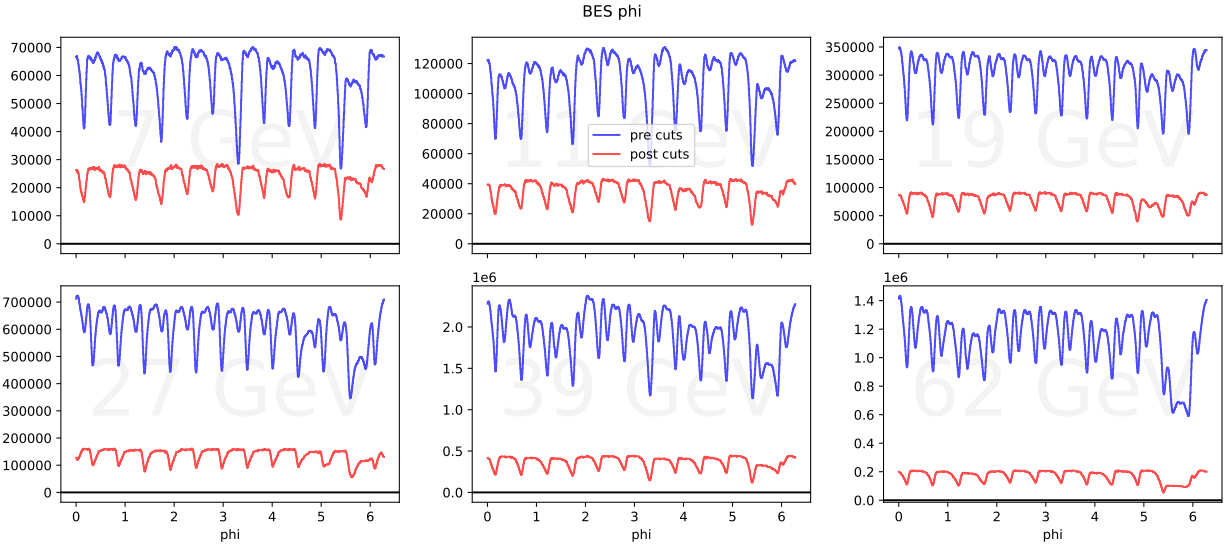


Figure 5.25: Azimuthal angle distributions of protons used in analysis. The TPC sector boundaries are clearly visible. The distributions in red correspond to the protons used in the default version of the analysis while the blue distributions correspond to a larger set of protons (including the red distribution) with much looser cuts, used for systematic variations. Most notably, the blue distribution includes protons with $|\eta| \leq 2.1$ while the $|y| < 0.5$ cut is performed to get the red distributions.

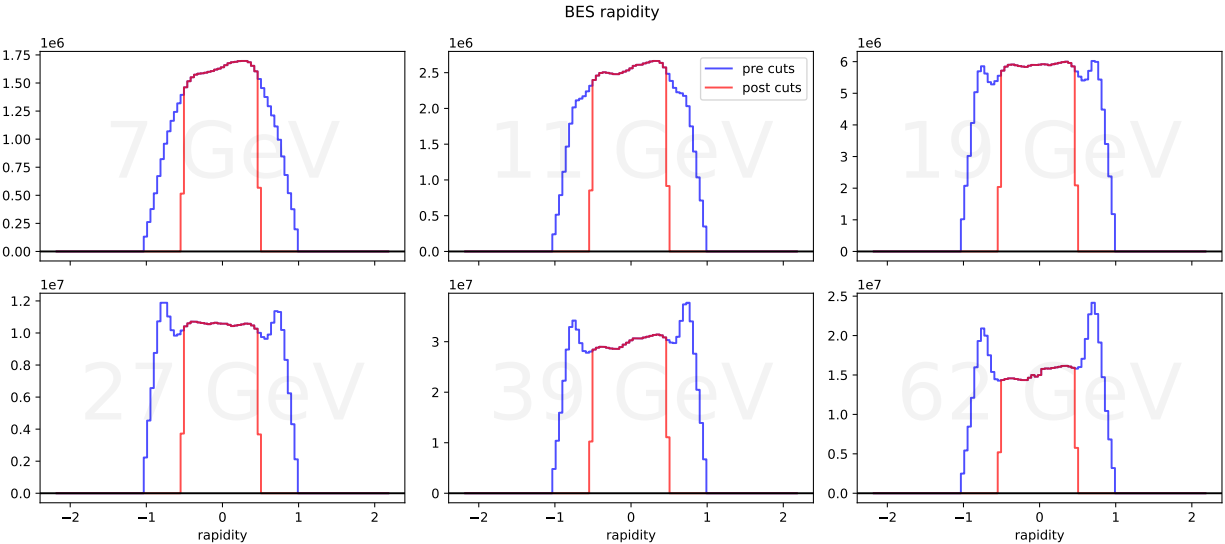


Figure 5.26: Rapidity distributions of a loose selection of protons in blue and the protons after all default cuts have been made in red.

5.4 Elliptic Flow Measurement

The correction for elliptic flow explained in Section 4.4 depends upon the measurement of v_2 for our protons of interest. For this measurement we utilize the sub-event method in the TPC. In STAR data, the particle and event plane azimuthal distributions must be flattened before the elliptic flow can be measured.

5.4.1 Azimuthal Distribution Flattening

Due to azimuthal inhomogeneities in the STAR detector, we must flatten our phi distributions before conducting any flow analysis. The assumption is that, over many events, any azimuthal multiplicity distribution should be flat since the event plane is random. When the azimuthal distribution is not flat, this is due to phi-dependent detector inefficiencies. To correct for these, we first integrate the azimuthal distribution over all events in our data set, splitting events into classes based on the time they were taken (run number), their pseudorapidity value, and the type of particle. We then decompose these distributions into a Fourier series and slightly rotate all angles based on these coefficients such that the resulting distribution is flat [49]. This flattening is also performed on the event plane distribution.

5.4.1.1 Particle Distribution Flattening

We first flatten the azimuthal distributions of particles. We select particles with $nHitsFit > 15$, $dca < 2$ cm, $|\eta| < 1$ and $0.2 < p_T < 2$ GeV. We separate particles into classes based on three criteria: whether the particle is a proton that will be used in our analysis or not, the pseudorapidity, and the run number. We identify protons that will be used in the actual analysis as those with $nHitsDedx > 5$, $dca < 1$ cm, $0.3 \text{ GeV} \leq p_T$, $|n\sigma_{\text{proton}}| < 2$, $|y| < 0.5$ and TOF $0.6 < m^2 < 1.2$ GeV if TOF is used. We also split the rapidity range from $-1 < \eta < 1$ into 4 sub-ranges in steps of 0.5. We finally split the data set on the run number, grouping all runs taken on the same day together (the thousands place of the run

number enumerates the day).

Once we have azimuthal particle distributions for all of these classes, we perform a Fourier decomposition up to order 12 for each. We then use these Fourier coefficients to flatten the azimuthal distributions in subsequent analysis steps.

5.4.1.2 Event Plane Angle Flattening

We use the flattened azimuthal distributions of the non-protons to calculate the event plane in each event. We calculate the second order event plane by summing the p_T weighted sin and cos of twice each azimuthal angle and then determine the angle of the resulting Q vector [49]:

$$\begin{aligned} q_x &= \sum_i p_{t_i} \cos(2\phi_i) \\ q_y &= \sum_i p_{t_i} \sin(2\phi_i) \end{aligned} \tag{5.6}$$

$$\psi = \frac{1}{2} \tan^{-1} \left(\frac{q_y}{q_x} \right) \tag{5.7}$$

We again split events into classes using their run number based on the day the data was taken. Instead of the four η bins, we calculate the event plane for the east ($\eta < -0.2$) and west ($\eta > 0.2$) side of the detector separately. We include a pseudorapidity gap of 0.4 to reduce the effects of short range correlation between protons on one side of the detector with particles used for determining the event plane on the other side. The Fourier coefficients of these event plane distributions up to order 12 are then extracted and used to flatten the event plane when measuring v_2 .




Figure 5.27: STAR event plane resolution (left) and v_2 of identified protons (right) as a function of centrality for six BES-I energies.

5.4.2 Measuring v_2

With the particle and event plane azimuthal distributions flattened, v_2 can be measured. In each event a v_2 value for each proton is calculated as the cosine of twice the difference between its azimuthal angle and the event plane angle. To avoid auto-correlation, if the proton is on the east side of the detector ($\eta < 0$) it is compared to the west side event plane angle. If it's on the west side ($\eta > 0$) it is compared to the east side event plane angle. The v_2 for all protons in the event is calculated and averaged. The v_2 of the data set is obtained by averaging over all event v_2 measurements.

$$v_2^{\text{obs}} = \langle \cos(2(\phi - \psi)) \rangle \quad (5.8)$$

This measured v_2 is slightly suppressed due to the finite resolution of the measured event plane angle. If the measured event plane angle is offset from the true reaction plane, the proton ϕ angles will be compared to the wrong angle and the resulting v_2 values will be smaller on average due to this misalignment. The event plane resolution can be calculated from the ψ_{east} and ψ_{west} sub-event planes with Equation 5.9. To correct for this resolution suppression, the “observed” v_2 of Equation 5.8 is divided by the resolution.




Figure 5.28: AMPT event plane resolution (left) and v_2 of identified protons (right) as a function of centrality.

$$\text{Resolution} = \sqrt{\langle \cos(2(\psi_{\text{east}} - \psi_{\text{west}})) \rangle} \quad (5.9)$$

$$v_2 = v_2^{\text{obs}} / \sqrt{\langle \cos(2(\psi_{\text{east}} - \psi_{\text{west}})) \rangle} \quad (5.10)$$

5.4.3 STAR Proton v_2 Measurement

The elliptic flow of protons was measured in STAR BES-I data in 9 centrality bins. The event plane resolution and v_2 (Figure 5.27) results are comparable to similar measurements published by STAR [50]. This STAR paper, however, did not contain measurements of protons for the specific centrality bins needed in this analysis.

5.4.4 AMPT Proton v_2 Measurement

Measuring flow in AMPT (Appendix B.1) and other models is considerably easier, as there are no detector effects and therefore no need for flattening azimuthal distributions. Beside the flattening procedure, the resolution and v_2 of protons is measured in the same way as




Figure 5.29: Event plane resolution (left) and v_2 of identified protons (right) for STAR, AMPT, and all three versions of the MUSIC+FIST model as a function of energy.

STAR data, using Equations 5.9, 5.8 and 5.10. The resulting resolution and v_2 measurements made on AMPT data are shown in Figure 5.28.

5.4.5 MUSIC+FIST Proton v_2 Measurement

The MUSIC+FIST model (Appendix B.2), similar to AMPT, doesn't need any azimuthal flattening. Only 0-5% most central events are generated for this model and the resolution and v_2 are therefore plotted as a function of energy in Figure 5.29 for each version of the model and compared to AMPT and STAR data. No MUSIC data was available at 11 GeV.

5.5 Signed Distance of Closest Approach Bad Events

The Distance of Closest Approach (DCA) is the closest distance a reconstructed track comes to the primary vertex of an event. If the track originated from the primary vertex and was reconstructed perfectly, the track would intersect the primary vertex and its DCA would be zero. In reality, there is a finite resolution for both the reconstruction of the primary vertex location (relatively small uncertainty when there are enough tracks in the event) and the reconstruction of each particle track. The track reconstruction is the larger source of uncertainty and leads to a finite DCA resolution, which is typically on the order of a few tenths of a centimeter.

In the BES-I data set, it was found in some events that the signed DCA variable (DCAxy) (slightly different from but similar to the DCA) was abnormal. Runs with average DCAxy values greater than 3 sigma (4.5 sigma for 19 GeV and 4 sigma for 62 GeV to adjust for smooth and significant run variations) from the mean are discarded, but this only occurred for one 27 GeV run. It is ok that almost no runs are cut, as the real QA work in this case is done on an event-by-event basis. The run averaged DCAxy variables for each energy are shown in Figure 5.30.

The more prominent issue was the deviation of the *event* averaged DCAxy variable. Especially at the beginning or end of some runs, the event averaged DCAxy variable could undergo a very large shift, possibly due to conditions caused by an unstable new beam or beam close to being dumped. To find and remove outliers, a set of moving averages were calculated on the event averaged DCAxy for the most central 20% of events, since the primary vertex of peripheral events is less precise. The trivial 1 point moving average is just the value of each event average. An outlier at this level is defined as any event 5σ away from the run average. These 5σ outliers and any subsequent events are removed until an event falls within a “relaxation window”, in this case 0.2σ from the run average. Next, a 5 event moving average was calculated in each run. Since the moving average is an average of 5 points, a 5.5σ

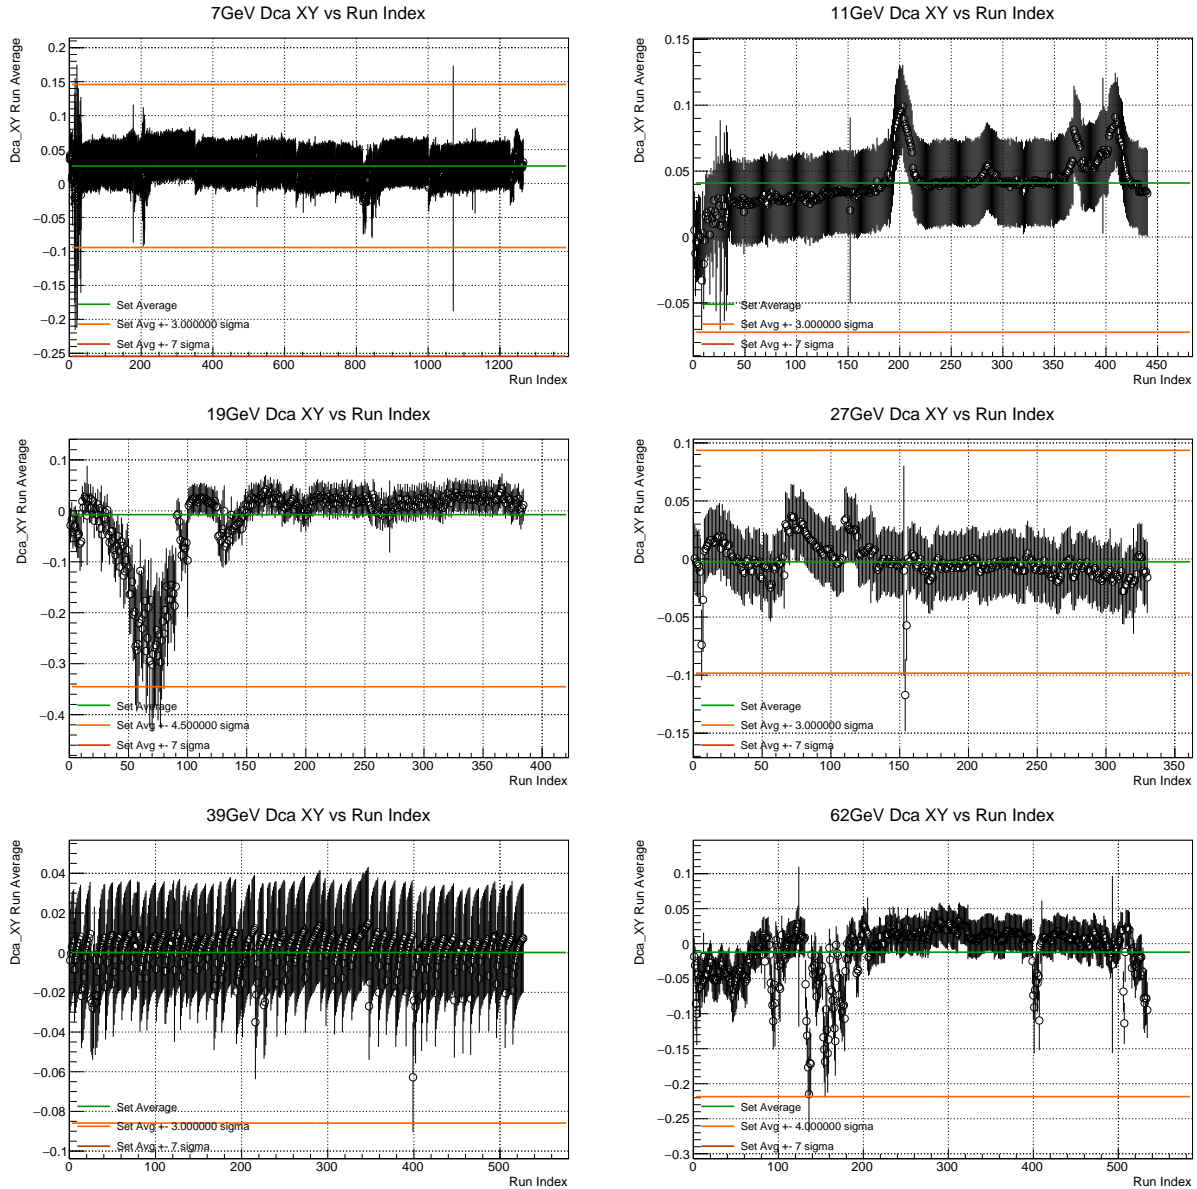


Figure 5.30: Dcaxy run average vs run index for all six BES-I energies analyzed. Only the most central 20% of events were averaged.




Figure 5.31: 7.7 GeV Dcaxy event average vs event index for run 11126001. The purple rectangles enclose bad event ranges which are excluded from the analysis.

of this average actually corresponds to a smaller deviation than the corresponding 5.5σ value for the 1 point moving average. This corresponding deviation value is calculated with the function below, where `new_points` is the number of moving average points and `old_points` is 1. The assumption is that each of the 5 points is independent, and the probability of average of these 5 points deviating from the mean is calculated. For the 5 point moving average, this corresponding number of σ s is calculated (smaller than 5.5, as a smaller deviation of an average is more significant) and outliers are removed. In addition, there is a relaxation value of 2σ , and all subsequent events before this 2 sigma relaxation threshold is reached are also removed. Finally, a 10 point moving average is calculated and treated similarly, with an 8σ trigger and 4σ relaxation.

In the discussion above, deviation of an event is taken with respect to a linear fit of the run, rather than an average. This is necessary as the DCAxy was found to slowly drift during runs. In addition, for an entire data set (energy) the average event DCAxy is calculated along

	default	very loose	loose	tight	very tight
1 point	5	7	6	4	3
5 point	5.5	7.5	6.5	4.5	3.5
10 point	8	10	9	7	6

Table 5.3: Dcaxy default and systematic trigger $n\sigma$ values.

with the 7σ deviation. If any individual event is greater than 7σ away from this set average, it is marked as bad.

The benefit of the 5 and 10 event moving averages is that they eliminate significant jumps in the DCaxy for groups of events. This allows the QA analysis to be sensitive to how the observable changes in time rather than just using uncorrelated event values. An example of a run at 7.7 GeV is shown in Figure 5.31. The purple rectangles indicate ranges of bad event flagged by at least one of the above criteria (1, 5, 10 event moving average or more than 7 sigma from the 7 GeV data set mean). Events within these rectangles are excluded from analysis.

For a systematic check, the $n\sigma$ trigger values are varied in the 1 point case from the default of 5 to 4 (tight), 3 (very tight), 6 (loose), and 7 (very loose). The variations for the other moving averages are listed in Table 5.3.

```
float convert_sigmas(float old_sigma, int new_points, int old_points) {
    double exponent = (double)old_points / (double)new_points;
    double prob = pow(ROOT::Math::normal_cdf(old_sigma) * 2, exponent)/2;
    return ROOT::Math::gaussian_quantile(prob, 1);
}
```




Figure 5.32: Left: Example of BTOFMatch vs Refmult correlation at 62.4 GeV. Right: The histogram on the left is rotated such that a linear fit is vertical. This should ensure that horizontal slices are Gaussian if the two detectors are linearly related.

5.6 Pile Up Rejection

As discussed in Section 2.7, out-of-time pile up may be rejected by comparing the multiplicities of a fast detector such as the TOF to those of the slower TPC. We compare a quantity representative of the total number of tracks reconstructed per event in the TPC to a similar quantity in the TOF. These quantities should be linearly related — when the TPC detects many tracks, so should the TOF. This is no longer true in a pile-up scenario, and events which deviate from linearity in this TPC/TOF comparison can be rejected as likely pile-up events.

For this TPC/TOF pile-up rejection, we utilize the refmult quantity from the TPC and the BTOFMatch quantity from the TOF. An example of the linearity between these two quantities is shown on the left of Figure 5.32. The pile-up events here are apparent — there are many events below the dense linear region. These correspond to events with refmult (TPC) values disproportionately large compared to the BTOFMatch (TOF) values. These are then the events we would like to reject as likely pile-up candidates.




Figure 5.33: Example horizontal slice of the rotated 62.4 GeV BTOFMatch vs refmult distribution. A Gaussian fit alone (left) is unable to model these distributions adequately. Fitting with a Gaussian plus two Woods-Saxon distributions multiplied by decaying exponential functions (right) seems to produce the most stable fitting results. From these fits, the $n\sigma$ x-axis locations with respect to the mean of the Gaussian distribution are calculated (yellow vertical lines).

To implement a pile-up rejection based on the BTOFMatch and refmult correlation, we first rotate the two dimensional histogram such that the linear fit (red line in the original histogram) is vertical, as shown on the right of Figure 5.32. The rotation isn't perfect, as the original histogram is discrete and this produces the artifacts seen in the rotated histogram (white dots where the cells are empty). Since the original histogram had integer binning on both axes and the max multiplicities was in the hundreds, the resolution is good enough for these artifacts to be unimportant.

With the rotated histogram, one dimensional horizontal slices are taken with slice widths of 5 vertical units. These slices (Figure 5.33) contain a large Gaussian portion corresponding to the linear response between the two detectors. There is also clearly a non-Gaussian tail to the right coming from the pile-up events we hope to reject. We would like to fit this distribution for each slice and make an adjustable $n\sigma$ cut on the right (and also the left). This fit requires that we somehow model the contribution of the pile-up events superimposed on




Figure 5.34: The $n\sigma$ Gaussian cuts for each slice are rotated back to the original coordinates and fit with fourth order polynomials (extrapolated at the high end). These red lines define the boundary between the good events between then and the pile-up events above and below. The original green fitting points are difficult to see as they are covered by the red fit lines.

the main Gaussian response. After exploring many functional forms, the pile-up was most stably modeled with a sum of two independent Woods-Saxon distributions multiplied by exponential decays. One of these (blue) would model the wider but decaying right tail of the Gaussian while the other (purple) would model the rather flat saturation of pile-up events. An example of one of these fits is shown on the right of Figure 5.33.

$$\text{gaus}(0) + [3] * \exp(-[4] * (x - [5])) / (1 + \exp(-(x - [5]) / [6])) \\ + [7] * \exp(-[8] * (x - [9])) / (1 + \exp(-(x - [9]) / [10]))$$

Each of the horizontal slices on the y-axis are fit with this function, starting from (20, 25) and continuing to the maximum slice which still has a converging fit. From these, the large Gaussian component is extracted. From the mean of this Gaussian, the x-axis location of $n\sigma$ from the mean is calculated and stored (3σ shown as vertical yellow lines in Figure 5.33). These points (both the left and right sets) are plotted on the rotated 2D histogram and then

rotated back to the original BTOFMatch vs Refmult. These points are shown in green in Figure 5.34. Each set (upper and lower) are fit to a fourth order polynomial, plotted in red. These red curves now define the cut boundaries in BTOFMatch vs refmult space. In each event, the BTOFMatch and refmult value are compared to these red curves. Any events outside of these curves are rejected as pile-up.

For a systematic check, the $n\sigma$ values of the (green) points that define the cut lines in Figure 5.34 are varied from the default of 3 to 2 (tight), 1.5 (very tight), 4 (loose), and 5 (very loose).

5.7 STAR Systematics

Nine types of systematic studies were conducted on the STAR data set and used to obtain systematic uncertainty estimates for this analysis. Each of these studies involves a variation of a cut used in the analysis and two or three variations were performed for each of these nine studies. The bottom panel of Figure 5.35 shows the measured $\langle\Delta\sigma^2\rangle$ for all of these variations for the most central 0-5% events of 7.7 GeV with 120° azimuthal partitions. The default value is shown on the far left and is also represented by the horizontal black line. The solid error bars on this point represent the statistical uncertainty while the transparent error band represents the combined systematic uncertainty from all of the sources shown.

$\langle\Delta\sigma^2\rangle$ for each systematic variation is then plotted for this data set, grouped by color according to the type of variation. The error bars represent the statistical uncertainty on the systematic variation's $\langle\Delta\sigma^2\rangle$. From these systematic variations, the total systematic uncertainty is calculated in the spirit of Rodger Barlow's oft-quoted paper [46]. For a given systematic variation, we first calculate the difference between the default value (d) and this variation value (v): $\delta = d - v$. We propagate uncertainties according to Barlow's paper, which assumes that the variation data set is simply a subset of the default data set (or vice-versa) and the error scales as $1/\sqrt{N}$ where N is the number of events. The squared uncertainty on the difference δ is then actually the *difference* of the squares of the default uncertainty (σ_d) and the variation uncertainty σ_v , since the two are highly correlated: $\sigma_\delta = \sqrt{\sigma_d^2 - \sigma_v^2}$.

If this difference δ is highly consistent with zero, the systematic variation had little effect on the measured $\langle\Delta\sigma^2\rangle$ value and we say that the systematic check passed the Barlow test. It is only systematic variations that fail the Barlow test that must be accounted for in the final systematic uncertainty. We must then arbitrarily choose a threshold at which a systematic check passes or fails the Barlow test. We choose the commonly used (though we note still arbitrary) 1σ deviation of δ from zero. The contribution to the total systematic uncertainty from a single source is then given by Equation 5.11 which uses the value of




Figure 5.35: Visualization of the systematic checks made for the 0-5% centrality class in 7.7 GeV data using 120° partitions. On the bottom panel, the default $\langle \Delta\sigma^2 \rangle$ value is shown on the left, with the solid error bars representing the statistical uncertainty and the transparent band representing the total systematic uncertainty. The $\langle \Delta\sigma^2 \rangle$ value of each systematic variation is plotted. Each type of systematic check is grouped together and displayed in the same color. The top panel shows the contributions of each systematic source to the total systematic uncertainty, plotted on the left in black with the statistical uncertainty represented as a red line for comparison.

δ as the contribution but subtracts σ_δ in quadrature in order to account for the expected variation due to statistical fluctuations. These contributions (σ_{sys_i}) are added in quadrature from each systematic variation to produce the final systematic uncertainty on $\langle \Delta\sigma^2 \rangle$.

$$\sigma_{\text{sys}_i} = \sqrt{(d - v)^2 - |\sigma_d^2 - \sigma_v^2|} \quad (5.11)$$

In Figure 5.35 we run two or three variations for each systematic. The number of variations we run are arbitrary, and allowing them all to contribute to the total systematic uncertainty would produce an arbitrarily large systematic uncertainty, depending upon how many checks were made. Instead, we choose the systematic variation with the largest contribution according to Equation 5.11 for each type of systematic check and add only this contribution to the total.




Figure 5.36: Visualization of the baseline values from the quadratic fit to $\langle \Delta\sigma^2 \rangle$ vs partition width for the 0-5% centrality class at 7.7 GeV.

$$\sigma_{\text{sys}_{\text{total}}} = \sqrt{\sum_i \sigma_{\text{sys}_i}^2} \quad (5.12)$$

The top panel of Figure 5.35 shows the contribution to the total systematic uncertainty, represented by the black bar on the far left, from each source. The red line on the total uncertainty represents the statistical uncertainty, for comparison.

A brief description of each of the systematic checks is provided below:

- v_z : The analysis is performed on events within three sub-volumes of the primary vertex v_z distribution.
- efficiency: 5% and 10% of the proton tracks are randomly discarded in each event.
- dca: The maximum dca cut is varied from the nominal 1 cm.
- $n\sigma$ proton: The maximum $|n\sigma_{\text{proton}}|$ cut is varied from the nominal 2.

- m2 range: The TOF m^2 range is varied from the nominal 0.6 GeV. The range remains centered on $m^2 = 0.9$ GeV.
- nhits fit: The minimum number of TPC hits fit is varied from the nominal 20.
- ref3 shift: The refmult3 reference multiplicity distribution which is used to define the centrality is shifted by plus or minus one unit. This checks the effect of shifting events between centrality bins.
- dcaxy qa: The threshold values for marking an event bad is shifted by $\pm 1\sigma$ for each of the moving averages, corresponding to the “loose” and “tight” variations in Table 5.3.
- pile-up qa: The $n\sigma$ values used to define the cut boundaries in Section 5.6 are shifted by $\pm 1\sigma$, corresponding to the “loose” and “tight” variations described in the Section.

Figure 5.36 shows the systematic uncertainties in the baseline of the $\langle \Delta\sigma^2 \rangle$ vs partition width for the 0-5% centrality class at 7.7 GeV. Using all partition widths in the fit effectively improves the statistical resolution though the systematics are similar to Figure 5.35 for 120° .

Figure 5.37 lists the contributions of each source to the total uncertainty of $\langle \Delta\sigma^2 \rangle$ for all data sets with 120° azimuthal partitions. Figure 5.36 shows these contributions for the baseline of the $\langle \Delta\sigma^2 \rangle$ vs partition width fit.

120° Partitions

	0-5% Centrality						5-10% Centrality						10-20% Centrality						20-30% Centrality					
default	11%	12%	5%	5%	7%	7%	10%	5%	6%	5%	5%	8%	7%	7%	6%	4%	3%	4%	4%	6%	1%	5%	3%	4%
vz range	88%	87%	16%	10%	12%	50%	89%	18%	78%	23%	52%	32%	72%	89%	67%	76%	19%	7%	59%	75%	0%	80%	35%	45%
efficiency	1%	0%	0%	7%	1%	13%	3%	7%	6%	1%	7%	15%	0%	0%	0%	0%	0%	18%	2%	2%	17%	1%	1%	1%
dca	1%	2%	31%	13%	12%	4%	2%	14%	2%	32%	12%	21%	12%	8%	21%	16%	62%	39%	2%	5%	39%	14%	53%	35%
nσ proton	0%	2%	2%	2%	0%	1%	0%	2%	0%	1%	0%	1%	0%	1%	1%	0%	1%	1%	3%	1%	0%	1%	1%	6%
m ² range	0%	1%	0%	1%	0%	0%	0%	0%	0%	0%	0%	3%	0%	0%	0%	2%	0%	8%	2%	0%	5%	0%	1%	8%
nHits fit	3%	5%	2%	11%	70%	16%	4%	24%	6%	6%	24%	7%	2%	0%	2%	1%	14%	0%	14%	3%	14%	0%	0%	1%
refmult3 shift	3%	0%	24%	15%	2%	6%	0%	26%	6%	21%	1%	4%	7%	1%	5%	1%	2%	2%	3%	5%	8%	0%	2%	0%
dcaxy qa	1%	2%	22%	11%	1%	8%	1%	0%	0%	14%	1%	7%	4%	1%	2%	2%	1%	13%	3%	2%	4%	0%	2%	4%
pile-up qa	2%	0%	2%	30%	2%	3%	1%	9%	1%	1%	3%	10%	2%	0%	2%	2%	1%	12%	13%	8%	12%	3%	5%	1%
	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4

120° Partitions

	30-40% Centrality						40-50% Centrality						50-60% Centrality						60-70% Centrality					
default	8%	5%	7%	5%	4%	4%	3%	11%	5%	7%	4%	3%	13%	18%	9%	11%	7%	16%	6%	12%	14%	17%	6%	27%
vz range	65%	30%	56%	59%	69%	6%	23%	80%	30%	41%	24%	35%	47%	80%	38%	70%	83%	53%	5%	23%	21%	37%	10%	35%
efficiency	20%	14%	1%	0%	2%	1%	2%	3%	15%	2%	1%	3%	1%	9%	2%	3%	0%	8%	1%	33%	2%	29%	18%	17%
dca	0%	8%	3%	21%	25%	74%	5%	0%	2%	12%	48%	4%	2%	1%	0%	8%	4%	0%	0%	0%	1%	3%	23%	5%
nσ proton	0%	0%	0%	0%	1%	5%	8%	1%	4%	2%	0%	10%	2%	1%	0%	3%	2%	1%	8%	4%	1%	0%	8%	0%
m ² range	0%	0%	2%	0%	0%	3%	1%	0%	1%	1%	0%	4%	0%	0%	0%	0%	1%	0%	0%	0%	0%	0%	0%	0%
nHits fit	1%	0%	2%	0%	1%	6%	0%	3%	9%	1%	1%	15%	2%	2%	3%	2%	1%	1%	9%	2%	5%	1%	1%	1%
refmult3 shift	7%	30%	10%	11%	2%	3%	6%	9%	3%	19%	14%	3%	23%	5%	37%	8%	3%	13%	35%	11%	37%	13%	25%	16%
dcaxy qa	4%	6%	19%	1%	0%	0%	7%	3%	16%	12%	5%	6%	8%	1%	5%	3%	3%	7%	26%	14%	21%	5%	7%	16%
pile-up qa	3%	11%	6%	6%	0%	1%	48%	1%	20%	10%	7%	20%	15%	1%	16%	4%	3%	15%	15%	13%	12%	12%	8%	11%
	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4

Figure 5.37: Table showing the contribution of each systematic source to the total systematic uncertainty of $\langle \Delta\sigma^2 \rangle$ at 120° for all data sets. The top row (default) shows the size of the systematic uncertainty compared to the magnitude of $\langle \Delta\sigma^2 \rangle$. Each of the succeeding rows represents a source of systematic uncertainty and its contribution to the total systematic uncertainty. Each column (excluding the default row) will sum to 100% within rounding error.

	0-5% Centrality						5-10% Centrality						10-20% Centrality						20-30% Centrality					
default	11%	14%	5%	5%	7%	7%	9%	6%	6%	5%	6%	8%	7%	7%	6%	4%	4%	4%	4%	6%	2%	5%	2%	5%
vz range	89%	90%	25%	25%	6%	59%	87%	37%	70%	21%	63%	31%	73%	90%	71%	75%	25%	9%	66%	81%	15%	74%	36%	67%
efficiency	1%	1%	3%	4%	1%	11%	4%	12%	4%	2%	3%	11%	0%	0%	3%	1%	0%	13%	2%	2%	32%	3%	1%	0%
dca	1%	1%	19%	12%	12%	6%	2%	10%	3%	30%	9%	18%	10%	4%	15%	17%	56%	45%	0%	10%	18%	13%	55%	16%
n σ proton	0%	2%	4%	2%	0%	2%	0%	2%	1%	1%	0%	1%	0%	1%	1%	0%	1%	2%	3%	0%	0%	1%	1%	2%
m ² range	0%	0%	0%	0%	0%	0%	0%	0%	0%	1%	0%	4%	0%	0%	0%	2%	0%	6%	1%	0%	5%	0%	1%	5%
nHits fit	6%	4%	3%	15%	75%	15%	3%	22%	9%	7%	22%	10%	4%	1%	1%	1%	15%	1%	10%	3%	10%	0%	0%	1%
refmult3 shift	1%	1%	22%	10%	2%	1%	0%	9%	12%	18%	1%	6%	9%	1%	5%	0%	1%	2%	3%	2%	14%	2%	3%	2%
dcaxy qa	0%	1%	23%	12%	1%	5%	2%	0%	1%	15%	0%	10%	2%	2%	3%	3%	1%	12%	3%	0%	7%	1%	0%	4%
pile-up qa	2%	1%	1%	19%	2%	0%	1%	8%	0%	5%	2%	9%	1%	1%	1%	2%	0%	9%	11%	1%	0%	4%	2%	2%
	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4
	30-40% Centrality						40-50% Centrality						50-60% Centrality						60-70% Centrality					
default	7%	4%	7%	4%	4%	4%	4%	13%	3%	7%	5%	4%	13%	21%	9%	12%	6%	16%	7%	11%	14%	17%	6%	25%
vz range	58%	24%	67%	59%	57%	13%	50%	94%	18%	49%	26%	36%	46%	86%	30%	71%	85%	54%	4%	26%	26%	43%	16%	40%
efficiency	26%	20%	7%	1%	4%	0%	0%	1%	7%	3%	0%	0%	1%	7%	12%	6%	1%	7%	22%	34%	0%	16%	13%	12%
dca	0%	14%	3%	25%	33%	65%	3%	0%	4%	10%	34%	3%	1%	2%	0%	7%	5%	0%	1%	0%	0%	2%	23%	6%
n σ proton	0%	0%	1%	0%	1%	5%	2%	0%	7%	2%	1%	11%	3%	1%	0%	3%	1%	1%	8%	3%	0%	0%	8%	1%
m ² range	0%	1%	2%	1%	1%	4%	1%	0%	2%	1%	0%	4%	0%	0%	0%	0%	1%	0%	1%	0%	0%	0%	0%	0%
nHits fit	1%	1%	2%	2%	1%	6%	1%	1%	20%	1%	1%	21%	3%	1%	2%	2%	1%	1%	6%	2%	6%	1%	1%	1%
refmult3 shift	11%	19%	5%	6%	2%	2%	14%	3%	17%	16%	16%	7%	24%	3%	36%	7%	3%	13%	41%	12%	37%	15%	30%	15%
dcaxy qa	1%	10%	13%	0%	0%	1%	4%	0%	18%	12%	8%	3%	9%	0%	3%	3%	2%	9%	2%	11%	19%	8%	3%	12%
pile-up qa	2%	12%	1%	5%	0%	4%	25%	0%	7%	7%	13%	15%	13%	0%	17%	1%	1%	14%	15%	11%	11%	15%	6%	14%
	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4	7.7	11.5	19.6	27	39	62.4

Figure 5.38: Table showing the contribution of each systematic source to the total systematic uncertainty of the $\langle \Delta\sigma^2 \rangle$ vs partition width fit baseline for all data sets.

CHAPTER 6

Conclusion and Outlook

Mapping the phase diagram of QCD is one of the key goals of experimental nuclear physics in the twenty-first century. RHIC has invested five years into two beam energy scans to search for any signal of the conjectured critical point. To this end, STAR has focused on higher moments of the net-proton multiplicity distribution at mid-rapidity. A tantalizing non-monotonicity as a function of energy was observed in the ratio of kurtosis to variance in these distributions, consistent with qualitative theoretical expectations when crossing the critical point. This trend will be revisited in the near future with higher statistics from the second beam energy scan.

Observing a first-order phase transition at STAR would be sufficient to infer the existence of a critical point in the phase diagram, as this point is simply the boundary between a cross-over and first-order transition. The formation of bubbles is a general characteristic of first-order phase transitions. In the QGP to hadron gas transition, bubbles of high baryon density may form in the cooling QGP volume and produce attractive correlation between the momenta of final state baryons. It is possible that this correlation could be detected in the azimuthal proton distributions in STAR data.

We have constructed an observable to measure this correlation via proton multiplicity distributions within azimuthal partitions in each event. The widths of these distributions are compared to the width of a binomial distribution that would be produced if the protons were uncorrelated. We define the $\Delta\sigma^2$ variable to quantify the properly normalized deviation of the variance from binomial and study methods for subtracting efficiency and flow related

backgrounds.

The $\Delta\sigma^2$ observable was measured in STAR BES-I data and a strong repulsive interaction between protons was observed. The magnitude of this repulsion increases rapidly as the event multiplicity decreases. A similar effect was found in the AMPT model and is likely a byproduct of global momentum conservation. This event multiplicity dependence of $\Delta\sigma^2$ in AMPT was found to be energy independent, while in STAR data the magnitude of repulsion decreased with decreasing colliding beam energy. It is possible that this energy dependence, which was not replicated by AMPT, is the result of a positive correlation between proton tracks superimposed on an energy independent repulsive momentum conservation background. In this scenario, it appears that the attractive correlation would be small at higher energies and increase as energy decreased, consistent with naive expectation of a clustering phenomena turning on at the onset of a first-order phase transition. Without a reliable understanding of the event multiplicity dependent background, however, we are unable to make any conclusion regarding the physical nature of the possible attractive correlation with the STAR measurements.

It may be possible to model or even correct the effect of global momentum conservation on this observable, and we hope to pique the interest of phenomenology in helping us to understand this background. From there, it may be possible to measure any residual correlation between proton tracks to determine if any signs of a first-order transition are present at RHIC energies.

APPENDIX A

Simulations

A collection of simple simulations were constructed in order to help test the response of our observable to various types of azimuthal correlation. These simulations can be divided into two categories: those modeling signal and those modeling background. Our goal is to measure the correlation between proton tracks in search of clustering phenomena associated with a first-order phase transition. Experimental data, however, is inevitably an admixture of multiple effects and much of the work of experimental physics is trying to separate these mixed contributions in an observable. A powerful tool for approaching this problem is the simulation of the known effects. If these effects can be faithfully represented in a simulation, the observable can be tested on the simulated data to help us understand its dependence on this effect alone. Further, it may be possible to combine these simulations to better understand how the effects interact and their combined effect on the observable. Through this reductionist approach, it may be possible to build from simple simulations a better understanding of the measurements made on experimental data.

This Appendix will outline the algorithms and procedures for generating each simulation used in the analysis. Background effects such as hydrodynamic flow and detector inefficiency are briefly described first, followed by a discussion of signal models of Gaussian clustering.

A.1 Rejection Sampling

Rejection sampling is one method of drawing random samples from a specific probability distribution. Two random numbers are needed per sample. The first, x , should be sampled from a distribution uniform within the support of the probability distribution. The second number, y , should be sampled from a uniform distribution between zero and the maximum of the probability distribution (the pdf here can have arbitrary normalization, as long as the maximum is known). If y is greater than the value of the pdf at x , this sample is rejected. If y is less than or equal to the value of the pdf at x , this sample is kept with x as the value.

This method is repeated until the desired number of successful samples is obtained. This method is favored for its simplicity and versatility, as any well-defined probability distribution can be sampled. In addition, the rejection can be added in series with any track generation model, making this sampling ideal for combining with other simulations or even model or experimental data. Rejection sampling can be computationally inefficient if the maximum value is much larger than other values in the distribution.

A.2 Homogeneous Efficiency

Detector inefficiency which is independent of any track characteristics (momentum, azimuthal angle, rapidity, etc) is rather easy to simulate. Given any input tracks, a homogeneous inefficiency simply removes the track from the analysis with probability $p \in [0, 1]$ via rejection sampling (Section A.1). This removes $p \cdot 100\%$ of the generated tracks on average which produces a data set which is said to have efficiency of $(1 - p) \cdot 100\%$.

A.3 Azimuthal Dependent Efficiency

ϕ dependent efficiency is implemented by rejection sampling. In this algorithm, the efficiency pdf is re-normalized such that the max value is 1 (the area under the pdf is no longer 1 in

this case). A track is randomly generated with azimuthal angle $\phi \in [0, 2\pi)$. The value of the renormalized efficiency pdf at this ϕ value is retrieved and will be between 0 and 1. Next a random number in $[0, 1)$ is generated. If the random number is less than the efficiency value at ϕ , the track is kept. If the random number is greater than the efficiency value, the track is rejected. This process is repeated until the desired N tracks have been generated.

The efficiency pdf will not be rotated between events, as its azimuthal orientation is fixed by the detector. This should allow mixed events to capture correlations due to efficiency while washing out signal correlations, whose orientations have no correlation between events.

A.4 Flow

Hydrodynamic flow is simulated by rejection sampling (Section A.1) using the Fourier decomposition of Equation A.1. Tracks are generated with azimuthal angle $\phi \in [0, 2\pi)$ and rejected based on the pdf which is characterized by the flow coefficients v_n . Typically, only elliptic flow is simulated, in which case v_2 is finite while all other flow coefficients are zero. The event plane angle ψ_R is randomly oriented in each generated event, emulating the random orientation of the event plane in heavy ion collisions.

$$\frac{dN}{d\phi} = \frac{1}{2\pi} \left[1 + \sum_{n=1}^{\infty} 2v_n \cos[n(\phi - \psi_R)] \right] \quad (\text{A.1})$$

A.5 Simple Clustering Model

While the Gaussian correlation model (Section 5.2) can generate repulsive correlations as well as attractive, it does so at the expense of a dynamic azimuthal probability distribution. Due to this dynamic pdf, it can not be treated analytically as described in Section 4.3. In addition, it will couple to additional correlation models when combined, as described in Section 5.2. For this reason, a simple Gaussian clustering model has proven to be a useful

tool for comparison.

This model is simply a static Gaussian probability distribution superimposed on a flat background:

$$p(\phi) = C \left(1 + Ae^{-\frac{1}{2}\left(\frac{\phi-\mu}{\sigma}\right)^2} \right) \quad (\text{A.2})$$

Here C is a normalizing constant such that the integral of the pdf over the azimuthal range of $[0, 2\pi)$ is 1. A is the relative amplitude of the Gaussian cluster to the flat background. When $A = 0$, the pdf reduces to a uniform distribution and produces randomly distributed tracks. For a finite A , tracks have a larger probability of being produced close to the Gaussian center, μ . The σ parameter determines how wide this increased probability is. A larger σ not only increases the width of the distribution but also decreases its height (due to the normalization), meaning that the produced clustering is less sharp.

When sampling this distribution for event-by-event simulations, each track is sampled from either a Gaussian distribution or a flat distribution. Each event has a randomly chosen μ for the Gaussian distribution along with a fixed σ . In addition, a $p \in [0, 1]$ parameter characterizes the probability which the Gaussian is sampled rather than the uniform distribution. If a random number $r \in [0, 1)$ is less than p the Gaussian is sampled and its values are wrapped into a $[0, 2\pi)$ range. Otherwise a flat distribution is sampled. This produces a slightly different parameterization from Equation A.2:

$$p(\phi) = \frac{1-p}{2\pi} + \frac{p}{\sigma\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{\phi-\mu}{\sigma}\right)^2} \quad (\text{A.3})$$

Equation A.3 can be rewritten to match the form of Equation A.2 and relate the p and A parameters:

$$p(\phi) = \frac{1-p}{2\pi} \left(1 + \frac{\sqrt{2\pi}}{\sigma} \frac{p}{1-p} e^{-\frac{1}{2}\left(\frac{\phi-\mu}{\sigma}\right)^2} \right) \quad (\text{A.4})$$

$$A = \frac{\sqrt{2\pi}}{\sigma} \frac{p}{1-p} \quad (\text{A.5})$$

$$p = \frac{A}{A + \frac{\sqrt{2\pi}}{\sigma}} \quad (\text{A.6})$$

$$C = \frac{1-p}{2\pi} \quad (\text{A.7})$$

A.6 Momentum Conservation Model

The global conservation of momentum in each event is expected to produce a negative correlation between particles. This is most easily understood in the extreme case of only two particles being produced. In this case, the two particles will be produced back to back, enforcing a perfect anti-correlation of their momenta vectors. As more particles are produced in the final state, each is less responsible for conserving the total momentum. As the number of particles becomes large, the anti-correlation between them should become smaller and approach zero as the number of particles is increased. While the behavior is understood in these extreme limits, the exact functional form of this effect is not known.

To test the effect of global momentum conservation on the $\Delta\sigma^2$ observable, we have constructed an extremely simple model of pure momentum conservation. M particles are generated in each event. Each is assigned a random 3-momentum in which each component is sampled from a uniform distribution of $[0, 2]$ GeV¹. The net momentum vector of the system is calculated as a sum of the individual momenta. Each particle is then slightly rotated away from this net momentum vector by an amount proportional to the particle's momentum magnitude and the event's net momentum magnitude. After all particles are slightly rotated, the event's net momentum is recalculated and the procedure is repeated.

¹This will not produce a uniform distribution on a sphere. In future investigations it would likely be better to generate three vectors in spherical coordinates.




Figure A.1: $\langle \Delta\sigma^2 \rangle$ is calculated for events generated with M total particles (unfortunately labeled as N in the legend) and plotted against M . The data is fit to power law in M plus a constant. The fit is plotted and the fit parameters are displayed.

Since the rotation is proportional to the event net momentum magnitude, the rotations should become smaller with iteration and the event net momentum should converge to zero. In practice, ensuring convergence requires some fine tuning of the rotation amount but generally works. The iteration stops once the event's net momentum is smaller than some predetermined value.

We then utilize this model to test the effect momentum conservation may have on the $\Delta\sigma^2$ observable. 250000 events of zero (very small) net momentum are generated for each value of M total particles. 40% of these particles are randomly labeled protons. The protons roughly within STAR acceptance are selected: $|y| < 0.5$, $p_T < 2$ GeV, $p < 2$ GeV. This produces a distribution of proton multiplicity per event for each fixed total number of particles, M . $\langle \Delta\sigma^2 \rangle$ is calculated as usual for these selected protons and is plotted in Figure A.1 as a function of the total number of particles in the event.

We find a significantly negative $\langle \Delta\sigma^2 \rangle$ for all points simulated which becomes more

negative as the total particles per event decreases. This trend is fit to a power law plus a constant and it is found that the best fit is consistent with a $1/M$ dependence with no offset. This trend is similar to that found in AMPT and STAR data and the simple functional dependence may indicate that the effect of momentum conservation on the $\langle\Delta\sigma^2\rangle$ observable is analytically tractable and that a correction could possibly be devised.

While the results are enticing, we acknowledge that the momentum conservation model outlined here is extremely simplistic but at the same time has quite a few parameters that may influence the $\langle\Delta\sigma^2\rangle$ values. Specifically, the percentage of protons per event and the distribution and magnitude of particle momenta. Sensitivity to these variables must be studied before drawing any firm conclusions from this model. In addition, much physics is missing from the model, though building a realistic model of heavy ion collisions is not the intention of this study.

APPENDIX B

Models

Two models were compared to STAR data in this analysis: AMPT and the MUSIC+FIST model.

B.1 AMPT

A multi-phase transport (AMPT) model [51] was developed to simulate the microscopic interactions in relativistic heavy ion collisions. The model begins with initial conditions from the HIJING model [52], [53], which produces hadrons, strings and minijet partons from initial scattering. In the original model, interactions between the minijet partons are described by Zhang's parton cascade (ZPC) [54] until parton freezeout, at which point both these partons and the excited strings hadronize according to the Lund string fragmentation model. Scattering amongst these hadrons is then described by a relativistic transport (ART) model.

To describe the large observed collective flow, the string melting version of AMPT [55] was developed. In this model, the initial state hadrons and minijets are converted into partonic degrees of freedom, whose interactions are described by the ZPC model. At partonic freezeout, a quark coalescence model describes hadronization before the ART model is used to describe the final hadronic scattering. A recently updated AMPT model with an improved quark coalescence algorithm [56] was used in this analysis.

B.2 MUSIC+FIST

The MUSIC+FIST model is a combination of the MUSIC hydrodynamic model [57]–[61] and the FIST sampler [62]. Initial conditions from the Glauber model [40] undergo hydrodynamic evolution in MUSIC. The simulated QGP volume expands until reaching a “switching” energy density, $\varepsilon_{\text{SW}} = 0.26 \text{ GeV}/\text{fm}^3$. This density is reached at some specific time for each hydrodynamic volume element of the expanding system, producing a particlization hypersurface in both space and time. The FIST sampler then samples final state particles from this hypersurface. This model pre-determines the number of each particle species to sample in each event, which ensures global conservation of conserved charges, at the cost of not conserving momentum.

An excluded volume effect was added to the original FIST model [63]. When sampling the hypersurface, the excluded volume model ensures that no two baryons are sampled within a volume of radius $r = 1 \text{ fm}$. This radius is a parameter which can be adjusted.

Bibliography

- [1] S. Weinberg, “A model of leptons,” *Phys. Rev. Lett.*, vol. 19, pp. 1264–1266, 21 Nov. 1967. DOI: [10.1103/PhysRevLett.19.1264](https://doi.org/10.1103/PhysRevLett.19.1264).
- [2] G. Aad, T. Abajyan, B. Abbott, *et al.*, “Observation of a new particle in the search for the standard model higgs boson with the atlas detector at the lhc,” *Physics Letters B*, vol. 716, no. 1, pp. 1–29, 2012. DOI: <https://doi.org/10.1016/j.physletb.2012.08.020>.
- [3] S. Chatrchyan, V. Khachatryan, A. Sirunyan, *et al.*, “Observation of a new boson at a mass of 125 gev with the cms experiment at the lhc,” *Physics Letters B*, vol. 716, no. 1, pp. 30–61, 2012. DOI: <https://doi.org/10.1016/j.physletb.2012.08.021>.
- [4] A. Takenaka, K. Abe, C. Bronner, *et al.*, “Search for proton decay via $p \rightarrow e^+\pi^0$ and $p \rightarrow \mu^+\pi^0$ with an enlarged fiducial volume in super-kamiokande i-iv,” *Phys. Rev. D*, vol. 102, p. 112011, 11 Dec. 2020. DOI: [10.1103/PhysRevD.102.112011](https://doi.org/10.1103/PhysRevD.102.112011).
- [5] B. P. Abbott, R. Abbott, T. D. Abbott, *et al.*, “Observation of gravitational waves from a binary black hole merger,” *Phys. Rev. Lett.*, vol. 116, p. 061102, 6 Feb. 2016. DOI: [10.1103/PhysRevLett.116.061102](https://doi.org/10.1103/PhysRevLett.116.061102).
- [6] R. L. Workman *et al.*, “Review of Particle Physics,” *PTEP*, vol. 2022, p. 083C01, 2022. DOI: [10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).
- [7] H. D. Politzer, “Reliable perturbative results for strong interactions?” *Phys. Rev. Lett.*, vol. 30, pp. 1346–1349, 26 Jun. 1973. DOI: [10.1103/PhysRevLett.30.1346](https://doi.org/10.1103/PhysRevLett.30.1346).
- [8] D. J. Gross and F. Wilczek, “Ultraviolet behavior of non-abelian gauge theories,” *Phys. Rev. Lett.*, vol. 30, pp. 1343–1346, 26 Jun. 1973. DOI: [10.1103/PhysRevLett.30.1343](https://doi.org/10.1103/PhysRevLett.30.1343).
- [9] C. Davies, *Lattice qcd - a guide for people who want results*, 2005. arXiv: [hep-lat/0509046](https://arxiv.org/abs/hep-lat/0509046) (hep-lat).

- [10] C. Ratti and R. Bellwied, “Introduction to lattice qcd,” in *The Deconfinement Transition of QCD: Theory Meets Experiment*. Cham: Springer International Publishing, 2021, pp. 3–22. DOI: [10.1007/978-3-030-67235-5_1](https://doi.org/10.1007/978-3-030-67235-5_1).
- [11] K. Aamodt, B. Abelev, A. Abrahantes Quintana, *et al.*, “Higher harmonic anisotropic flow measurements of charged particles in pb-pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV,” *Phys. Rev. Lett.*, vol. 107, p. 032301, 3 Jul. 2011. DOI: [10.1103/PhysRevLett.107.032301](https://doi.org/10.1103/PhysRevLett.107.032301).
- [12] H. Niemi, K. J. Eskola, and R. Paatelainen, “Event-by-event fluctuations in a perturbative qcd + saturation + hydrodynamics model: Determining qcd matter shear viscosity in ultrarelativistic heavy-ion collisions,” *Phys. Rev. C*, vol. 93, p. 024907, 2 Feb. 2016. DOI: [10.1103/PhysRevC.93.024907](https://doi.org/10.1103/PhysRevC.93.024907).
- [13] U. Heinz and R. Snellings, “Collective flow and viscosity in relativistic heavy-ion collisions,” *Annual Review of Nuclear and Particle Science*, vol. 63, no. 1, pp. 123–151, 2013. DOI: [10.1146/annurev-nucl-102212-170540](https://doi.org/10.1146/annurev-nucl-102212-170540).
- [14] A. Bzdak, S. Esumi, V. Koch, J. Liao, M. Stephanov, and N. Xu, “Mapping the phases of quantum chromodynamics with beam energy scan,” *Physics Reports*, vol. 853, pp. 1–87, 2020, Mapping the phases of quantum chromodynamics with beam energy scan. DOI: <https://doi.org/10.1016/j.physrep.2020.01.005>.
- [15] S. Blundell and K. M. Blundell, *Concepts in thermal physics*. Oxford University Press, 2008.
- [16] M. A. Halasz, A. D. Jackson, R. E. Shrock, M. A. Stephanov, and J. J. M. Verbaarschot, “Phase diagram of qcd,” *Phys. Rev. D*, vol. 58, p. 096007, 9 Sep. 1998. DOI: [10.1103/PhysRevD.58.096007](https://doi.org/10.1103/PhysRevD.58.096007).
- [17] J. Berges and K. Rajagopal, “Color superconductivity and chiral symmetry restoration at non-zero baryon density and temperature,” *Nuclear Physics B*, vol. 538, no. 1, pp. 215–232, 1999. DOI: [https://doi.org/10.1016/S0550-3213\(98\)00620-8](https://doi.org/10.1016/S0550-3213(98)00620-8).

- [18] X. An, M. Bluhm, L. Du, *et al.*, “The best framework for the search for the qcd critical point and the chiral magnetic effect,” *Nuclear Physics A*, vol. 1017, p. 122 343, 2022. DOI: <https://doi.org/10.1016/j.nuclphysa.2021.122343>.
- [19] P. Parotto, M. Bluhm, D. Mroczek, *et al.*, “Qcd equation of state matched to lattice data and exhibiting a critical point singularity,” *Phys. Rev. C*, vol. 101, p. 034 901, 3 Mar. 2020. DOI: [10.1103/PhysRevC.101.034901](https://doi.org/10.1103/PhysRevC.101.034901).
- [20] S. Tang, S. Wu, and H. Song, “Dynamical critical fluctuations near the qcd critical point with hydrodynamic cooling rate,” *Phys. Rev. C*, vol. 108, p. 034 901, 3 Sep. 2023. DOI: [10.1103/PhysRevC.108.034901](https://doi.org/10.1103/PhysRevC.108.034901).
- [21] S. Gupta, X. Luo, B. Mohanty, H. G. Ritter, and N. Xu, “Scale for the phase diagram of quantum chromodynamics,” *Science*, vol. 332, no. 6037, pp. 1525–1528, 2011. DOI: [10.1126/science.1204621](https://doi.org/10.1126/science.1204621).
- [22] M. A. Stephanov, “Sign of kurtosis near the qcd critical point,” *Phys. Rev. Lett.*, vol. 107, p. 052 301, 5 Jul. 2011. DOI: [10.1103/PhysRevLett.107.052301](https://doi.org/10.1103/PhysRevLett.107.052301).
- [23] J. Steinheimer and J. Randrup, “Spinodal density enhancements in simulations of relativistic nuclear collisions,” *Phys. Rev. C*, vol. 87, p. 054 903, 5 May 2013. DOI: [10.1103/PhysRevC.87.054903](https://doi.org/10.1103/PhysRevC.87.054903).
- [24] I. Brown, “Vacuum Arc Ion Sources,” pp. 311–329, 2013, Comments: 19 pages, contribution to the CAS-CERN Accelerator School: Ion Sources, Senec, Slovakia, 29 May - 8 June 2012, edited by R. Bailey, CERN-2013-007. DOI: [10.5170/CERN-2013-007.311](https://doi.org/10.5170/CERN-2013-007.311).
- [25] J. ALESSI, D. BARTON, E. BEEBE, D. GASSNER, and E. AL., “Electron beam ion source preinjector project (ebis) conceptual design report.,” Feb. 2005. DOI: [10.2172/15016043](https://doi.org/10.2172/15016043).
- [26] G. Zschornacka, M. Schmidt, and A. Thorn, “Electron Beam Ion Sources,” pp. 165–201, 2013, Comments: 37 pages, contribution to the CAS-CERN Accelerator School:

- Ion Sources, Senec, Slovakia, 29 May - 8 June 2012, edited by R. Bailey. DOI: [10.5170/CERN-2013-007.165](https://doi.org/10.5170/CERN-2013-007.165).
- [27] M. Cornacchia, “Lattices for synchrotron radiation sources,” pp. 30–58, May 1994.
- [28] E. D. Courant, M. S. Livingston, and H. S. Snyder, “The strong-focusing synchrotron—a new high energy accelerator,” *Phys. Rev.*, vol. 88, pp. 1190–1196, 5 Dec. 1952. DOI: [10.1103/PhysRev.88.1190](https://doi.org/10.1103/PhysRev.88.1190).
- [29] B. N. Laboratory, “Booster Design Manual,” 1986.
- [30] B. N. Laboratory, “Alternating Gradient Synchrotron Project Construction Completion Report,” 1966.
- [31] B. N. Laboratory, “RHIC Design Manual,” 2006.
- [32] H. Hahn, E. Forsyth, H. Foelsche, *et al.*, “The rhic design overview,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 499, pp. 245–263, Mar. 2003. DOI: [10.1016/S0168-9002\(02\)01938-1](https://doi.org/10.1016/S0168-9002(02)01938-1).
- [33] W. Herr and B. Muratori, “Concept of luminosity,” 2006. DOI: [10.5170/CERN-2006-002.361](https://doi.org/10.5170/CERN-2006-002.361).
- [34] K. Ackermann, N. Adams, C. Adler, *et al.*, “Star detector overview,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 499, no. 2, pp. 624–632, 2003, The Relativistic Heavy Ion Collider Project: RHIC and its Detectors. DOI: [https://doi.org/10.1016/S0168-9002\(02\)01960-5](https://doi.org/10.1016/S0168-9002(02)01960-5).
- [35] M. Anderson, J. Berkovitz, W. Betts, *et al.*, “The star time projection chamber: A unique tool for studying high multiplicity events at rhic,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 499, no. 2, pp. 659–678, 2003, The Relativistic Heavy Ion

- Collider Project: RHIC and its Detectors. DOI: [https://doi.org/10.1016/S0168-9002\(02\)01964-2](https://doi.org/10.1016/S0168-9002(02)01964-2).
- [36] H. Bichsel, “A method to improve tracking and particle identification in tpcs and silicon detectors,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 562, no. 1, pp. 154–197, 2006. DOI: <https://doi.org/10.1016/j.nima.2006.03.009>.
- [37] B. Bonner, H. Chen, G. Eppley, *et al.*, “A single time-of-flight tray based on multi-gap resistive plate chambers for the star experiment at rhic,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 508, no. 1, pp. 181–184, 2003, Proceedings of the Sixth International Workshop on Resistive Plate Chambers and Related Detectors. DOI: [https://doi.org/10.1016/S0168-9002\(03\)01347-0](https://doi.org/10.1016/S0168-9002(03)01347-0).
- [38] J. Landgraf, M. LeVine, A. Ljubicic, J. Nelson, D. Padrazo, and M. Schulz, “An overview of the star daq system,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 499, no. 2, pp. 762–765, 2003, The Relativistic Heavy Ion Collider Project: RHIC and its Detectors. DOI: [https://doi.org/10.1016/S0168-9002\(02\)01973-3](https://doi.org/10.1016/S0168-9002(02)01973-3).
- [39] D. Smirnov, J. Lauret, V. Perevoztchikov, G. Van Buren, and J. Webb, “Vertex Reconstruction at STAR: Overview and Performance Evaluation,” *J. Phys. Conf. Ser.*, vol. 898, no. 4, R. Mount and C. Tull, Eds., p. 042058, 2017. DOI: [10.1088/1742-6596/898/4/042058](https://doi.org/10.1088/1742-6596/898/4/042058).
- [40] M. L. Miller, K. Reygers, S. J. Sanders, and P. Steinberg, “Glauber modeling in high-energy nuclear collisions,” *Annual Review of Nuclear and Particle Science*, vol. 57, no. 1, pp. 205–243, 2007. DOI: [10.1146/annurev.nucl.57.090506.123020](https://doi.org/10.1146/annurev.nucl.57.090506.123020).

- [41] X. Luo, “Error estimation for moment analysis in heavy-ion collision experiment,” *Journal of Physics G: Nuclear and Particle Physics*, vol. 39, no. 2, p. 025 008, Jan. 2012. DOI: [10.1088/0954-3899/39/2/025008](https://doi.org/10.1088/0954-3899/39/2/025008).
- [42] B. Efron, “Bootstrap Methods: Another Look at the Jackknife,” *The Annals of Statistics*, vol. 7, no. 1, pp. 1–26, 1979. DOI: [10.1214/aos/1176344552](https://doi.org/10.1214/aos/1176344552).
- [43] T. D. Tibbe and A. K. Montoya, “Correcting the bias correction for the bootstrap confidence interval in mediation analysis,” en, *Front Psychol*, vol. 13, p. 810 258, May 2022.
- [44] P. Hall, “Resampling a coverage pattern,” *Stochastic Processes and their Applications*, vol. 20, no. 2, pp. 231–246, 1985. DOI: [https://doi.org/10.1016/0304-4149\(85\)90212-1](https://doi.org/10.1016/0304-4149(85)90212-1).
- [45] N. Chamandy, O. Muralidharan, A. Najmi, and S. Naidu, “Estimating uncertainty for massive data streams,” Google, Tech. Rep., 2012.
- [46] R. Barlow, “Systematic errors: Facts and fictions,” Aug. 2002.
- [47] R. A. Fisher and J. H. Jeans, “The general sampling distribution of the multiple correlation coefficient,” *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, vol. 121, no. 788, pp. 654–673, 1928. DOI: [10.1098/rspa.1928.0224](https://doi.org/10.1098/rspa.1928.0224).
- [48] A. Bilandzic, R. Snellings, and S. Voloshin, “Flow analysis with cumulants: Direct calculations,” *Phys. Rev. C*, vol. 83, p. 044 913, 4 Apr. 2011. DOI: [10.1103/PhysRevC.83.044913](https://doi.org/10.1103/PhysRevC.83.044913).
- [49] A. M. Poskanzer and S. A. Voloshin, “Methods for analyzing anisotropic flow in relativistic nuclear collisions,” *Phys. Rev. C*, vol. 58, pp. 1671–1678, 3 Sep. 1998. DOI: [10.1103/PhysRevC.58.1671](https://doi.org/10.1103/PhysRevC.58.1671).

- [50] L. Adamczyk, G. Agakishiev, M. M. Aggarwal, *et al.*, “Inclusive charged hadron elliptic flow in au + au collisions at $\sqrt{s_{NN}} = 7.7\text{--}39$ gev,” *Phys. Rev. C*, vol. 86, p. 054908, 5 Nov. 2012. DOI: [10.1103/PhysRevC.86.054908](https://doi.org/10.1103/PhysRevC.86.054908).
- [51] Z.-W. Lin, C. M. Ko, B.-A. Li, B. Zhang, and S. Pal, “Multiphase transport model for relativistic heavy ion collisions,” *Phys. Rev. C*, vol. 72, p. 064901, 6 Dec. 2005. DOI: [10.1103/PhysRevC.72.064901](https://doi.org/10.1103/PhysRevC.72.064901).
- [52] X.-N. Wang and M. Gyulassy, “Hijing: A monte carlo model for multiple jet production in pp, pA, and AA collisions,” *Phys. Rev. D*, vol. 44, pp. 3501–3516, 11 Dec. 1991. DOI: [10.1103/PhysRevD.44.3501](https://doi.org/10.1103/PhysRevD.44.3501).
- [53] M. Gyulassy and X.-N. Wang, “Hijing 1.0: A monte carlo program for parton and particle production in high energy hadronic and nuclear collisions,” *Computer Physics Communications*, vol. 83, no. 2, pp. 307–331, 1994. DOI: [https://doi.org/10.1016/0010-4655\(94\)90057-4](https://doi.org/10.1016/0010-4655(94)90057-4).
- [54] B. Zhang, “Zpc 1.0.1: A parton cascade for ultrarelativistic heavy ion collisions,” *Computer Physics Communications*, vol. 109, no. 2, pp. 193–206, 1998. DOI: [https://doi.org/10.1016/S0010-4655\(98\)00010-1](https://doi.org/10.1016/S0010-4655(98)00010-1).
- [55] Z.-w. Lin and C. M. Ko, “Partonic effects on the elliptic flow at relativistic heavy ion collisions,” *Phys. Rev. C*, vol. 65, p. 034904, 3 Feb. 2002. DOI: [10.1103/PhysRevC.65.034904](https://doi.org/10.1103/PhysRevC.65.034904).
- [56] Y. He and Z.-W. Lin, “Improved quark coalescence for a multi-phase transport model,” *Phys. Rev. C*, vol. 96, p. 014910, 1 Jul. 2017. DOI: [10.1103/PhysRevC.96.014910](https://doi.org/10.1103/PhysRevC.96.014910).
- [57] B. Schenke, S. Jeon, and C. Gale, “(3+1)d hydrodynamic simulation of relativistic heavy-ion collisions,” *Phys. Rev. C*, vol. 82, p. 014903, 1 Jul. 2010. DOI: [10.1103/PhysRevC.82.014903](https://doi.org/10.1103/PhysRevC.82.014903).

- [58] B. Schenke, S. Jeon, and C. Gale, “Elliptic and triangular flow in event-by-event $D = 3 + 1$ viscous hydrodynamics,” *Phys. Rev. Lett.*, vol. 106, p. 042301, 4 Jan. 2011. DOI: [10.1103/PhysRevLett.106.042301](https://doi.org/10.1103/PhysRevLett.106.042301).
- [59] J.-F. G. Paquet, C. Shen, G. S. Denicol, *et al.*, “Production of photons in relativistic heavy-ion collisions,” *Phys. Rev. C*, vol. 93, p. 044906, 4 Apr. 2016. DOI: [10.1103/PhysRevC.93.044906](https://doi.org/10.1103/PhysRevC.93.044906).
- [60] G. S. Denicol, C. Gale, S. Jeon, A. Monnai, B. Schenke, and C. Shen, “Net-baryon diffusion in fluid-dynamic simulations of relativistic heavy-ion collisions,” *Phys. Rev. C*, vol. 98, p. 034916, 3 Sep. 2018. DOI: [10.1103/PhysRevC.98.034916](https://doi.org/10.1103/PhysRevC.98.034916).
- [61] P. Huovinen and H. Petersen, “Particlization in hybrid models,” *The European Physical Journal A*, vol. 48, no. 11, 2012. DOI: [10.1140/epja/i2012-12171-9](https://doi.org/10.1140/epja/i2012-12171-9).
- [62] V. Vovchenko, V. Koch, and C. Shen, “Proton number cumulants and correlation functions in au-au collisions at $\sqrt{s_{NN}} = 7.7\text{--}200$ gev from hydrodynamics,” *Phys. Rev. C*, vol. 105, p. 014904, 1 Jan. 2022. DOI: [10.1103/PhysRevC.105.014904](https://doi.org/10.1103/PhysRevC.105.014904).
- [63] V. Vovchenko, “Cooper-frye sampling with short-range repulsion,” *Phys. Rev. C*, vol. 106, p. 064906, 6 Dec. 2022. DOI: [10.1103/PhysRevC.106.064906](https://doi.org/10.1103/PhysRevC.106.064906).