

**Neutral Pion Double Helicity Asymmetry in Polarized
Proton-Proton Collisions at \sqrt{s} =200 GeV at STAR**

by

William Leight

B.A., Physics and Mathematics, Yale University (2004)

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

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2012

© Massachusetts Institute of Technology 2012. All rights reserved.

Author.....
Department of Physics
May 10, 2012

Certified by.....
Bernd Surrow
Associate Professor of Physics
Thesis Supervisor

Accepted by.....
Krishna Rajagopal
Associate Department Head for Education

Neutral Pion Double Helicity Asymmetry in Polarized Proton-Proton Collisions at $\sqrt{s}=200$ GeV at STAR

by

William Leight

Submitted to the Department of Physics
on May 10, 2012, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

One of the primary goals of the spin physics program at the STAR experiment is to constrain the polarized gluon distribution function, $\Delta_g(x, Q^2)$, by measuring the double longitudinal spin asymmetry, A_{LL} , of various final-state channels. Neutral pions provide a potentially powerful final state because they are copiously produced in p+p collisions and have few backgrounds.

In 2009, STAR took 14 pb^{-1} of integrated luminosity of 200 GeV p+p collisions, with average beam longitudinal polarization of 59%. Neutral pions produced in these collisions can be identified using STAR's large-acceptance electromagnetic calorimeter, with help from tracking from the STAR Time Projection Chamber. This work presents a measurement of the inclusive neutral pion A_{LL} from this data, based on a new π^0 reconstruction algorithm. A comparison to theoretical predictions and other experimental results suggests that the current best-fit value of ΔG , the gluon contribution to the proton spin, is too small and that ΔG is actually comparable in magnitude to the quark contribution to the proton spin $\Delta\Sigma$.

Thesis Supervisor: Bernd Surrow

Title: Associate Professor of Physics

Acknowledgments

First of all, thanks to those who directly contributed to this thesis, most importantly Bernd Surrow, who has been a constant source of support throughout my graduate career. Thanks are also due to (especially) Joe Seele, Jan Balewski, and Renee Fatemi, who answered my (occasionally stupid) questions and provided important help in developing my analysis. Thanks also to MIT RHIC Spin graduate students past and present, including Julie Millane, Alan Hoffman, Adam Kocoloski, Chris Jones, Jim Hays-Wehle, Ross Corliss, Matt Walker, and Mike Betancourt, who helped with both my analysis and my sanity. Thanks as well to the many non-MIT members of the STAR collaboration, especially those with the spin program.

Many friends, including several of the people listed above, were also essential in my getting my PhD: a short and almost certainly incomplete list would include Rahul, Arghavan, Sian, Ian, Greg, Joe, Seth, Jim, Parthi, Jean, Ross, Michael, all the GQ guys, and past and current roommates Madhu, Amanda, and Albert. I also would like to thank Ashdown House Housemasters Ann and Terry Orlando for fostering the kind of environment that made it so much fun to be a resident of Ashdown for seven years.

Finally, I'd like to thank my family: my siblings Elias, Ethan, and Jess, my future brother-in-law Amir, and above all my parents, without whose love, encouragement, support, and advice this thesis would almost certainly never have been written.

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

1	Introduction and Theoretical Background	17
1.1	Proton Structure and QCD	18
1.2	The Proton Spin	21
1.2.1	Quark Polarization and the “Spin Crisis”	22
1.3	The Gluon Polarization	24
1.3.1	Measuring ΔG at a Hadron Collider	25
1.3.2	The π^0	26
1.3.3	Note on Units	26
2	The RHIC Accelerator and the STAR Experiment	29
2.1	RHIC	29
2.1.1	Polarization	30
2.2	STAR	34
2.2.1	The TPC	35
2.2.2	The BBC and ZDC	38
2.2.3	The BEMC	39
2.2.4	The BSMD	43
3	π^0 Finding	47
3.1	Calibration	47

3.1.1	BSMD Calibration	47
3.1.2	BEMC Calibration	54
3.2	BSMD Clustering	57
3.3	π^0 Candidates	60
3.3.1	Decay Topologies	61
3.3.2	Ghost Candidates	66
3.3.3	General Cuts	68
4	Measuring A_{LL}	73
4.1	Event Selection	74
4.1.1	Run Selection	75
4.1.2	Triggers	76
4.2	Backgrounds and Fits	77
4.2.1	Combinatorial Background	78
4.2.2	Split-Photon Background	81
4.2.3	Simulating the π^0 peak	82
4.2.4	Fitting the Data to Simulation+Backgrounds	84
4.3	Relative Luminosity	85
4.4	Polarization	88
4.5	Single-Spin Asymmetries	88
4.6	Uncertainties	89
4.6.1	Statistical Uncertainties	90
4.6.2	Systematic Uncertainties	91
4.7	Results	99
5	Interpretation and Conclusion	103

List of Figures

1-1	Feynman diagram for deep inelastic scattering	19
1-2	Neutral pion production in a pp collision	21
1-3	Polarized gluon density at $Q^2=10$ GeV obtained from polarized DIS from de Florian, Navarro, and Sassot: the green band is the $\Delta\chi^2 = 1$ uncertainty, and the yellow the $\Delta\chi^2 = 2\%$ uncertainty.	25
1-4	At left, several predictions of the value of $x\Delta g(x, Q^2)$, plotted vs. x ; at right, the corresponding predictions for the $\pi^0 A_{LL}$	27
2-1	The RHIC Complex, with an emphasis on the pieces that are important for polarized operations.	30
2-2	The change in the spin vector of a proton as it passes through a full Siberian snake: the blue arrow shows the direction of the beam.	32
2-3	The RHIC CNI polarimeter setup, looking along the beam.	33
2-4	Schematic of the RHIC H-jet polarimeter, with the purple cylinder representing the hydrogen jet target.	33
2-5	The STAR experiment. Note that the ZDCs are too far from the interaction point to be visible in this diagram. The BSMD is also not visible as it is inside the BEMC. The second BBC is just visible on the opposite side from the labeled one.	36
2-6	The TPC	37

2-7	Schematic of a BBC, showing the scintillator tiles	38
2-8	Diagram showing the location of the ZDCs	39
2-9	The BEMC, with all other detectors removed.	40
2-10	Cross-sectional view of two adjacent BEMC modules. The grey layers are lead, the white scintillator. The BSMD is visible after the 5th lead layer. . .	42
2-11	Cross-sectional view of the BSMD	44
2-12	Schematic of the BSMD, demonstrating how an electromagnetic shower would appear in the two layers.	45
3-1	Two strips with significantly different fitting ranges: pedestal-subtracted ADC values are shown in blue, with the red line being the result of the fit. The dashed green lines show the location of three and five times the pedestal RMS. Note that since this data is zero-suppressed, the full pedestal doesn't appear in the ADC distribution.	50
3-2	A strip with its fitting range artificially shortened to keep the low end off the pedestal. See the caption for Figure 3.1 for more details.	50
3-3	A strip with its fitting range entirely on the pedestal.	51
3-4	BSMDE Relative Gains	55
3-5	BSMDP Relative Gains	56
3-6	Cartoon of the first stage of BSMD clustering: green strips are seed strips, red strips are not added to clusters, boxes represent the clusters that emerge from this stage.	57
3-7	GEANT deposited energy (in GeV) in the BSMDE plane vs. GEANT de- posited energy in the BSMD plane for simulated electrons.	60
3-8	Cartoon of a π^0 with its decay photons in two non-adjacent towers. The orange blobs represent the photons, the blue lines the BSMDP strips, the red lines the BSMDE strips, and the green bands the BSMD clusters (horizontal ones in the BSMDP plane, vertical ones in the BMSDE plane).	62

3-9	Cartoon of a π^0 with its decay photons in adjacent towers that are in different pods. The orange blobs represent the photons, the blue lines the BSMDP strips, the red lines the BSMDE strips, and the green bands the BSMD clusters (horizontal ones in the BSMDP plane, vertical ones in the BMSDE plane).	63
3-10	Cartoon of a π^0 with its decay photons in adjacent towers in the same pod. The orange blobs represent the photons, the blue lines the BSMDP strips, the red lines the BSMDE strips, and the green bands the BSMD clusters (horizontal ones in the BSMDP plane, vertical ones in the BMSDE plane). . .	64
3-11	Cartoon of a π^0 with both its decay photons in the same tower. The orange blobs represent the decay photons, the blue lines the BSMDP strips, the red lines the BSMDE strips, and the green bands the BSMD clusters (horizontal ones in the BSMDP plane, vertical ones in the BMSDE plane).	65
3-12	Probability for a pion to decay with some opening angle (actually $\sin(\theta/2)$, which goes into the invariant mass calculation) and p_T (in GeV) combination. The red rectangle in the lower left shows the region that is excluded by the cut.	69
3-13	π^0 candidates from data with $p_T > 4$ GeV and decay photons in non-adjacent towers. The η peak is clearly visible on top of the combinatorial background.	70
3-14	π^0 candidates from data with $p_T < 6$ GeV and decay photons in the same tower.	71
4-1	Vertex z distribution of events in our sample.	75
4-2	p_T distribution of events that passed each trigger (with the overall distribution in black).	78
4-3	Illustration of the event rotating procedure used in creating combinatoric background events whose underlying structure is similar to that of actual events.	80

4-4	Illustration of two possible ways in which a single cluster can be reconstructed as two separate clusters. The green strips and red strips represent different reconstructed clusters.	82
4-5	Illustration of the effect of applying the cut in p_T -opening angle space (here for π^0 candidates with $4 < p_T < 6.5$). The blue is the split-photon background, the black is the data. The split-photon background has been normalized so as to produce the best possible data-peak+background fit. . . .	83
4-6	Location of the invariant mass peak (obtained from a gaussian fit) vs. p_T for data (black), full simulations (blue), and single-pion simulations (red). .	84
4-7	For each p_T bin, the data invariant mass distribution (black) is fit to a sum of the simulated pion peak (purple), the split-photon background (blue), and the combinatorial background (green): the fit is in red.	86
4-8	(Data-Fit)/Data as a function of invariant mass for each p_T bin.	87
4-9	Polarization for both beams plotted by fill.	88
4-10	Single-spin asymmetries plotted vs. run index and fit to a constant: left, the yellow beam; right, the blue beam.	90
4-11	Number of pions per event, by p_T bin.	91
4-12	Difference between triggered and untriggered A_{LL} 's, with the latter calculated from the method of asymmetry weights using different $\Delta g(x, Q^2)$ scenarios, for each p_T bin. The $\Delta G = -.3$ scenario is in black, the $\Delta G = .3$ scenario is in red, and the DSSV scenario is in blue.	94
4-13	(Data-Embedding)/Data for each p_T bin. The embedding is scaled separately in each p_T bin such that the integrals of the data and the embedding over the range $.4 < m_{inv} < .5$ are the same.	96

4-14	Mass windows, shown with data and simulation/background distributions for each p_T bin. The black dotted lines are the original mass window boundaries, the red dotted lines show the new mass window boundaries, and the green line indicates the location of the π^0 mass peak.	98
4-15	A_{LL} vs p_T for each mass window. Black is the original mass window, $.08 < m_{inv} < .25$; red is $.1 < m_{inv} < .25$; green is $.08 < m_{inv} < .18$; and blue is $.1 < m_{inv} < .18$	98
4-16	A_{LL} with relative luminosity from the BBC (in black) and ZDC (in red).	99
4-17	$\pi^0 A_{LL}$ results, with systematic errors shown as green bands.	102
5-1	Comparison of the 2009 $A_{LL} \pi^0$ result to the 2006 result and several theoretical predictions and global fits.	104
5-2	The published STAR neutral pion A_{LL} result using data taken in 2005.	105
5-3	The 2005 and 2006 $\pi^0 A_{LL}$ results from the PHENIX experiments, compared to several theory curves.	106
5-4	The 2006 and 2009 STAR single-inclusive jet A_{LL} results, compared with a number of theory curves, including old DSSV (in green) and new DSSV (in magenta).	107
5-5	The 2009 STAR preliminary dijet result.	107

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

2.1	BSMD Parameters	43
3.1	Status Tests	52
3.2	Gain Status Tests	53
3.3	π^0 Candidate Cuts	71
4.1	Background Fractions	100
4.2	A_{LL} Uncertainties	101
4.3	A_{LL} Results	101

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 1

Introduction and Theoretical Background

The question of how the proton's spin arises out of its components is an interesting unsolved problem in hadron structure. But it is more than that: it is also an investigation of the nature of QCD, the theory of the strong force. Understanding how quarks and gluons, the fundamental particles of QCD, combine to form the proton, one of the fundamental particles of the world we see, has always been one of the main goals of the study of QCD, and attempting to determine the origin of the proton spin is one way in which we can make progress on this problem.

This thesis will present a measurement, using data taken by the STAR experiment of 200 GeV longitudinally polarized proton-proton collisions at the Relativistic Heavy Ion Collider at Brookhaven National Laboratory, of the double helicity asymmetry in the production of neutral pions, a measurement that can be used to constrain the value of the gluon polarization in the proton and thus bring us closer to an understanding of the proton's internal structure. The remainder of this chapter will discuss the theoretical background to this measurement. The second chapter will discuss the accelerator – RHIC – and detector – STAR – needed to do this measurement. The third chapter will explain how neutral pions

are found in the data, and the fourth how the double helicity asymmetry is constructed. Finally, the last chapter will discuss results and conclusions.

1.1 Proton Structure and QCD

The proton has been known to be a composite object since the late 1960's, when deep inelastic scattering (DIS) experiments carried out at SLAC revealed the existence of pointlike, spin-1/2 particles within the proton. DIS involves the scattering of an energetic lepton (in this case, an electron) off the nucleons in a fixed target. If the momentum transfer Q^2 is sufficiently large, the electron no longer interacts with the proton's charge distribution but instead sees the individual nearly free quarks that make up that distribution: this is seen experimentally by the fact that the proton structure functions cease to depend on Q^2 (because no matter how much more closely the electron probes the proton, all it sees are the pointlike quarks), a phenomenon known as Bjorken scaling. Additionally, the fact that the quark has spin 1/2 could be derived from the relationship between the two proton structure functions [10] [17].

Quark theory itself, however, dates to the early 1964, when Murray Gell-Mann attempted to explain the proliferation of elementary particles by proposing that all (excluding leptons) were bound states of two or three quarks, with the proton being made up of two charge-2/3 up quarks and one charge-1/3 down quark. In a way, the SLAC results confirmed Gell-Mann's hypothesis: however, Gell-Mann had proposed that the proton was made up of three bound quarks, while the objects (referred to as partons, a term that has persisted to refer to any constituent of a proton) that the electrons in DIS were scattering off of were free particles. This discrepancy was reconciled by quantum chromodynamics (QCD), the theory of the strong force.

QCD is a fully relativistic quantum field theory, analogous in many ways to quantum electrodynamics (QED). Just as "electro" in "electrodynamics" refers to the electric charge, the "chromo" in "chromodynamics" refers to the three charges that exist in QCD, known




Figure 1-1: *Feynman diagram for deep inelastic scattering*

as red, blue, and green (because particle physicists have a strange sense of humor). In QED, charged particles interact by exchanging photons; in QCD, they exchange gluons, although unlike the photon, the gluon also carries charge and so gluons can interact with each other. And just as in QED, there is a fundamental coupling constant, α_s rather than α . However, there is a very important difference in the behavior of the coupling constants in QED and QCD. In QED, the coupling constant increases at smaller and smaller distance scales (although the scales have to be quite small for this effect to become important). In QCD, the opposite occurs: as the distance scale becomes smaller (or the momentum transfer Q^2 increases), α_s falls as well. Conversely, as the distance scale increases (and Q^2 falls), α_s also gets larger. The result is a theory that incorporates both the SLAC results and Gell-Mann's quark theory: at the high Q^2 of DIS, the quark appears to be free, while if one wishes to regard a proton as a single object, one must look at a lower Q^2 value, where the large value of the α_s means that the quarks are bound. This feature of QCD in which bound quarks can almost achieve freedom inside a hadron is known as asymptotic freedom. QCD makes a further prediction: gluon interactions with the quarks would cause Bjorken scaling to be violated, as will be discussed in more detail later [29] [20].

A few additional notes are necessary on QCD in proton-proton colliders. In DIS, the experiment is relatively simple: an electron interacts with a quark, and the subsequent

behavior of the electron gives the experimenter almost all the necessary information. Unfortunately, DIS is limited in two ways: the electron can never interact with a gluon, because the gluon carries no electric charge, and the center-of-mass energy of the electron-quark interaction is necessarily limited. Ideally, we would like to accelerate quarks and collide them with each other: unfortunately, thanks to QCD quarks are bound permanently inside hadrons, so we have to settle for colliding two hadrons, e.g. two protons. Unfortunately, proton-proton collisions are considerably more complex than DIS events. In a pp collision, we are interested in the hard scattering of two partons (they could be quarks or gluons), but both protons contain a large number of partons, with various momenta, and we have no way to know which will be involved. Additionally, we cannot examine the partons after the scattering is complete: instead, we have to look at the results of their fragmentation into hadrons, without knowing how exactly that process will take place.

A couple of further features of QCD simplify the problem somewhat. The first is factorization, which allows us to break down the pp cross section for the production of some final state, say a π^0 , in terms of some set of kinematic variables P , $\frac{d\sigma}{dP}$, into a product of three probabilities, those for a particular initial state, for some scattering process to occur, and for the final state that you want to be produced. The initial state is described by the parton distribution functions (PDFs) $f(x, \mu^2)$, which describe the likelihood of finding a parton in a proton with some fraction x of the proton's momentum (μ is the factorization scale); the scattering process comes in via the partonic hard scattering cross-section $\frac{d\hat{\sigma}}{dP}$; and the final state is described by the fragmentation function (FF) $D(z, \mu^2)$, which gives the likelihood of a parton fragmenting into a hadron with some fraction z of the parton's momentum. Thus we can write the overall cross-section as (summing over possible initial states and end products of the partonic hard scattering)

$$\frac{d\sigma^{pp \rightarrow \pi^0 X}}{dP} = \sum_{f_1, f_2, f} \int dx_1 dx_2 dz f_1(x_1, \mu^2) f_2(x_2, \mu^2) \times \frac{d\hat{\sigma}^{f_1 f_2 \rightarrow f X'}}{dP} D_f^{\pi^0}(z, \mu^2)$$




Figure 1-2: *Neutral pion production in a pp collision*

The partonic hard scattering cross-section can usually be calculated using perturbative QCD, but the other two pieces of the puzzle are non-perturbative (low-energy) and so non-calculable (at least, not in perturbative QCD). However, there is a second feature of QCD, universality, which solves this problem. We can measure the PDFs and FFs in any other experiment (in particular, in theoretically simpler experiments such as DIS, or for fragmentation functions even simpler experiments like $e^+ - e^-$ collisions), and then apply the results to our experiment [2].

1.2 The Proton Spin

Spin (which for the proton, a fermion, has the value of $\frac{1}{2}\hbar$: as per usual, the \hbar will be dropped throughout) does not make an appearance in the above discussion, even though it has been important in investigations of proton structure since 1933, when Otto Stern measured the proton magnetic moment to be about 2.5 times as large as predicted, suggesting the possibility of some internal structure. Once the compositeness of the proton was experimentally confirmed, that left the question of how its spin arises, as it must, out of its components. In the simple bound-quark theory of Gell-Mann, it's obvious: two of the quarks have their spin in one direction while one has its spin in the opposite direction.

This simple model even does a relatively good job of predicting the proton anomalous magnetic moment. In the fully relativistic world of QCD, though, it is clearly insufficient. Gluons have spin, too, and their spin may contribute to the proton’s spin. The quark sea – quark-antiquark pairs that pop out of the vacuum before annihilating – may also contribute from its spin. Additionally, the quarks and gluons are not bound within the proton, so they may have some orbital angular momentum, which could also contribute to the overall spin. Working on the light cone and in the $A^+ = 0$ gauge, we can write a sum rule for the proton spin that includes all of these components [22]:

$$\frac{1}{2} = \int_0^1 dx \left\{ \frac{1}{2} \Delta q(x, Q^2) + L_q(x, Q^2) + \Delta g(x, Q^2) + L_g(x, Q^2) \right\}$$

Or, as it is more usually represented,

$$\frac{1}{2} = \frac{1}{2} \Delta \Sigma + L_g + \Delta G + L_q$$

$\Delta \Sigma$ is the quark polarization, the contribution of the quark spin to the proton spin, ΔG is the same for gluons, and L_g and L_q are the quark and gluon angular momentum contributions respectively. Note that It is also possible to write a relation in the laboratory frame,

$$\frac{1}{2} = \hat{L}_q + \frac{1}{2} \Sigma + \hat{J}_g.$$

In this frame, however, there is no gauge-invariant representation of the gluon polarization, and so it is subsumed into \hat{J}_g , the total gluon contribution to the proton spin [24].

1.2.1 Quark Polarization and the “Spin Crisis”

The first attempts to break down the proton spin in terms of its constituents focused on the quark polarization and angular momentum. $\Delta \Sigma$ can be written as

$$\Delta\Sigma = \int_0^1 dx \sum_{q=u,d,s\dots} (\Delta q(x) + \Delta\bar{q}(x))$$

where $\Delta q(x)$ is, for a given quark flavor,

$$\Delta q(x) = q_+(x) - q_-(x).$$

$q_+(x)$ is the polarized quark PDF, the probability of finding a quark not only with momentum fraction x , as previously discussed, but also with its spin aligned with that of the proton, and $q_-(x)$ is the same except for a quark with spin anti-aligned to the proton spin.

It is possible to use the parton model to predict $\Delta\Sigma$ by means of a sum rule, the Ellis-Jaffe sum rule, for the spin-dependent structure function g_1 ,

$$g_1(x) = \frac{1}{2} \sum_i e_i^2 \Delta q_i(x),$$

where the i 's run over the quark flavors. Making some assumptions, including that only u and d quarks contribute (i.e., that $\Delta s = 0$, as heavier quarks wouldn't be expected to contribute anyway), the integral of $g_1(x)$, and hence also $\Delta\Sigma$, can be written in terms of quantities that are calculable from β -decay experiments. This yields the Ellis-Jaffe sum rule prediction that $\Delta\Sigma \sim .6$, with the remainder of the proton spin arising from the quark angular momentum [6].

The first test of this prediction was carried out by the European Muon Collaboration (EMC) experiment at CERN. The EMC experiment was a polarized DIS experiment: i.e., just like an ordinary DIS experiment, except that both the beam (here muons, rather than electrons as at SLAC) and the target are polarized. By measuring the difference in cross-sections when the longitudinally polarized muons have their spin aligned and anti-aligned with the spin of the longitudinally polarized target nucleon (actually the asymmetry, which

is just the difference divided by the total cross-section), the EMC experiment was able to probe $g_1(x)$ and so derive a value for $\Delta\Sigma$. To everyone's surprise, the value was much smaller than expected, $\Delta\Sigma(Q^2 = 10.7\text{GeV}^2) = .13 \pm .19$ [13] (with subsequent experiments suggesting that the value is closer to .25 or .3). This experimental determination that the quark polarization did not actually make a major contribution to the spin of the proton gave rise to what was then known as the “spin crisis” and produced considerable speculation that the origins of proton spin might lie in the gluon polarization instead [23].

1.3 The Gluon Polarization

The gluon polarization, ΔG , is just the first moment of $\Delta g(x, Q^2)$, which is defined by analogy to $\Delta q(x, Q^2)$ as $g_+(x, Q^2) - g_-(x, Q^2)$, the difference between the PDF for a gluon with spin aligned with the proton and one with spin anti-aligned. Because photons do not couple to gluons, ΔG cannot be probed directly in polarized DIS experiments. It's possible to use polarized DIS to access ΔG indirectly by taking advantage of the violation of Bjorken scaling. The presence of gluons in the nucleus adds additional possibilities to the standard electron-quark interaction: the quark may also radiate a gluon, or absorb one, or radiate a gluon before the interaction that it absorbs afterwards, or engage in even more complicated interactions. Effectively, these quark-gluon interactions mean that the electron is not actually scattering off of a single, well-defined point particle: instead, it is seeing a structure composed of a quark and a field of gluons, and the more strongly the electron probes this structure (i.e., the higher the Q^2), the more gluons it sees. Therefore the structure functions, including g_1 , are not only dependent on x but also on Q^2 , and the extent of their dependence on Q^2 is dependent on the gluon PDF, either polarized or unpolarized depending on the experiment. Unfortunately, this technique, while very successful at measuring unpolarized gluon PDFs, provides only poor constraints on $\Delta g(x)$, due to the fact that data on g_1 is currently only available at a limited range of Q^2 values (see Figure 1.3 [5]). One way to surmount this obstacle is to avoid it by using a hadronic probe




Figure 1-3: Polarized gluon density at $Q^2=10$ GeV obtained from polarized DIS from de Florian, Navarro, and Sassot: the green band is the $\Delta\chi^2 = 1$ uncertainty, and the yellow the $\Delta\chi^2 = 2\%$ uncertainty.

which can interact directly with gluons. Due to QCD confinement, such a probe must actually be a parton within a polarized hadron, scattering off of a gluon inside another polarized hadron.

1.3.1 Measuring ΔG at a Hadron Collider

Just as the EMC experiment used cross-section asymmetries to measure $\Delta\Sigma$ in polarized DIS, we can use cross-section asymmetries to constrain ΔG in longitudinally polarized hadron-hadron collisions. Here, we use the double-spin asymmetry A_{LL} , defined as

$$A_{LL} = \frac{\sigma_{++} - \sigma_{+-}}{\sigma_{++} + \sigma_{+-}}$$

where σ_{++} is the cross-section for the production of, in the case of this thesis, neutral pions with the helicities (helicity is the projection of spin onto momentum, which is essentially the same as spin at high energies) of both proton beams aligned, while σ_{+-} is the cross-section with the polarizations of the beams anti-aligned. We can factorize this expression for A_{LL} into its components:

$$A_{LL} = \frac{\sum_{f=q,\bar{q},g} \Delta f_a \times \Delta f_b \times d\Delta \hat{\sigma}^{f_a f_b \rightarrow f^X} \times D_f^{\pi^0}}{\sum_{f=q,\bar{q},g} f_a \times f_b \times d\hat{\sigma}^{f_a f_b \rightarrow f^X} \times D_f^{\pi^0}}.$$

Here Δf_a is the polarized PDF for the parton f_a , $d\Delta \hat{\sigma}^{f_a f_b \rightarrow f^X}$ is the polarized hard-scattering cross-section, f_a is the unpolarized PDF for that parton, $\hat{\sigma}^{f_a f_b \rightarrow f^X}$ is the unpolarized hard-scattering cross-section, and $D_f^{\pi^0}$ is the fragmentation function for the final-state parton f to fragment into a π^0 . Clearly if either hard scattering parton f_a or f_b is a gluon, A_{LL} will access to $\Delta g(x, Q^2)$. As the hard scattering cross-sections can be calculated from perturbative QCD and the unpolarized PDFs and the FF can be measured in other experiments, $\Delta g(x, Q^2)$ can be extracted from the experiment, and hence also ΔG . However, it is difficult, and in practice theoretical predictions of A_{LL} are generated based on different values of ΔG , with experimental results used to constrain ΔG , rather than calculate it directly: see Figure 1.4.

1.3.2 The π^0

The π^0 is a neutral particle, composed of a mixture of $u\bar{u}$ and $d\bar{d}$ quark-antiquark pairs, $\pi^0 = \frac{1}{\sqrt{2}}(|u\bar{u}\rangle - |d\bar{d}\rangle)$. At a mass of $\sim 135 \text{ MeV}/c^2$, the π^0 is the lightest known meson. It decays electromagnetically and hence for our purposes instantaneously (that is, it will be treated as if the decay vertex and the collision vertex are the same), with a lifetime of 8.4×10^{-17} seconds, or $c\tau = 25.1 \text{ nm}$. The primary decay, with branching fraction 98.8%, is $\pi^0 \rightarrow \gamma\gamma$: this is the decay channel that will be used to identify π^0 s [18].

1.3.3 Note on Units

As is typical in particle physics, factors of the speed of light, c , will usually be suppressed. Thus, energy, momenta, and mass will all generally be expressed in units of GeV, instead of GeV, GeV/ c , and GeV/ c^2 respectively.




Figure 1-4: At left, several predictions of the value of $x\Delta g(x, Q^2)$, plotted vs. x ; at right, the corresponding predictions for the $\pi^0 A_{LL}$.

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 2

The RHIC Accelerator and the STAR Experiment

2.1 RHIC

The Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL) was, as its name suggests, originally intended for the purpose of colliding heavy ions in order to study the behavior of quarks and gluons at very high temperatures. It was predicted that at such temperatures a transition from hadronic to partonic degrees of freedom would occur, resulting in the formation of a new state of matter, the quark-gluon plasma. However, this thesis is on data from RHIC's other program: as the world's only polarized proton-proton collider, RHIC provides unique access to the proton spin structure. The RHIC complex (see figure 2.1) consists of a polarized proton source, a linear accelerator (LINAC), a Booster accelerator, the Alternating Gradient Synchrotron (AGS), and the RHIC accelerator itself. The LINAC accelerates protons to 200 MeV, after which they are injected into the booster which further accelerates them to 2 GeV, and then to the AGS which takes them to approximately 23 GeV. Finally the proton beam is split and injected into the RHIC rings (the two rings, carrying proton beams in opposite

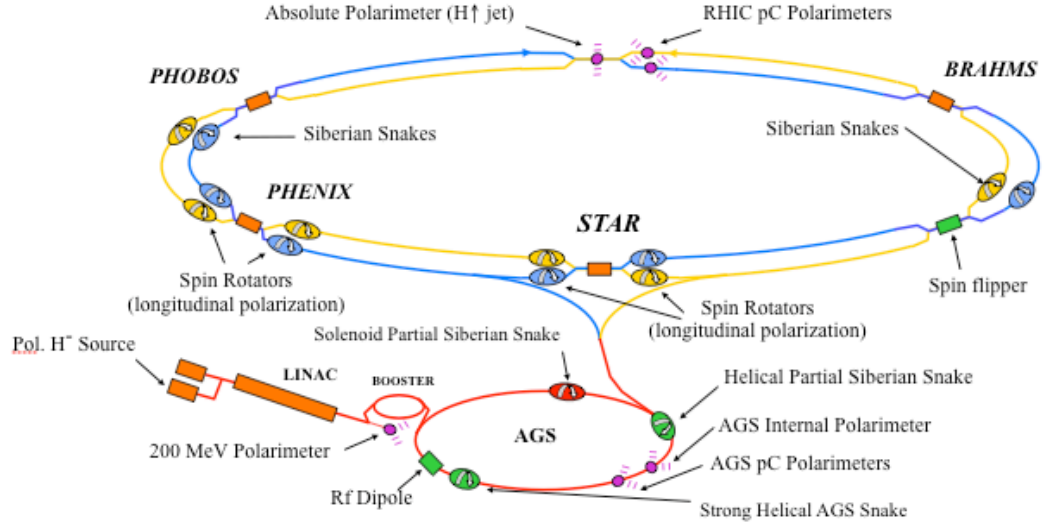


Figure 2-1: *The RHIC Complex, with an emphasis on the pieces that are important for polarized operations.*

directions, are referred to as yellow and blue) where they are accelerated to the final energy of 100 GeV (RHIC is also capable of accelerating to 250 GeV but that data is not used in this thesis), for collisions at center-of-mass energies of $\sqrt{s}=200$ GeV [11].

2.1.1 Polarization

The most important feature of RHIC for our purposes is its ability to generate beams of polarized protons and maintain their polarization through acceleration and collisions. An optically pumped polarized H^- source (OPPIS) is used to create polarized protons. First, hydrogen gas is ionized to obtain unpolarized protons: then the protons acquire electrons from a rubidium gas which is polarized by a continuous wave laser. The polarization of the rubidium electrons is transferred to the protons, and then the resulting hydrogen, with $\sim 80\%$ polarization at this point, acquires another electron to form an H^- ion, which allows it to be accelerated through the LINAC. At the end of the LINAC the ions are stripped

of their electrons by passage through a foil, creating a polarized proton beam for injection into the Booster [32] [12].

Maintaining the polarization of this beam is not easy, however. The polarization will remain stable as long as the proton spin is aligned or anti-aligned with the magnetic bending field of the accelerator, i.e. is transverse to the beam direction. The spin will precess around this vertical axis as the proton moves through the magnetic field. But the presence of horizontal magnetic fields, which can arise either accidentally from misaligned dipole magnets or be introduced deliberately by focusing quadrupole magnets, will perturb the spin, and if the frequency of the perturbation matches the frequency of the spin's precession, a depolarizing resonance occurs and the protons will start to lose polarization. The main tool to prevent depolarization from such resonances is a magnet known as a Siberian Snake. In essence, a Siberian Snake works by providing a strong enough magnetic field that the stable spin direction for all particles rotates by 180° , thus absorbing the loss of polarization into a simple switch of polarization directions (for more details, see [27]): see figure 2.2. The AGS has a partial Siberian Snake, which the protons pass through every rotation: it only partly rotates the proton spin, but by enough to cause a total spin flip when crossing a resonance. The AGS also deploys a pulsed RF dipole magnet, which is used for stronger depolarization resonances caused by horizontal focusing magnets. Like the Siberian Snake, the RF dipole magnet causes a complete spin flip: it is pulsed each time the beam crosses a strong resonance. In addition, each RHIC ring has two full snakes [32] [1].

It is also vitally important that the polarization of the beams be monitored. At RHIC, this is done using two types of polarimeter, a carbon ribbon one which works on the principal of Coulomb-nuclear interference (CNI) and a hydrogen-gas jet one (see Figures 2.3 and 2.4). The CNI polarimeter consists of a very thin carbon target and a set of silicon strip detectors. When the target is inserted into the beam, proton-carbon elastic scattering occurs. At very low momentum transfer squared, coulomb-nuclear interference effects, arising from the interference between electromagnetic interactions (which are spin-




Figure 2-2: *The change in the spin vector of a proton as it passes through a full Siberian snake: the blue arrow shows the direction of the beam.*

dependent) and hadronic interactions (which are not), give rise to a left-right asymmetry in the scattered carbon nuclei. The magnitude of this asymmetry is directly related to the polarization of the beam. Because of the high event rates, the asymmetry can be measured rapidly, allowing for periodic monitoring of polarization in both the blue and yellow rings as well as the AGS [32].

In order to convert the CNI asymmetry measurement into a beam polarization measurement, the CNI polarimeter must be calibrated, which is done using the H-jet polarimeter. The H-jet polarimeter measures absolute polarization: the beam is scattered off of a polarized proton target of known polarization, so the left-right asymmetry in the scattered protons can be used to calculate the beam polarization directly. As the target is a gas jet, rather than a solid ribbon as in the CNI polarimeter, the event rate is very low, which is why the H-jet polarimeter is used only for calibration rather than for ongoing polarization monitoring [28].

In 2009, STAR took roughly 14 pb^{-1} of p-p collisions at $\sqrt{s} = 200 \text{ GeV}$, with an average




Figure 2-3: *The RHIC CNI polarimeter setup, looking along the beam.*




Figure 2-4: *Schematic of the RHIC H-jet polarimeter, with the purple cylinder representing the hydrogen jet target.*

polarization of $\sim 59\%$, compared to 7.5 pb^{-1} and 55% polarization in the previous longest polarized p-p run in 2006. As the uncertainty in A_{LL} goes as

$$\sigma_{A_{LL}} \sim \frac{1}{P^2 * \sqrt{\int L dt}}$$

(where P is the average polarization and the integral over L is the integrated luminosity) the figure of merit (FOM) for a polarized proton-proton data set is

$$FOM = P^4 * \int L dt$$

The combination of more data and slightly better polarization gives a FOM for 2009 that is about twice that of 2006.

2.2 STAR

The Solenoidal Tracker at RHIC (STAR) detector, shown in Figure 2.5 as it was during Run 9, is more or less a standard particle physics detector. Its center is a Time Projection Chamber (TPC) which provides charged track identification. Around the TPC is the Barrel Electromagnetic Calorimeter (BEMC) which measures the energy of neutral particles and contains a shower-maximum detector (the BSMD) to provide fine position resolution: these detectors are the most important for this thesis. The Beam-Beam Counter (BBC) and Zero-Degree Calorimeter (ZDC) both consist of two detectors located on either side of the main detector which are used for triggering and luminosity monitoring. Other detector subsystems, mainly in the forward region, are not used in this thesis. A full description of STAR and all its subsystems can be found in [14] and its references.

A note on terminology: for the accelerator, the overall run is divided into fills, each corresponding to a beam store. The beam is injected into RHIC and left there until it degrades to the point that it is no longer useful, which takes roughly 8 hours if everything

is going well. Polarization measurements are done by fill. The detector then further divides the fills into runs. A run is a dataset taken at one time, usually lasting 30 minutes to an hour and comprising on the order of a million events if everything is going right. Relative luminosity is measured by run.

2.2.1 The TPC

The STAR TPC is a 4.2 m long by 4 m in diameter cylinder of P10 gas (90% Argon and 10% methane) that sits in a .5 T magnet. A well-defined and uniform electric field of about 135 V/cm is created by a thin and conductive central membrane, concentric field-cage cylinders, and the readout endcaps. Charged particles passing through the gas create ionization electrons, which drift to the endcaps, which have essentially a multi-wire proportional chamber readout system. The drift electrons avalanche near the anode wires: the resulting positive ions induce an image charge on the readout pads. As the induced charge is shared over several pads, the track position can be reconstructed to a small fraction of the pad width. The z coordinate of the track points is provided by the timing of the collision, combined with the known drift velocity (~ 5.5 cm/ μ s) of electrons in the drift gas under known electric and magnetic fields. The gating grid, a wire plane that acts as a shutter to control the entry of electrons into the readout (as well as preventing positive ions from entering the drift volume where they would distort the electric field) by having all wires at the same potential when the grid is ‘open’ and alternative wires at opposite potentials when it is ‘closed’, ensures that only one event is read out at a time. The TPC provides charged-particle tracking over $\eta = \pm 1.3$ and the full 2π azimuth [15].

In this analysis, the TPC is primarily used as a vertex finder: the TPC is able to find a vertex by extrapolating backwards from reconstructed tracks, typically with an error of 1 mm or less. The TPC also provides a veto: photon candidates reconstructed in the BEMC are only accepted if they do not have a charged track pointing to them.




Figure 2-5: *The STAR experiment. Note that the ZDCs are too far from the interaction point to be visible in this diagram. The BSMD is also not visible as it is inside the BEMC. The second BBC is just visible on the opposite side from the labeled one.*




Figure 2-6: *The TPC*




Figure 2-7: *Schematic of a BBC, showing the scintillator tiles*

2.2.2 The BBC and ZDC

There are two of both the BBC and the ZDC, located on either side of STAR. Each BBC is composed of segmented scintillator detectors surrounding the beam pipe located 374 cm from the interaction region. Each BBC provides full azimuthal coverage and has a pseudorapidity range of $2.1 < |\eta| < 5.0$. The BBC acceptance is roughly 53% of the total proton-proton cross-section of 51 mb. BBC coincidences (signals from both the east and west side BBCs) are used to measure the luminosity (both the total luminosity and the relative luminosity between different spin patterns) and also as the basis for some triggers [25].




Figure 2-8: *Diagram showing the location of the ZDCs*

The ZDC are small hadronic calorimeters located about 18 m from the interaction point: they provide a cross-check on the BBC luminosity measurements. Because the ZDCs are located behind magnets which sweep charged particles away from them, they measure only the production of neutral particles, as opposed to the BBCs which measure mostly the impingement of charged particles. This allows them to provide a somewhat independent luminosity measurement: however, the hit rate on the ZDCs is much lower, which is why the BBCs are the primary luminosity detector. A ZDC is composed of three modules, each of which consists of layers of tungsten and wavelength-shifting fibers, with a total of roughly one nuclear interaction length. As with the BBC, coincidences are used to measure both total and relative luminosity [9].

2.2.3 The BEMC

The BEMC, including the BSMD (which will be discussed in detail in the next section), is the most important detector for this analysis. The BEMC is a lead-scintillator sampling calorimeter, covering from -1 to 1 in η and the full azimuth in ϕ . It is located immediately outside the TPC, at an inner radius of ~ 225 cm and an outer radius of ~ 265 cm. The calorimeter is divided into 120 modules, each covering 6 degrees in ϕ and 1.0 in η . Each module is further subdivided into 40 towers – the tower being the basic unit of the




Figure 2-9: *The BEMC, with all other detectors removed.*

calorimeter – 2 in ϕ and 20 in η , for a total of 4800 towers, each .05x.05 in $\eta \times \phi$. The towers are furthermore projective, with each tower pointing back to the interaction region. Each tower consists of a stack of alternating 5 mm thick sheets of lead and scintillator, with the BSMD located at a depth of roughly 5 radiation lengths (at $\eta = 0$). There are 20 layers of lead and 21 of scintillator, for a total of roughly 20 radiation lengths (again, measured at $\eta = 0$: at higher values of $|\eta|$ a particle may take an angled path through the tower, which could change the amount of material it sees) [16].

Like any electromagnetic calorimeter, the BEMC measures the energy (and, with a very poor resolution, the position) of electrons, positrons, and, most importantly for this measurement, photons. An energetic photon will interact with lead by pair-producing: the resulting electron and positron will then interact with lead by bremsstrahlung, producing more photons that pair produce in their turn, and so on until the energies of the produced particles drop to the point where ionization and the photo-electric effect start to dominate instead and the shower dies out. As the energetic particles in the cascade pass through the scintillator layers (being plastic, they are much less dense than the lead and so are

unlikely to induce either pair-production or bremsstrahlung), they produce scintillation light, which is read out. The longitudinal size of the shower is governed by the radiation length (X_0): each lead tile is roughly one radiation length thick, giving a total of 20 (X_0), which should be enough to contain the whole shower for all the photons considered in this analysis. The transverse size of the shower is governed by the Molière radius $R_M \sim 1.6$ cm in lead – making the tower roughly 9 Molière radii in both the η and ϕ directions. As on average, 90% of the shower energy falls within a cylinder of radius R_M , the tower should contain almost all of the energy of most particles that fall within it. The downside of this large transverse size is, of course, poor position resolution, which is why the BSMD is necessary [18].

The BEMC is not designed to measure the energy of hadrons. The nuclear interaction length of lead is ~ 17 cm, larger than the amount of lead (about 10 cm) in a tower: to have any hope of fully containing a hadronic shower, the BEMC would have to be far larger than it is. Instead, hadronic particles generally pass through the BEMC without showering, and are referred to as minimum ionizing particles (MIP) because they deposit essentially the minimum possible amount of energy due to ionization as they pass through. MIPs are useful in calibrating the detector, as will be discussed in the next chapter.

Because the BEMC resides inside the STAR magnet, the light from the scintillating tile is carried several meters by wavelength-shifting (WLS) fiber to the photomultiplier tubes (PMTs) that read out the tower. The light from all the tiles in one tower is merged onto a single PMT using a light mixer which ensures that the light from all the fibers shines on the same part of the PMT photo-cathode, thus improving detector uniformity. Cosmic ray and test beam results showed that approximately 3 photo-electrons were generated per MIP, leading to an expected energy resolution of $\delta E/E \sim 14\%/\sqrt{E(\text{GeV})} \oplus 1.5\%$, though tower non-uniformities and between-tower cross-talk make the actual energy resolution slightly worse [16].




Figure 2-10: *Cross-sectional view of two adjacent BEMC modules. The grey layers are lead, the white scintillator. The BSMD is visible after the 5th lead layer.*

Table 2.1: *BSMD Parameters*

Depth (in BEMC)	$5X_0$ at $\eta=0$
Occupancy (p+p)	$\sim 1\%$
Chamber Depth	20.6 mm
Wire Diameter	50 μm
Gas Amplification	~ 3000
Signal Length	110 ns
BSMDE Strip Width (Pitch)	1.46 (1.54) cm for $ \eta < 0.5$ 1.88 (1.96) cm for $ \eta > 0.5$
BSMDP Strip Width (Pitch)	1.33 (1.49) cm
Strips per 2x2 Tower Group (Pod)	15
Strips per Module	300
Modules	120
Total Channels	36000

2.2.4 The BSMD

To provide the fine spatial resolution that is needed to distinguish between energetic neutral pions and photons, there is a shower-maximum detector located at a depth of roughly $5 X_0$, where the shower is expected to peak. The BSMD is a two-layer multi-wire proportional chamber (MWPC) read out with two sets of strips, one (the inner) providing resolution in η and the second (outer) in ϕ . There are a total of 36000 strips in the BSMD: each 2x2 group of towers, covering 0.1×0.1 in $\eta \times \phi$, contains 15 strips in each plane (referred to as the η and ϕ planes, after the coordinate in which they provide discrimination). Each strip extends over .1 in the opposite coordinate, while the strip pitch in each plane is roughly .007 in η/ϕ , or ~ 1.5 cm (see Table 2.1 for details) [16].

The BSMD gas is 90% argon and 10% carbon dioxide. Charged particles going through the detector ionize the gas: the electrons produced drift under an electric field towards anode wires. Close to an anode wire, they start to accelerate, resulting in additional ionization and yet more electrons heading towards the wire. The resulting avalanche induces an image charge on the pad, which is read out. A MIP traversing the barrel calorimeter




Figure 2-11: *Cross-sectional view of the BSMD*

will leave a relatively small amount of ionization, but an electromagnetic shower, which will consist of many electrons at its maximum, should leave a large and roughly gaussian signal in both planes. Matching the signals in the two planes allows for a relatively accurate determination of the position of the shower. Test beam data gives a position resolution of $5.6/\sqrt{E(\text{GeV})} \oplus 2.4\text{mm}$ for the η plane and $5.8/\sqrt{E(\text{GeV})} \oplus 3.2\text{mm}$ for the ϕ (the additional layer of material between the two planes degrades the ϕ resolution somewhat). Unfortunately, the BSMD energy resolution is not as good, being roughly $\delta E/E \sim 12\% \oplus 86\%/\sqrt{E(\text{GeV})}$ for the η plane and slightly worse for the ϕ . Additionally, the dynamic range of the BSMD readout is relatively limited, leading to saturation of the central strip at photon energies of ~ 8 GeV and above.




Figure 2-12: *Schematic of the BSMD, demonstrating how an electromagnetic shower would appear in the two layers.*

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 3

π^0 Finding

The first step in a neutral pion measurement, identifying neutral pions in the dataset, is intrinsically non-trivial because the π^0 decays immediately into two photons. Therefore finding neutral pions is actually a matter of reconstructing them from their decay products. The fact that these products are photons adds an extra layer of difficulty: as neutral particles, photons leave no tracks in the TPC or any other tracking detector. Therefore there is no useful track pointing back to a vertex and so no way to tell if any two photons are due to a single π^0 or not. However, it is possible to overcome these difficulties using the STAR BEMC and BSMD. Before we can use these detectors, however, we need to determine what the electronic responses we get from the detectors actually correspond to in terms of the particles.

3.1 Calibration

3.1.1 BSMD Calibration

Calibrating the BSMD is a tricky proposition since it is inside the BEMC. Calibrating a detector essentially involves finding the factor of proportionality between the detector output in ADCs and the deposited energy in GeV. For the BEMC, we can use charged

particles whose momentum is measured by the TPC: as the particle should deposit all of its energy in the BEMC, this makes it possible to relate energy and ADC (as will be discussed in more detail in the next section). The situation isn't so simple for the BSMD because the particle only deposits a fraction of its energy in the detector, and there's no other detector to provide an independent measurement of the amount of deposited energy. Theoretically, we could calibrate to the particle's energy, rather than just the deposited energy, but due to saturation effects such a calibration would start to become useless at particle energies of 6-8 GeV. Instead, we use the fact that this analysis is not sensitive to the absolute BSMD calibration. BSMD energies are used only to find the center of a cluster through a weighted average of energy or to calculate (in some cases) the ratio $z_{\gamma\gamma}$ (which will be discussed below): in both cases, we are interested in a ratio of BSMD strip energies and not an absolute BSMD strip energy. Therefore, all that is necessary for the analysis is to perform a relative calibration, in which we equalize the response of BSMD strips over a large block of strips (for the BSMDE, all the strips with $i/10 < \eta < (i+1)/10$, for $-10 \leq i < 9$; for the BSMDP, the whole detector). Any absolute calibration factor applied for that block of strips to obtain a real energy value would simply drop out of the ratios anyway.

The relative gain for a strip j in η bin i is defined as $Avg_i/Slope_j$, where Avg_i is the average slope of all strips in the η bin and $Slope_j$ is the slope of the strip. The slope defines the response of the strip: thus, we are correcting all the strips in the η bin to have the same response. The slope α is found by fitting the ADC distribution of a BSMD strip to a falling exponential $Ce^{-\alpha x}$. However, precisely because we expect the response to differ for different strips, we cannot use a fixed fitting range to find the slope. As the same ADC range will correspond to different energy ranges for different strips, we would be comparing the responses of the strips to different energy deposits. Instead, we find a fixed energy range by integrating (or rather, since we have a discrete distribution, summing) from the end of the ADC range to a fixed fraction of the total number of events in the sample to find the

low and high ends of the fitting range (see below).

$$\sum_{1024}^{x_{high}} N(\text{ADC}) = .000045 \cdot \text{nEv}$$

$$\sum_{1024}^{x_{low}} N(\text{ADC}) = .00031 \cdot \text{nEv}$$

Note that the calibration is not sensitive to the exact values of the parameters 4.5×10^{-5} and 3.1×10^{-4} . This allows us to require that x_{low} always be at least five times the pedestal RMS to avoid having the low end of the fitting range be on the pedestal itself, which would distort the result of the fit (see Figure 3.2, for instance, where the fitting range is shortened by this method). Figure 3.1 shows an example of how two different strips could have quite different fitting ranges from this method.

Before we can begin fitting, there is one further step to be taken: we must identify and remove strips that had significant electronic or other problems during the run. To do this, we look at 500 GeV events from a set of runs that were carefully analyzed for high quality. Because the pp500 run occurred before the pp200 run and no settings changed in between the two runs, using 500 GeV events here should not present a problem. Two separate datasets are required to find the status. In order to look at the tails of the strip ADC distributions, we use events taken with the L2W trigger, which requires a high-energy deposit in at least one BEMC tower (it was designed to look for W bosons via a high- p_T electron). However, because the BSMD has so many channels, it is usually zero-suppressed to save space: only strips that are at least 1.5σ above pedestal are written out. This makes it impossible to QA the pedestal of strips using the L2W data. Instead, we use the small fraction of events which are read out without BSMD zero-suppression. Since these events are randomly selected, they lack the statistics that we would need to analyze the tails, as well as the pedestals, of the ADC distributions. Therefore the high-energy and low-energy

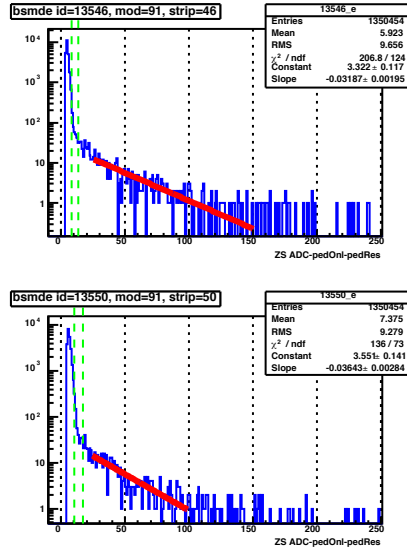


Figure 3-1: Two strips with significantly different fitting ranges: pedestal-subtracted ADC values are shown in blue, with the red line being the result of the fit. The dashed green lines show the location of three and five times the pedestal RMS. Note that since this data is zero-suppressed, the full pedestal doesn't appear in the ADC distribution.




Figure 3-2: A strip with its fitting range artificially shortened to keep the low end off the pedestal. See the caption for Figure 3.1 for more details.

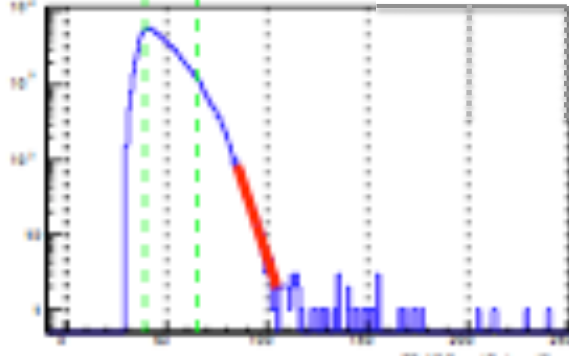


Figure 3-3: *A strip with its fitting range entirely on the pedestal.*

response are examined separately.

Table 3.1 lists the tests applied to each strip: the first four to the pedestal, ensuring that its location and width are in a reasonable range and that it has only one peak, and the last three to the tail, ensuring that it exists and that it is not a “hot” strip, i.e. one that is producing high amounts of ADC counts without a correspondingly large energy deposition. As with the parameters determining x_{low} and x_{high} above, the calibration is not sensitive to the exact values of the parameters used in the table. Fatal tests are those which mark a strip as bad if it fails them.

Additionally, there are several other tests which are applied to strips before their slopes are calculated: these are listed in Table 3.2. While relative gains cannot be calculated for these strips, they are still potentially useful in other applications, so they are given a bad ‘calibration’ status but not a bad ‘overall’ status. Dead strips are eliminated, as are strips from modules which have too many bad strips, or in two cases have other problems which were not caught by the standard QA. The other tests have to do with the shape of the pedestal, and are largely intended to eliminate strips in which the entire fitting range is on the pedestal (see Figure 3.3, for example). After all the status tests were applied, 30412, or 84%, of the BSMD strips had failed none of the tests, and a further 1914 had failed only one or more non-fatal tests, so that $\sim 90\%$ of the strips were usable.

Table 3.1: *Status Tests*

Status Bit	Failure Condition	Test Fatality	Notes
1	$ \text{MPV} > 50$	fatal	MPV is Most Probable (ADC) Value
2	$\text{RMS} < 1$ or $\text{RMS} > 15$	fatal	RMS calculated over range $\text{MPV} \pm 23$
3	Integral Ratio $\text{IR} < .95$	fatal	<p>If $\text{RMS} < 7.7$:</p> $\text{IR} = \frac{\int_{\text{MPV}-23}^{\text{MPV}+23}(\text{ADC})}{\int_{\text{all ADC values}}(\text{ADC})}$ <p>If $\text{RMS} > 7.7$:</p> $\text{IR} = \frac{\int_{\text{MPV}-8.5\text{RMS}}^{\text{MPV}+8.5\text{RMS}}(\text{ADC})}{\int_{\text{all ADC values}}(\text{ADC})}$
4	$\text{RMS} > 11.5$	Not fatal	RMS calculated as in bit 2
5	$\text{Par1} < .00003$ or $\text{Par1} > .02$	Not fatal	$\text{Par1} = \frac{\int_{300}^{500}(\text{ADC})}{\int_{\text{all ADC values}}(\text{ADC})}$
6	$\text{Par2} < .00003$ or $\text{Par2} > .02$	Not fatal	$\text{Par2} = \frac{\int_{500}^{800}(\text{ADC})}{\int_{\text{all ADC values}}(\text{ADC})}$
7	$\text{Par3} < .00003$ or $\text{Par3} > .02$	Not fatal	$\text{Par3} = \frac{\int_{\text{ADC} > 800}(\text{ADC})}{\int_{\text{all ADC values}}(\text{ADC})}$
	Conditions 5, 6, and 7 all fail	fatal	If conditions 5-7 fail, the strip is bad.

Table 3.2: *Gain Status Tests*

Status Bit	Failure Conditions	Test Fatality	Notes
1	$\sum_{\text{ADC} > 50} N(\text{ADC}) < 10$	fatal	
2	$\text{Int2}/\text{Int1} < .07$	fatal	$\text{Int1} = \sum_{50}^{100} N(\text{ADC})$ $\text{Int2} = \sum_{100}^{150} N(\text{ADC})$
3	$\text{Int3}/\text{Int2} < 1$	fatal	$\text{Int3} = \sum_0^{50} N(\text{ADC})$
4	$\text{Int2}/\text{Int3} > 1$	fatal	
6	Bad module	fatal	A module is bad if it has more than 50 bad strips; additionally, modules 116e and 112 are marked bad. Bit 5 is no longer used but bit 6 was not changed to 5 for consistency.

In order to maximize the number of events in our sample, we use a different data set to calculate the slopes, namely events (from a good set of, again, 500 GeV runs) which were taken with a minbias trigger, which simply takes a random sample of events. These events will be zero-suppressed, but as the zero suppression only removes hits on the strip pedestal it should not effect the slope calculation. It is worth noting that the calibration is not particularly sensitive to the average slope in each η bin, as any deviation from the “correct” average slope is just a multiplicative factor that will cancel out whenever we use

the BSMD energies. The resulting relative gains for both BSMD planes are plotted vs. η and ϕ in Figures 3.4 and 3.5: the BSMD crates, as well as every fifth BSMD module, are clearly indicated, so that bad modules or areas of bad strips are visible.

3.1.2 BEMC Calibration

The BEMC tower calibration includes an absolute calibration in addition to the relative calibration, since it is necessary to know the actual energy values of the towers. It proceeds in two steps, starting with a relative calibration. While the BSMD uses the high-energy tail of the ADC distribution to judge the response of each strip, the tower response is judged from the lowest-energy part of the distribution, the MIP response. Test-beam results and simulations indicate that the MIP peak should occur at roughly 20 ADC counts above pedestal [31]. By looking at tracks with momentum of at least 1 GeV (measured by the TPC) that are projected to enter and exit the same tower, if that is the only track entering the tower and none of the neighboring towers have any energy, we can get a good sample of MIP tracks, and obtain a MIP peak by fitting the region around 20 ADC to a gaussian. The location of the MIP peak is used as a measure of tower response: the relative calibration is then done for all the towers in a crate at a given η value. The presence or absence of the MIP peak is also used to determine tower status: if there is no MIP peak, the tower is not usable.

The absolute calibration is done using electrons and positrons. Electrons can be identified using the TPC: the dE/dx values for an electron are significantly different than those for a MIP (usually a charged pion). For each group of towers, the value of electron E/p is plotted for all the electrons that fall into the towers (using the nominal calibration which assumes that the maximum ADC corresponds to 60 GeV), after the energy has been corrected by the previously measured relative gain and by a factor derived from simulation to account for loss of energy in the material between the TPC and the BEMC and into the space between the towers. The absolute calibration factor is then the correction which

bsmde relative gain vs. eta and phi (excellent modules only)




Figure 3-4: BSMDE Relative Gains

bsmdp relative gain vs. eta and phi




Figure 3-5: BSM DP Relative Gains

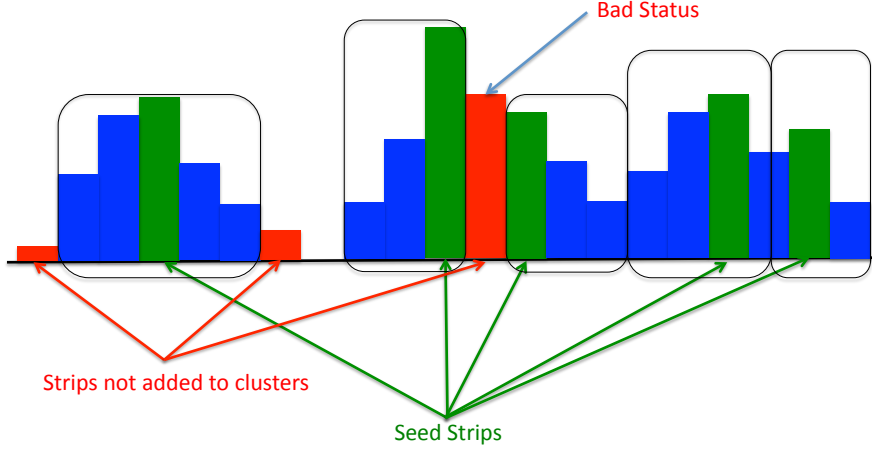


Figure 3-6: *Cartoon of the first stage of BSMD clustering: green strips are seed strips, red strips are not added to clusters, boxes represent the clusters that emerge from this stage.*

needs to be applied such that the peak of the E/p distribution for each group of towers is at 1. The product of the relative and absolute calibrations for each tower is that tower's calibration factor [31].

3.2 BSMD Clustering

Photons leave signals only in the STAR calorimeter. The towers themselves are too large for our purposes: once a π^0 is sufficiently energetic, both of its decay photons will fall in the same tower. However, the photons should also deposit energy in the BSMD, and we can use this to identify and separate photons up to a reasonably high energy.

Figure 3.6 shows an example of how the BSMD clustering process works. It starts with one BSMD plane (clustering is done separately in each layer). A BSMD plane can be thought of as a giant cylinder: if the cylinder is “unrolled” it becomes a rectangle, divided into rows (the η plane, or the BSMDE for η) or columns (the BSMDP), which are then further divided into strips. Clustering starts by finding for each row or column all the peak strips, defined as strips that have good status, more than 100 ADC of deposited energy

(all ADC values quoted are those obtained after the relative calibration discussed above is applied), and have more energy than the strips to either side of them (strips with bad status, or dead strips, are taken to have no energy). Then, for each peak strip, starting from the most energetic, a cluster is constructed by adding strips to both sides of the cluster until one of three stop conditions is reached: either the strip that would be added is dead/has fewer than five ADC of deposited energy, has more energy than the strip that was just added, or would make the third strip on that side of the cluster (so that a cluster has a maximum of five strips: in almost all cases, that is enough to catch the vast majority of the cluster's energy).

The perfect case can be seen on the left of Figure 3.6, where a nice five-strip cluster is formed, leaving off only two strips that have very little energy. However, in most cases this is not seen: even cases where there are no problems reconstructing the cluster usually don't have the full five strips, and there are many possible problems that can arise in photon reconstruction. For one thing, a photon may fall such that its shower shares energy between two columns or rows of a BSMD plane: this is dealt with by looking for clusters that are in adjacent rows or columns and are quite close together and merging them. Although merging sounds quite final, it is actually merely the forming of a hypothesis that two clusters are from a single photon: a new cluster is created, but the two original clusters are maintained. Merging can also happen within a row or column: if two clusters are separated by a dead strip that lies between their peaks, as in the center of Figure 3.6, then they may also be merged. And if a cluster containing just one strip lies adjacent to, or separated by just one empty strip from, another cluster, the two will also be merged. Finally, we have the case that is shown on the right of Figure 3.6, in which two clusters are separated by a single strip, but that strip is not dead and has energy. In this case, the two clusters could be from a single photon that merely had a strip that was expected to have a low deposited energy value fluctuate high, or from two photons that are quite close together. The former forms one of the backgrounds that will be discussed in the next

chapter: the latter is dealt with by assigning part of the energy of the strip to each cluster (in proportion to the amount of energy in the peak strip).

A cluster has two important properties: the total amount of energy in its strips, and its location. For BSMDE (BSMDP) clusters, the cluster's position in η (ϕ) is just the energy-weighted average of the positions of the strips it is made from. As all the strips making up the cluster are at one constant value of ϕ (η) for clusters in the BSMDE (BSMDP), the cluster's location in ϕ (η) is just that constant value. The exception is the case when a cluster is composed of two clusters merged across a row or column, when the position in ϕ (η) is just the energy-weighted average of the center positions of the two rows or columns. The energy and location are used in the next step, which is combining clusters in the BSMDE plane and clusters in the BSMDP plane to form photon candidates. Clusters in the two planes are combined only if they are within .1 in $\eta - \phi$ and if they are located in the same pod (a pod being defined as a 2x2 group of towers that share the same group of 15 strips in both BSMD planes). Additionally, the clusters are combined only if the ratio of their energies is between .25 and 4. More precise energy matching is impossible due to the poor energy resolution of the BSMD: Figure 3.7 plots the energy deposited in the BSMDE plane (as given by the GEANT simulation of the detector) against that deposited in the BSMDP plane for simulated single-electron events (an electron should deposit its energy in the BSMD planes in a manner very similar to a photon), showing that the correlation is too poor to support anything stricter. A photon candidate's location is simply the η location of its BSMDE cluster and the ϕ location of its BSMDP cluster.

Once we have photon candidates reconstructed, two cuts are imposed on them. The first of these is the charged track veto. Electrons, and charged pions that happen to shower in the BEMC (which are just as common, if not more so, than electrons, since charged pions are produced in vastly larger quantities) would produce BSMD clusters that would mimic that of a photon: we dispose of these spurious photons by discarding the photon candidate closest to a charged track, if it is within .045 of that track in $\eta - \phi$. Additionally,

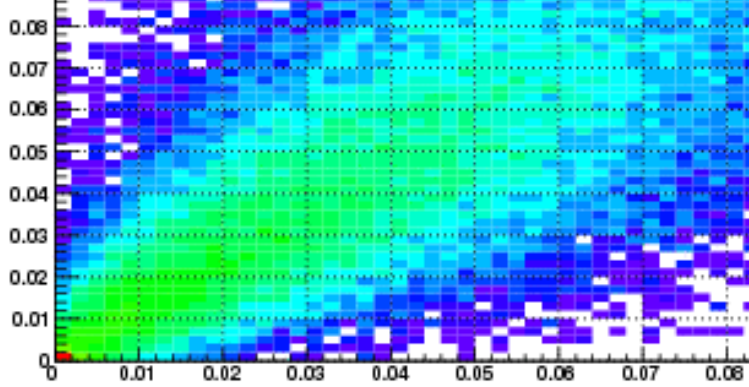


Figure 3-7: *GEANT deposited energy (in GeV) in the BSMDE plane vs. GEANT deposited energy in the BSMD plane for simulated electrons.*

to avoid having to deal with detector edge effects, any photon candidate with $|\eta| > .95$ is discarded.

Finally, because we have no way of associating photon candidates with each other, all possible combinations of photon candidates are taken to form π^0 candidates.

3.3 π^0 Candidates

For the π^0 candidates, it is necessary to take the BEMC towers, which we have to this point ignored, into account. The only way to evaluate the quality of a π^0 candidate is by calculating its invariant mass: the closer this is to .135 GeV, the better the candidate (with the proviso that we expect considerable width in the invariant mass distribution). The π^0 invariant mass is calculated from the following equation:

$$m_{inv} = E_{tot} * \sin(\theta/2) * z_{\gamma\gamma}$$

where E_{tot} is the pion energy, θ is the opening angle of the decay, and $z_{\gamma\gamma}$ is the energy asymmetry between the two decay photons, defined as

$$z_{\gamma\gamma} = \frac{|E_{\gamma,1} - E_{\gamma,2}|}{E_{\gamma,1} + E_{\gamma,2}}$$

Invariant mass is usually calculated directly from the 4-vectors of the decay particles: this (entirely mathematically equivalent) method was chosen because it isolates the problem of assigning energy to the decay photons into one piece of the calculation, $z_{\gamma\gamma}$.

The opening angle θ can be determined from the positions of the two photon candidates and the position of the collision vertex. The other two elements of the invariant mass are more difficult to determine, however: for them, it is useful to sort π^0 candidates by the topology of their decay in the BEMC.

3.3.1 Decay Topologies

Figures 3.8-11 shows examples of each of the four decay topologies we deal with. In each, the lines in the background represent BSMD strips, in both planes; the orange blobs are the decay photons; and the green bands represent the BSMD clusters that result from the photons depositing BSMD energy (though this shouldn't be taken too literally: clusters will often have more than the two or three strips depicted here, and no merged clusters are shown). The easiest case to deal with is that of Figure 3.8. Here the two decay photons fall in non-adjacent towers. This means that they are well separated, so that in this case the total pion energy is calculated as the sum of individual photon energies, with each photon's energy being the energy of the highest-energy 2x2 group of towers containing the tower that the photon candidate falls into (the "struck tower"). Note that for this purpose, towers containing photon candidates are treated as if they had zero energy, to avoid (as far as possible) including energy deposited in the BEMC by other particles: this exclusion of towers containing photon candidates (except for the struck towers, of course) in the energy calculation is done for all decay topologies. Additionally, the energy is corrected for the case in which the two 2x2 tower groups overlap to avoid doublecounting. The $z_{\gamma\gamma}$




Figure 3-8: *Cartoon of a π^0 with its decay photons in two non-adjacent towers. The orange blobs represent the photons, the blue lines the BSMDE strips, the red lines the BSMDE strips, and the green bands the BSMDE clusters (horizontal ones in the BSMDE plane, vertical ones in the BSMDE plane).*

is calculated using the energies of the struck towers.

The case of Figure 3.9 is a little more complicated: here, the photons are still in different pods but are in adjacent towers, which means that they are too close to make it feasible to calculate the energy of the photons separately. Additionally, there are now two subcases to consider: the first is that of the figure, in which the towers are directly adjacent, and the second, not shown, is when the towers only touch at a corner. In the first subcase, the energy of the pion is taken to be the sum of the energies of the struck towers and that of the two highest-energy towers of the ten that form the outside of the 3x4 (or 4x3) group of towers immediately surrounding the two struck towers. In the second subcase, the energy is that of the struck towers, plus the energy of the other two towers in the 2x2




Figure 3-9: *Cartoon of a π^0 with its decay photons in adjacent towers that are in different pods. The orange blobs represent the photons, the blue lines the BSMDP strips, the red lines the BSMDE strips, and the green bands the BSMD clusters (horizontal ones in the BSMDP plane, vertical ones in the BMSDE plane).*




Figure 3-10: *Cartoon of a π^0 with its decay photons in adjacent towers in the same pod. The orange blobs represent the photons, the blue lines the BSMDP strips, the red lines the BSMDE strips, and the green bands the BSMD clusters (horizontal ones in the BSMDP plane, vertical ones in the BSMDE plane).*

group (assuming, as always, no photon candidates). Finally, to this energy is added the energy of the two highest-energy of six towers, three from each of the 2x2 groups of towers that contains one of the struck towers and extends away from the other struck tower. As in the previous case, $z_{\gamma\gamma}$ is calculated from the struck tower energy.

From the point of view of calculating the energy and $z_{\gamma\gamma}$, the case of Figure 3.10, in which the two photons fall in the same pod but different towers, is much the same as Figure 3.9. There are two subcases, one in which the struck towers are diagonally adjacent and one in which they are not (cases a and c in the diagram are really the same), which are treated the same way as the subcases in Figure 3.9 are, and the calculation of $z_{\gamma\gamma}$ does not change. However, the fact that the photons fall in the same pod means that there will be




Figure 3-11: *Cartoon of a π^0 with both its decay photons in the same tower. The orange blobs represent the decay photons, the blue lines the BSMDP strips, the red lines the BSMDE strips, and the green bands the BSMD clusters (horizontal ones in the BSMDP plane, vertical ones in the BSMDE plane).*

ghost pion candidates, as will be discussed in the next section.

Finally, there is the case of Figure 3.11, in which both photons fall in the same tower. The energy of candidates of this decay type is defined as the energy of the struck tower plus the energy of the two highest-energy of the eight towers that make up the rest of the 3x3 group of towers centering on the struck tower. Additionally, the way in which $z_{\gamma\gamma}$ is calculated clearly has to change: with only one struck tower for both photons, it is not possible to use the struck tower energy as a stand-in for the photon energy. Instead, the energy deposited in the BSMD planes is used. For reasons that will be explained in the next section, rather than summing the BSMDE and BSMDP energies of the photons and using them to calculate $z_{\gamma\gamma}$, we calculate separate $z_{\gamma\gamma}$ values for each BSMD plane and

then take their geometric mean, so that

$$z_{\gamma\gamma} = \sqrt{\frac{|E_{\eta,1} - E_{\eta,2}|}{E_{\eta,1} + E_{\eta,2}} * \frac{|E_{\phi,1} - E_{\phi,2}|}{E_{\phi,1} + E_{\phi,2}}}$$

Note that, in the case where the two photons share a cluster in one of the BSMD planes (which is not uncommon: because a strip extends over the entire width of two towers, two photons at more or less the same value of η or ϕ in the same tower will deposit all their energy in one plane into the same few strips), the $z_{\gamma\gamma}$ is calculated using only the energy deposited in the other BSMD plane. Because of the poor energy resolution of the BSMD, the correlation between the BSMD energy and the particle energy is not expected to be particularly strong, making this a less than ideal method of measuring $z_{\gamma\gamma}$ for any individual pion. However, overall it should not have any negative effect on the measurement beyond widening the mass peak.

3.3.2 Ghost Candidates

The ghost π^0 candidates – so called because they are purely a detector effect, rather than resulting from background, flaws in the π^0 reconstruction algorithm, or even detector problems – arise from the fact that BSMDE(P) strips extend over two towers, and thus 15 strips in the BSMDP(E) plane, in the ϕ (η) direction. In the situation of Figures 3.10 and 3.11, when both decay photons fall in the same pod, finding a photon candidate essentially means looking for the intersection of a cluster in the BSMDE plane with one in the BSMDP plane. But due to the length of strips in the opposite direction to their plane, the BSMDE cluster from one photon will intersect not only its own photon's BSMDP cluster, but also with that from the other photon's BSMDP cluster. The result is that two photon candidates becomes four. And since pion candidates are obtained combinatorially from photon candidates, the original single π^0 gives rise to six π^0 candidates. Since same-pod decays become dominant at higher energies, as the opening angle shrinks, this threatens

to drown the actual data in ghosts.

In the diagonal (b) subcase of Figure 3.10, there's a fairly easy way to deal with the ghosts. In this case, each photon candidate falls in its own BEMC tower: therefore, in almost all cases the two real photon candidates will be the pair that fall in the towers that have more energy. In particular, by removing any diagonal same-pod π^0 candidates whose struck towers have less energy in them than in the other two towers of the pod, we can eliminate the ghosts pretty well. Unfortunately, in the other same-pod cases the coarse-graining of the towers means that this option is not available.

Instead, we have to use the distinctive rectangular shape of the same-pod decay to distinguish ghosts from real candidates. Four of the ghost candidates will be “edge” candidates: that is to say, they will be made from photons that are two adjacent corners of the decay rectangle, or, to put in a way that is easy for an algorithm to check, their decay photons will share a cluster in one of the BSMD planes. Therefore, such a π^0 candidate will make up three sides of the rectangle, with two clusters in one plane and one cluster in the other. We cannot discard this candidate simply on suspicion, though, since it's possible that this is one of those cases in which both photons simply happen to overlap in one of the BSMD planes. Instead, we look to find the fourth side of the rectangle: in practice, this means if the π^0 candidates' decay photons share a cluster in the BSMDP plane, checking to see if there exists a photon candidate made up of one of the π^0 candidate's BSMDE clusters and another BSMDP cluster. If so, we have constructed the full rectangle, and hence in almost all cases the candidate we are dealing with is a ghost candidate. It's possible, of course, that the cluster that should be there in the opposite plane isn't: if one of the photons is fairly low energy, or the BSMD plane doesn't respond as strongly as expected, it may be that the peak strip of the cluster is too low energy and so the cluster is not reconstructed (if the shower deposits its energy in two strips evenly, for instance, you could be left with two strips at roughly 90 ADCs and no reconstructed cluster). To prevent ghosts of this type from slipping by the cuts, a list of peak strips between 50 and 100 ADCs

is maintained, and we check to make sure that such a peak could not be responsible for a cluster that would complete the rectangle.

Unfortunately, even after eliminating the edge ghost candidates, we are left with two π^0 “diagonal” candidates, one of which is real and one of which is not, and there is absolutely no way to distinguishing between them. In this case, we cut the Gordian knot by not trying to distinguish between them. If the candidate is a same-pod different non-diagonal tower decay, or a same-tower decay, the total energy of both diagonal candidates will be the same. Their opening angles will also be essentially the same (bar a small difference due to non-zero vertex z position). And the $z_{\gamma\gamma}$ will also be the same for both candidates: if they don’t fall in the same tower, each has the same struck tower energies (interchanging the two tower energies in the $z_{\gamma\gamma}$ equation doesn’t change anything), and if they do, then thanks to our use of the geometric mean of the two plane $z_{\gamma\gamma}$ ’s to define the same-tower $z_{\gamma\gamma}$, the two candidates will still have the same $z_{\gamma\gamma}$. As both candidates are constructed from the same two BSMDE and BSMDP clusters, they will have the same two BSMDE and BSMDP energies, and so the same $z_{\gamma\gamma}$ in each BSMD plane. Thus we can expect the two diagonal candidates to have the same invariant masses, in which case there is no particular reason to try to figure out which is the ghost: we simply discard one and keep the other.

3.3.3 General Cuts

Once the ghost candidates have been eliminated, some further cuts are applied (see Table 3.3 for the complete list). The first are to the photon candidate’s struck tower(s), which must both have had good status during the run. Additionally, we require that each photon candidate’s struck tower have at least .25 Gev of energy (bearing in mind that a tower’s energy is set to zero if its ADC is less than three σ above pedestal or if it has fewer than 8 ADC).

For the photon candidates themselves, it’s necessary to test to make sure that no

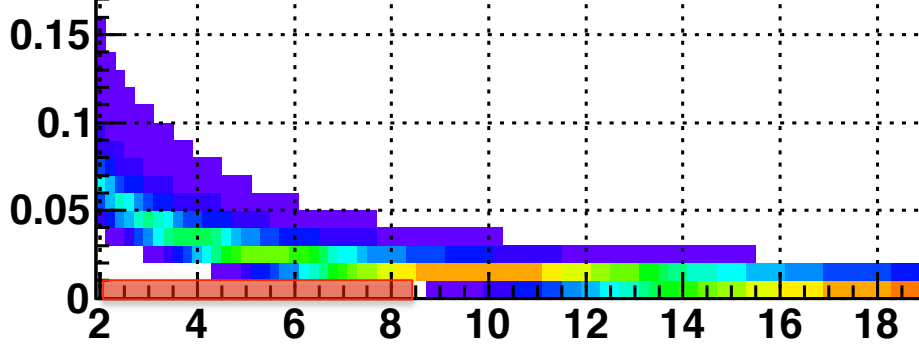


Figure 3-12: *Probability for a pion to decay with some opening angle (actually $\sin(\theta/2)$, which goes into the invariant mass calculation) and p_T (in GeV) combination. The red rectangle in the lower left shows the region that is excluded by the cut.*

photon candidate containing a cluster that is the result of a merger has been combined with a photon candidate that contains one of the clusters that were merged (the parent clusters). Clearly, either the merged cluster is real, or the parent clusters are, but both of them cannot be, and so such a candidate is unphysical. Furthermore, once the π^0 candidates have been constructed we have to decide for each case in which two clusters were merged whether the merger was justified (because, again, candidates that include a merged cluster and candidates that include the merged cluster's parents cannot both be correct). This is done by comparing a π^0 candidate containing a merged cluster with the π^0 candidate(s) that are identical, except that the merged cluster is replaced with one of its parent clusters. If the invariant mass of the candidate containing the merged cluster is closer to .135 GeV than the geometric mean of the invariant masses of the candidates containing the parent clusters, then the candidate with the merged cluster is kept and the others dropped: if not, the candidate with the merged cluster is dropped and the others kept.

We can also use the relationship between energy and opening angle to impose a cut. Figure 3.12 shows a plot of $\sin(\theta/2)$ vs. p_T for simulated π^0 decays: the red rectangle




Figure 3-13: π^0 candidates from data with $p_T > 4$ GeV and decay photons in non-adjacent towers. The η peak is clearly visible on top of the combinatorial background.

in the lower left-hand corner, where $p_T < 8.5$ GeV and $\sin(\theta/2) < .01$, is excluded (the discussion of the split-photon background in the next chapter will have more detail about this cut). Additionally, we can apply a somewhat cruder but quite useful cut using the decay topologies, as sorting by decay topology is similar to, if not exactly the same as, sorting by opening angle. Figure 3.13 plots the probability of a π^0 decaying in each of the four decay topologies vs. p_T . As expected, while non-adjacent tower decays are dominant at low p_T , the probability of this decay mode drops dramatically as p_T increases: in particular, above p_T of 4 GeV, π^0 candidates with this decay topology are almost all background (see the discussion of combinatorial background in the next chapter for more detail), with some η mesons, which can also decay to two photons, mixed in (Figure 3.14 shows the invariant mass distribution for these candidates, with the η peak clearly visible around .5 GeV). On the other hand, same-tower decays, which are dominant at high p_T , drop to nothing at lower p_T : below p_T of 6 GeV, same-tower decays are practically nonexistent, and any such candidate is almost certainly background. Therefore, both categories of candidates can be discarded, although they will still be useful when calculating the final A_{LL} result.




Figure 3-14: π^0 candidates from data with $p_T < 6$ GeV and decay photons in the same tower.

Table 3.3: π^0 Candidate Cuts

Cut	Explanation
Photon struck tower status=1	The struck tower has no electronics or other issues
Photon struck tower energy> .25	There is energy in the tower; no very-low-energy photons
$p_T > 8.5 \parallel \sin(\theta/2) > .01$	Eliminates some split-photon background
$p_T > 6 \parallel$ decay topology> 0	Eliminates some split-photon background
$p_T < 4 \parallel$ decay topology< 3	Eliminates some combinatorial background and all η mesons

THIS PAGE INTENTIONALLY LEFT BLANK

Chapter 4

Measuring A_{LL}

In Chapter 1, A_{LL} was defined as

$$A_{LL} = \frac{\sigma_{++} - \sigma_{+-}}{\sigma_{++} + \sigma_{+-}}$$

From an experimental point of view, though, this definition is not very useful. Experimentalists measure yields (of neutral pions, in this case): to convert the yield to a cross-section requires multiplying it by efficiencies, absolute luminosities, etc. The beauty of an A_{LL} , however, is that it is a ratio. Therefore, almost all of these scale factors, which affect both σ_{++} and σ_{+-} equally, will cancel out of the numerator and denominator. Not only that, but a number of important systematic errors associated with these scale factors will also drop out. Just as with the efficiencies and other such terms, these systematics, such as the BEMC calibration uncertainty, affect both the like- and opposite-sign cross-sections in the same fashion and so do not affect their ratio.

Once all the terms that appear in both σ_{++} and σ_{+-} have been canceled out, we are left with the experimental definition of A_{LL} :

$$A_{LL} = \frac{1}{P_1 P_2} \frac{N^{++} - RN^{+-}}{N^{++} + RN^{+-}}$$

Here, the P_i 's are the polarizations of the two beams (measured per fill); R is the relative luminosity (the ratio of the number of times the beams were collided with the same helicities to the number of times they were collided with opposite helicities); and N^{++} and N^{+-} are the π^0 yields for the same and opposite beam helicities, respectively.

A brief overview of the chapter is in order. The central issue involved in calculating N^{++} and N^{+-} , finding the actual π^0 's, was discussed in the previous chapter, but we still need to determine which events we want to use to look for π^0 's, which will be covered in section 4.1, and how many of the π^0 's we find are actually background events, which will be discussed in section 4.2. Section 4.3 will cover the measurement of the relative luminosity, and section 4.4 that of the polarization. Once all the pieces are in place, we can calculate the A_{LL} : however, we must also check that there are no false asymmetries in place (section 4.5) and check for systematic errors that cannot be dismissed thanks to A_{LL} being a ratio (section 4.6). Finally, everything will be put together in section 4.7.

4.1 Event Selection

Not all of the p-p collisions at $\sqrt{s}=200$ GeV that STAR recorded in 2009 are usable for this analysis. In many cases, a hardware problem renders the events unusable. In some cases, there may have been a problem with one of the detector subsystems which prevented enough information from being recorded to allow us to successfully search for π^0 's. Even if STAR was performing well, the polarization measurement may be bad, or the spin pattern may not have been recorded correctly, making it impossible to use any pions that were found in the A_{LL} calculation.

Furthermore, we only expect a small fraction of the collisions in STAR to actually produce a π^0 in our momentum range of 4 GeV-21 GeV. A particle with such a relatively high p_T will only be produced if there is a hard parton-parton collision, but most proton-proton collisions are glancing or elastic, with no hard scattering and so no high- p_T particles. STAR uses triggers to identify collisions with hard parton scattering, producing datasets




Figure 4-1: *Vertex z distribution of events in our sample.*

which should be enriched in π^0 's when compared to the full set of all 2009 pp200 events.

In addition to selections based on run and trigger, discussed in detail below, there is one additional criterion imposed on events, having to do with the event's vertex position. By extrapolating TPC tracks backwards to a common origin, we can calculate the z position of the vertex. While the beam interaction point is nominally at the center of STAR, the beams are of course not simple points, and proton-proton collisions actually take place over a large range of values of z , as show in Figure 4.1 (the z axis is taken to be along the beamline). The further away the collision is from the center of the detector, the less likely that both decay photons from any π^0 that were produced will actually fall in the detector: therefore we require that the decay vertex fall inside the TPC.

4.1.1 Run Selection

The first stage of the run selection involves performing a QA check on the performance of each detector subsystem. Detector performance is not steady: detectors are liable to electronics failures, high-voltage trips, and other problems. If the performance of one or

more of the key detectors for the BSMD analysis for a given run is sufficiently degraded, the run is discarded. For the BSMD, the key detectors are the TPC, the BEMC, and the BSMD: if a sufficiently large portion of any of those detectors was out of commission, the run is not used.

Additionally, it is necessary that we know the spin pattern, relative luminosity, and polarization for each run in the analysis. The spin pattern should theoretically be recorded for every event: however, if the pattern is off for one event, that may propagate to subsequent events, causing entire runs to be rendered useless. The relative luminosity measurement is also dependent on the spin pattern being recorded correctly: however, it is also necessary that the BBCs be functioning, so if there is a serious problem with the BBCs then the run is not usable. Finally, if there is no polarization value then the run cannot be used. This is mainly a problem for short fills, in which there is not time to make at least two polarization measurements with the beam at flattop. Runs from such fills have to be discarded.

4.1.2 Triggers

A trigger is a set of conditions that an event must meet for it to be written to disk. The simplest possible trigger is a minbias one, which randomly selects some fraction of the total events. However, most minbias events will be physically uninteresting. In order to save disk space and analysis time, triggers are designed to accept only events that we are actually interested in analyzing. Generally speaking, this involves looking for a large deposit of transverse (to the z axis) energy in the calorimeter: lots of energy in the transverse plane means a hard scatter in the original collision (since otherwise all the energy is along the z axis) and an interesting event.

This analysis uses two types of triggers. The first are high-tower triggers, which look for energy deposits from single particles via a large amount of energy in a single tower or a small area around a tower: these are well suited for our purpose, as most π^0 's in our p_T range fall in one tower or two adjacent towers. The BHT3 trigger simply requires that

some tower in the BEMC have at least 6 GeV. The $L2\gamma$ trigger – the “L2” in the name indicates that it is a level two trigger, and so uses a more sophisticated algorithm – requires only 4.5 GeV in the tower, but at least 6 GeV in the 3×3 cluster of towers centering on the tower. Events that pass one of these triggers are quite rare and so, paradoxically, make up the largest part of our dataset, as we are careful to accept all of them that arise.

The second type of trigger used are jet triggers, which look for energy deposits from a number of particles traveling in roughly the same direction, the signature of a jet (see below for a more detailed explanation of jets). The main JP1 trigger looks at the twelve fixed jet patches (from which come the trigger’s name) in the BEMC, each 1.0×1.0 in $\eta \times \phi$ and containing 400 towers, and requires that one have at least 5.4 GeV. There is also a level two jet trigger, L2Jet, which has a more sophisticated sliding jet patch (though it is the same size as the fixed jet patches) and a higher energy threshold (there were actually two energy thresholds, 6.5 GeV and 6.8 GeV, which were used at different times during the 2009 run). Because events that pass the jet trigger are quite common, a prescale is applied to them, meaning that only one out of every n JP1 triggered events is accepted (n varies from 5 to 12, depending on the JP1 trigger rate at different points during the run): therefore jet triggers are outnumbered by high-tower triggers in our dataset, as can be seen in Figure 4.2. While these triggers are not quite as useful as the high-tower triggers, they are still much more likely to yield pions than the minbias triggers. One additional cut is imposed on pion candidates from jet trigger events: they must be able to pass the jet trigger all by themselves, so as to reduce sub-leading pions that are created from jet fragmentation.

4.2 Backgrounds and Fits

Once we have a sample of neutral pion candidates from data events, we need to make sure that we understand the composition of our data. We expect to get three types of candidates, signal candidates (i.e. those corresponding to actual pions) and candidates

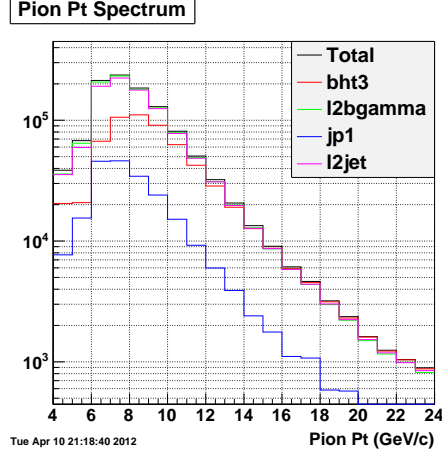


Figure 4-2: p_T distribution of events that passed each trigger (with the overall distribution in black).

from two categories of irreducible background (there will also be η mesons, but since their mass is 547 MeV we do not expect them to contribute at all in our mass window). To verify that this is indeed what we see in the data, we model the invariant mass distribution we expect for all three types of candidates and then attempt to fit the data invariant mass distribution to a sum of the three component distributions (by allowing the normalizations of the latter to vary). This fit will not only confirm our understanding of the data but will also allow us to calculate what fraction of the data is actually background.

4.2.1 Combinatorial Background

The first and largest type of irreducible background is the combinatorial background. This arises from the fact that we have no way to tell whether two photons arose from a π^0 decay or not and so are forced to create pion candidates by combining each photon with every other photon. The result, of course, will be that numerous pion candidates will be formed which consist of photons from one π^0 combined with photons from a different π^0 or even with photons that have nothing to do with any π^0 .

We can model this background of pion candidates formed from uncorrelated photons using a data-driven process known as event mixing. In order to model a background from combining uncorrelated photons, we need groups of uncorrelated photons to combine. While it is not possible to determine whether or not photons are correlated within events, it's quite clear that photons from two different events cannot possibly be correlated. Therefore we can model the combinatorial background by combining photons from one event with photons from another.

It's not quite that simple, though. First, we need to ensure that the events have vertex z positions that are sufficiently close to another. Combining a photon from a vertex at $z=50$ cm with one from a vertex at $z=-50$ cm is just not feasible due to the impossibility of determining the opening angle of the decay. Therefore, events are only combined if their vertices are within 10 cm. Additionally, it is not entirely true to say that the combinatorial background arises from uncorrelated photons. The presence of jets in proton-proton collisions introduces a degree of correlation that must be accounted for.

A jet is a group of particles traveling in roughly the same direction within a reasonably confined space (STAR's jet-finding algorithm, for instance, uses a radius of .7 in $\eta-\phi$ for its jet cone). Jets are a common product of hard parton-parton scatterings. As was explained in Chapter 1, one of the fundamental properties of QCD is that quarks and gluons can never be observed by themselves, but are always confined within a hadron. As the parton flies away from the collision, then, it attempts to hadronize by pulling quark-antiquark pairs out of the vacuum. The result is that the parton fragments into a spray of hadrons traveling outwards in a cone: i.e., instead of a quark or gluon, the detector sees a jet.

The important point for our purposes here is that the jet structure of an event clearly introduces a degree of correlation. If two π^0 's are both part of the same jet (or one is part of the jet from one of the scattered partons and the other is part of the jet from the other scattered parton) then we cannot say that their decay photons are entirely uncorrelated. In order to account for this, we rotate events before they are mixed so that their jet axes




Figure 4-3: *Illustration of the event rotating procedure used in creating combinatoric background events whose underlying structure is similar to that of actual events.*

line up (where the jet axis is the axis on which the two partons from the hard scattering left the collision). Because rotating STAR in η amounts to shifting it along the beam axis, the more we rotate the less realistic the result will be: if we rotate by η of .5, for instance, fully a quarter of the rotated event will fall outside of the actual detector. Therefore events are only rotated if their jet axes are within .1 of each other in η . Rotation in ϕ simply amounts to rotating STAR about the z axis so it presents no problem. Figure 4.3 illustrates the result of rotating two events to align their jet axes [19].

Once the rotation has been done, photons from one event can be combined with rotated photons from the other event to form pion candidates. An extra cut needs to be applied here to ensure that the photons don't get too close together. In a real event, two photons that were within two strip widths in both planes of the BSMD would appear to be a single photon: therefore, we must reject any combinatorial background candidates whose photons are that close together. Once this cut is applied, the combinatorial background candidates are treated in exactly the same fashion as the data candidates and run through the exact

same cuts. The result is an invariant mass distribution that, as we will see, matches the data quite well over the range where we expect the data to consist entirely of combinatorial background events.

4.2.2 Split-Photon Background

The second type of irreducible background we expect to see arises from cases in which the two photons that form the pion candidate are quite closely correlated, in that they are actually the result of a single photon being split in two by the BSMD. As was discussed in the section on BSMD clustering, clusters can be split in two main ways: either by dead strips, or by secondary strips that fluctuate high (Figure 4.4 provides a cartoon of both situations). The procedures for merging clusters attempt to compensate for these problems, but they will not always be successful. For instance, if the secondary peak occurs two strips away from the main peak, no merger will be made. Alternatively, in some cases fluctuations in the value of $z_{\gamma\gamma}$ – which will almost always be calculated using BSMD energies, as both fake photons are likely to be in the same tower – could cause the merging hypothesis to be rejected, even though it is true. The angle- p_T cut described in Section 3.3.3. will eliminate many of the split-photon candidates, especially the ones that are split by a dead strip, as they are often low energy – they are missing the energy of one photon – and low-angle – the fake photons will be closer together than the real ones – but some will survive. Figure 4.5 shows the result of applying the opening angle- p_T cut to both the data (in black) and the simulated split-photon background (in blue): the low-mass peak in the data disappears, but this is just due to a significant reduction in the amount of split-photon background. Note that the split-photon background has been scaled by a number that emerges from the overall data fit, as described below.

This background is modeled using single-photon simulations. Using GEANT, we can generate events consisting only of one photon (and a few muons to provide a vertex). Then these events are put through the pion finder: the resulting π^0 candidates are used to model




Figure 4-4: *Illustration of two possible ways in which a single cluster can be reconstructed as two separate clusters. The green strips and red strips represent different reconstructed clusters.*

the split-photon background [21].

4.2.3 Simulating the π^0 peak

In addition to modeling our backgrounds, we also need to model the signal distribution. At first glance, the obvious way to do this would be by using simulated single-pion events. However, it turns out that this is not possible due to the movement of the pion mass peak. In theory, the invariant mass distribution of the π^0 candidates should peak at .135 GeV in all p_T bins, but in practice the peak location increases fairly steadily with p_T . This is partially due to detector and algorithm effects. The detector's finite minimum position resolution is one cause: since clusters must be separated by at least one BSMD strip width, pion reconstruction is biased towards larger opening angles, and at higher p_T values, where the opening angles will tend to be smaller, this will tend to increase the mass peak. The BEMC energy resolution will also tend to increase the mass. Because the π^0 spectrum is steeply falling, it is more significant when a π^0 's energy is measured to be larger than it is than if the energy is measured to be smaller than it is, and if the measured energy is too high while the other parameters stay the same, the mass will be too large. These factors should be the same in single-particle simulations as in full simulations or in data, but that

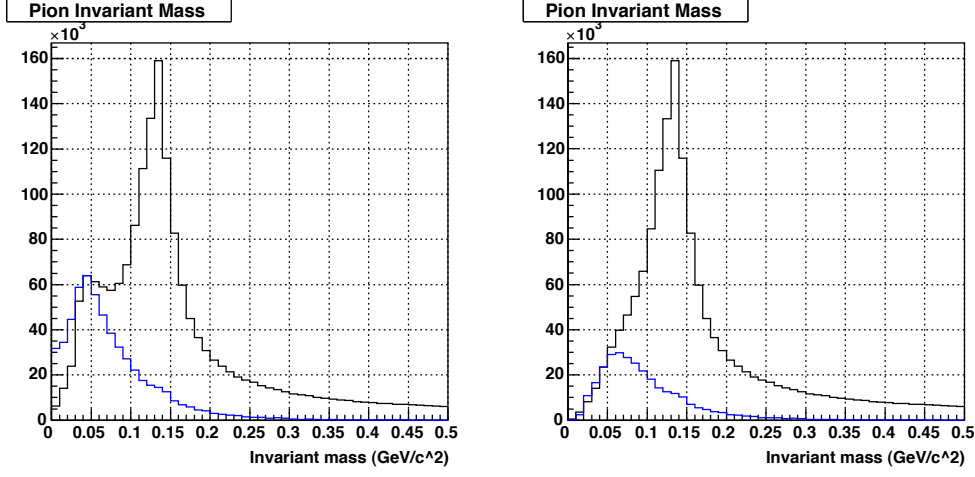


Figure 4-5: *Illustration of the effect of applying the cut in p_T -opening angle space (here for π^0 candidates with $4 < p_T < 6.5$). The blue is the split-photon background, the black is the data. The split-photon background has been normalized so as to produce the best possible data-peak+background fit.*

is not the case for the final contributor to the fact that the mass peak increases with p_T , jet background. The kind of hard scattering events that produce high- p_T pions are likely to have greater total E_T than lower-energy events. This means that there will be more particles created in such an event, and so a greater likelihood that low-energy particles will deposit part or all of their energy in the same tower(s) as our pion candidate, increasing its energy and hence its mass. This phenomenon should be seen in full simulations of proton-proton collisions, but it will not occur in single-pion simulations, and hence we should expect that the pion mass peak in the single-pion simulations will be lower than it is in data or full simulations. Figure 4.6 shows that this is in fact the case: while the data and full simulation mass peaks match each other quite well over a large p_T range, the single-pion simulation mass peak is smaller and growing more slowly.

Therefore, we must simulate the pion signal distribution using full simulations. But we clearly cannot simply run these simulations through the pion finder and take the result as our signal: the full simulations should contain the exact same backgrounds as the data.




Figure 4-6: Location of the invariant mass peak (obtained from a gaussian fit) vs. p_T for data (black), full simulations (blue), and single-pion simulations (red).

Instead, we need some way to determine which of the π^0 candidates we find are true pions. Since in finding true reconstructed pions we are more concerned with purity (our signal distribution must contain very little background, since any background present will disturb the fit) than efficiency (we can afford to throw out pions if we're not sure about them), we can use a relatively crude method of identifying true reconstructed pions based on the fact that in simulation we know everything about every particle produced in the event. For each reconstructed pion candidate, we can calculate the distance between its decay photons and the decay photons of each simulated pion from the simulated collision (taking the pair of distances that has the smaller sum to be the correct one), and find the simulated pion that the reconstructed pion is closest to. If, for this closest pion, the two decay photon distances are both less than 0.02 (in $\eta - \phi$) then this pion is taken to be a signal pion.

4.2.4 Fitting the Data to Simulation+Backgrounds

Finally, once we have invariant mass distributions for all three pieces of the data, we can fit the data to them by allowing their normalization to vary. The result is shown for

each p_T bin in Figure 4.7: the data are in black, the simulated pion peak in purple, the combinatorial background in green, the split-photon background in blue, and the fit in red. Figure 4.8 plots (data-fit)/data, giving a better idea of the fit quality. While the fit describes the data fairly well, there are clearly significant divergences from the data on the shoulders of the peak, where the data is greater than the result of the fit. Because this discrepancy on the shoulders of the π^0 peak was also seen in the 2006 result, which used an entirely different algorithm, we can be fairly confident that it is not an artifact of the algorithm. Further investigation indicated that the problem does not lie only in the simulation of the low-energy background, but persists in the simulation of the high-energy component of the event. This strongly suggests that the issue has to do with the way that the BSMD is simulated in GEANT. As we will use the results of the fit to calculate the fraction of events that fall into the invariant mass region of interest that are actually background, which will be important in determining the final result, it is necessary that the data-simulation discrepancy be accounted for somehow. We will do so with an additional systematic uncertainty: see the discussion of the fit systematic in Section 4.6.2.

4.3 Relative Luminosity

The relative luminosity R is defined as $\frac{L^{++}}{L^{+-}}$, where the L 's are the measured luminosity for like-sign and opposite-sign helicity states, respectively. R is determined by the BBC's, which count the number of coincident signals (i.e., signals in both BBC's) and combine that with the preset spin pattern (which is set as the bunches are injected into the accelerator) to determine the numbers N^{++} , N^{+-} , N^{-+} , and N^{--} of each possible spin configuration. A number of possible ratios can be constructed from these four quantities: the one we want is

$$R_3 = \frac{N^{++} + N^{--}}{N^{+-} + N^{-+}}$$

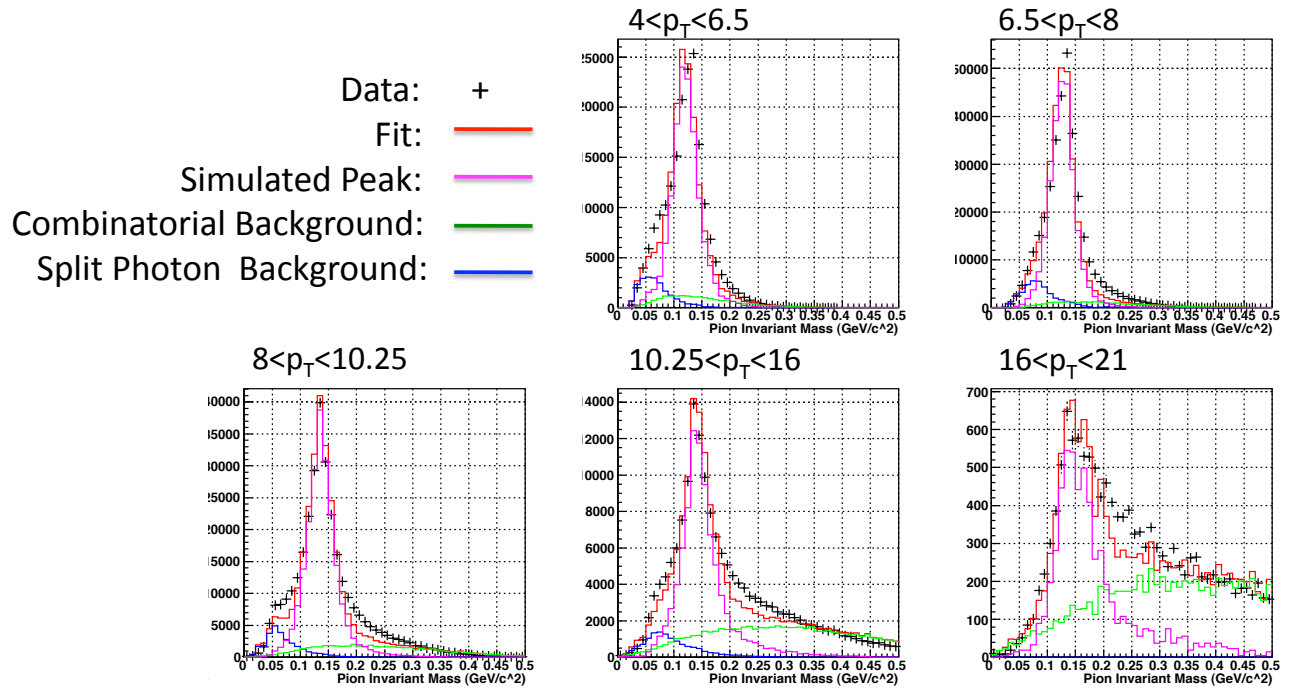


Figure 4-7: For each p_T bin, the data invariant mass distribution (black) is fit to a sum of the simulated pion peak (purple), the split-photon background (blue), and the combinatorial background (green): the fit is in red.

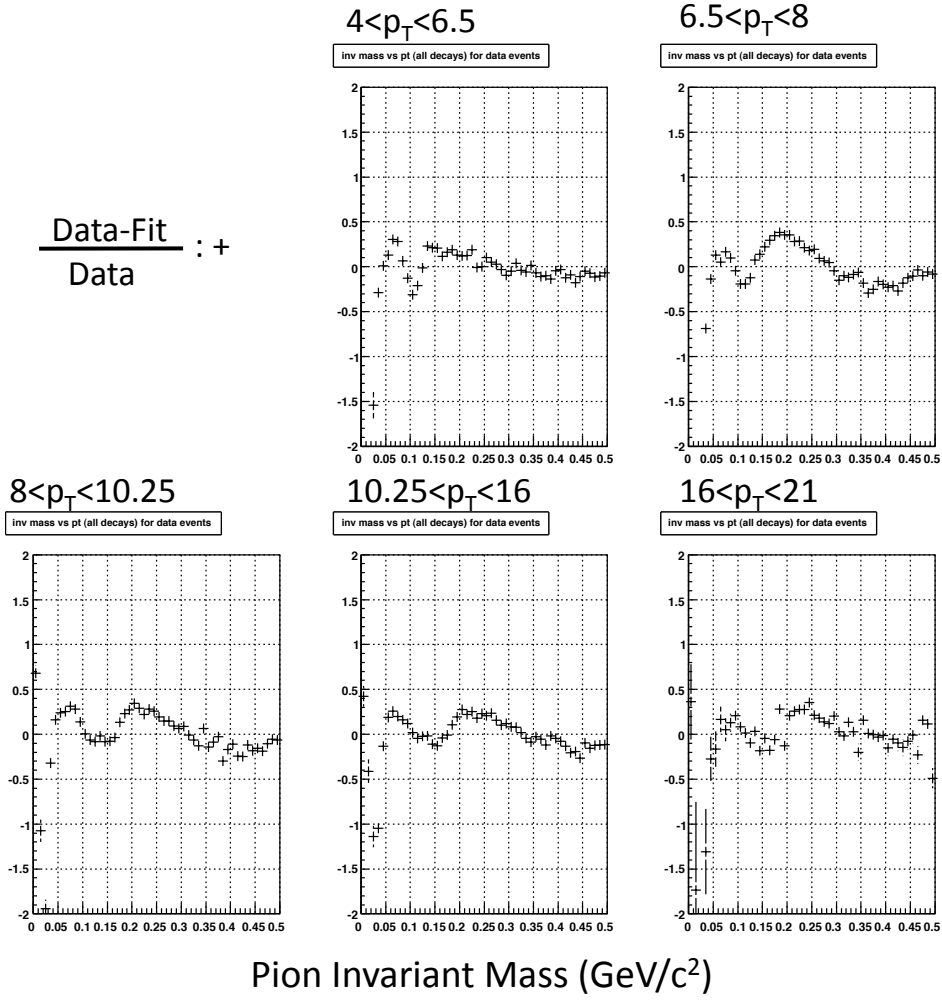


Figure 4-8: $(\text{Data-Fit})/\text{Data}$ as a function of invariant mass for each p_T bin.




Figure 4-9: *Polarization for both beams plotted by fill.*

It is calculated for every fill.

4.4 Polarization

The RHIC CNI and H-jet polarimeters were described in detail in Chapter 2. The CNI polarimeters measure the polarization of both beams at the beginning of each fill and every two hours or so during it. An average polarization is calculated for the entire fill and is used for each run in the fill.

4.5 Single-Spin Asymmetries

Before calculating A_{LL} , we must make a basic crosscheck by calculating the single-spin asymmetries A_L . These asymmetries are defined for the yellow and blue beams as

$$A_L^{Y(B)} = \frac{1}{P_{Y(B)}} \frac{N^+ - R_{Y(B)} N^-}{N^+ + R_{Y(B)} N^-},$$

where N^+ and N^- are the yields for the yellow or blue beam plus and minus helicity configurations (averaging over the helicity configuration of the other beam) and the R 's are relative luminosities (note that these are not the R_3 that was defined in section 4.3 but are instead the ratios of the positive to negative helicity configurations for the blue and yellow beams). These asymmetries are parity-violating, as under the action of the parity operator N^+ becomes N^- and vice versa, and since QCD is not parity violating, they must be zero (it's important to note here that A_{LL} is not parity violating, as the parity operator just transforms the two components of N^{++} and N^{+-} into each other). If they are not zero, that would suggest the presence of some outside source of asymmetry, which could also contaminate the A_{LL} measurement.

Rather than calculating a single A_L for each beam for the entire dataset, we calculate the two single-spin asymmetries for each run, which would allow us to see if any time-varying asymmetry was present. The results are shown vs. run index (i.e. 1 is the first run used in the analysis, 2 the second, etc.) in Figure 4.10, with the blue beam on the left and the yellow beam on the right. Also shown on the plot is the result of fitting the asymmetries to a constant value: that value, the statistical error, and the χ^2/NDF of the fit are shown on the plot. For both beams, the fit is consistent with zero, with a χ^2/NDF close to one.

4.6 Uncertainties

In terms of the quantities that we actually measure, A_{LL} (in a p_T bin j) is defined as

$$A_{LL,j} = \frac{\sum_I P_{B,i} P_{Y,i} (N_{i,j}^{++} + N_{i,j}^{--} - R_i (N_{i,j}^{+-} + N_{i,j}^{-+}))}{\sum_I P_{B,i}^2 P_{Y,i}^2 (N_{i,j}^{++} + N_{i,j}^{--} + R_i (N_{i,j}^{+-} + N_{i,j}^{-+}))}$$

where the sum over i is the sum over the runs. This quantity has both statistical and systematic uncertainties: we will deal with each in turn.




Figure 4-10: *Single-spin asymmetries plotted vs. run index and fit to a constant: left, the yellow beam; right, the blue beam.*

4.6.1 Statistical Uncertainties

The statistical uncertainty for A_{LL} as defined above is given by, if D_j and U_j are the denominator and numerator respectively for A_{LL} in bin j ,

$$\sigma_{A_{LL},j}^2 = \sum_i \left(\frac{D_j P_{B,i} P_{Y,i} - U_j P_{B,i}^2 P_{Y,i}}{D_j^2} \right)^2 (N_{i,j}^{++} + N_{i,j}^{--} + R_i^2 (N_{i,j}^{+-} N_{i,j}^{-+})).$$

But in the case where the asymmetry is small (as it is here), the numerator term can be neglected and the error simplifies to

$$\sigma_{A_{LL},j}^2 = \sum_i \left(\frac{P_{B,i}^2 P_{Y,i}^2 (N_{i,j}^{++} + N_{i,j}^{--} + R_i^2 (N_{i,j}^{+-} N_{i,j}^{-+}))}{D_j^2} \right).$$

In this form, the ROOT software package, which we use to calculate the final answer, will propagate it correctly.

There is one further important point: we must take into account the presence of events

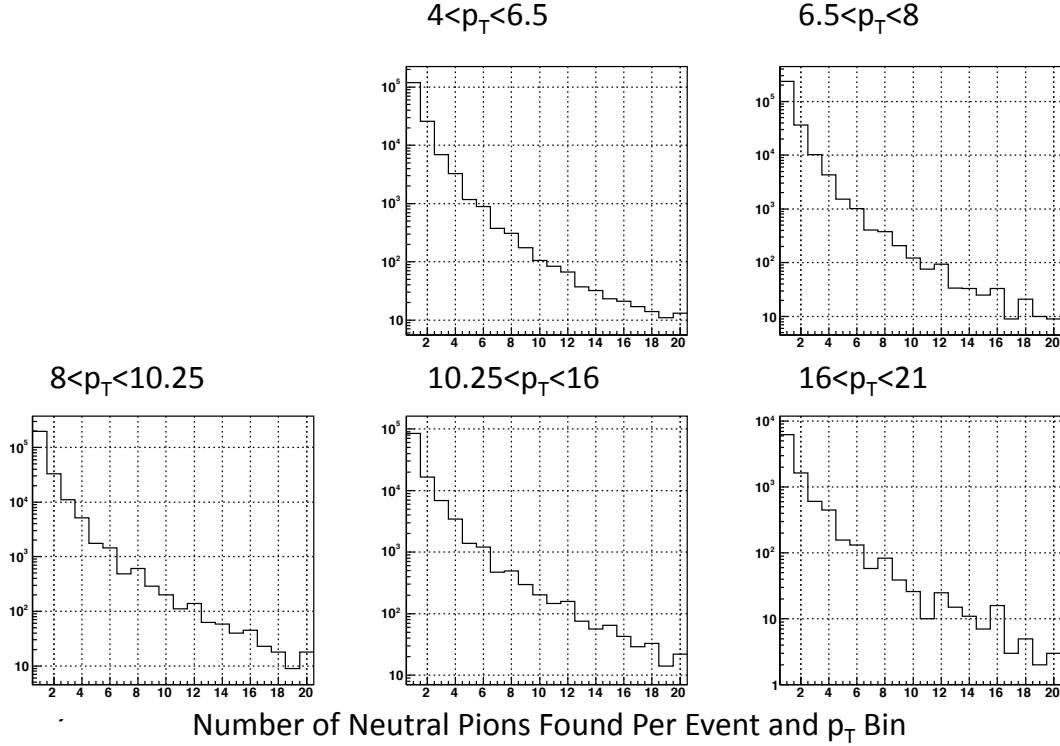


Figure 4-11: *Number of pions per event, by p_T bin.*

in which we find more than one pion in a given p_T bin (see Figure 4.11). Under these circumstances, to naively treat each pion as an independent event and calculate the resulting uncertainties using Poisson statistics as \sqrt{N} would be incorrect. For this analysis, we follow [30] and increment each bin in our histograms only once per event, using a weight equal to the number of pions in that bin in that event.

4.6.2 Systematic Uncertainties

There are four main systematic uncertainties to deal with. We expect the systematics to be less significant than the statistical uncertainty, with the exception of the fit (which could also be described as data-simulation) systematic.

Trigger Bias

As discussed above, the data used in this analysis all comes from one of four triggers: therefore, we must be alert to possible biases introduced by relying on these triggers. For instance, events that result in a mid-rapidity jet which hadronizes with a strong electromagnetic component – i.e. events that pass the JP1 trigger – are more likely to be events that contain a quark jet, since quark jets tend towards harder fragmentation than gluon jets. This means that these events are more likely to have a quark in the hard scattering process, and since quarks are more likely to be polarized than gluons, this may introduce a bias into our A_{LL} measurement. In order to estimate the size of the bias from this or any other trigger effect, we use the Method of Asymmetry Weights [26].

The method involves making a leading-order Monte Carlo simulation of A_{LL} using the Pythia event generator. We can use Pythia to determine the partonic a_{LL} (the ratio of the polarized to unpolarized hard scattering cross-sections) of the process and then determine an asymmetry weight for the event by sampling parton distribution functions as follows:

$$w = \frac{\Delta f_1(x_1, Q^2) * \Delta f_2(x_2, Q^2) * a_{LL}}{f_1(x_1, Q^2) * f_2(x_2, Q^2)}.$$

The Monte Carlo A_{LL} is then just the ratio of the asymmetry-weighted and unweighted distributions. To determine the systematic uncertainty, we simply take the difference between the Method of Asymmetry Weights A_{LL} for all events and that for just those events that pass a simulated trigger condition.

There are a few additional subtleties here. First, while the partonic a_{LL} is calculable and the quark PDFs, both polarized and unpolarized, and the unpolarized gluon PDF have been measured reasonably well, the polarized gluon PDF has not been, and in fact constraining it is the goal of this analysis. Rather than calculating just one pair of A_{LL} 's, then, we calculate three, each with a different version of $\Delta g(x, Q^2)$. Based on past analyses, it seems unlikely that the absolute value of ΔG is greater than .3, so we use two scenarios

of $\Delta g(x, Q^2)$ that give values of $\Delta G = .3$ and $\Delta G = -.3$ as an envelope (as calculated by GRSV). Additionally, we calculate one pair of A_{LL} 's based on the DSSV global fit value of $\Delta g(x, Q^2)$, that being the best global fit currently available. Then we take as the systematic the largest value of the triggered-untriggered difference, regardless of which $\Delta g(x, Q^2)$ scenario it came from.

Additionally, we have to be careful to ensure that the trigger mix in the simulations matches that of the data. As discussed above, JP1 && !BHT3 triggered events, being the most common, are heavily suppressed with a prescale during data-taking. No such suppression is performed when events are being generated in simulation, though: as a result, while such events make up only about 10% of our dataset, they are nearly half of our simulation sample. In order to ensure that we are measuring the bias introduced by the correct group of triggers, we weight the JP1 && !BHT3 events in the simulation so that their ratio to the total size of the simulation sample matches the ratio seen in the data.

Finally, it may be the case, especially in the higher p_T bins, that the statistical errors on the A_{LL} difference are larger than the difference itself. In such a case, the value of the statistical error will be taken as the systematic. As the trigger bias systematic uncertainty turns out to be about an order of magnitude less than the statistical uncertainty, such an approximation doesn't really effect the final uncertainties.

Figure 4.12 shows the results of this calculation: the A_{LL} differences and the statistical error on the differences are shown in each p_T bin for each $\Delta g(x, Q^2)$ scenario. The numerical values of the systematic are given in Table 4.2 below along with the other uncertainties.

Fit Systematic

The most important potential source of uncertainty in the measurement is that arising from the problems with the fit of the data invariant mass distribution to the sum of the two backgrounds and the simulated signal peak. Figure 4.8, which plots the difference




Figure 4-12: *Difference between triggered and untriggered A_{LL} ’s, with the latter calculated from the method of asymmetry weights using different $\Delta g(x, Q^2)$ scenarios, for each p_T bin. The $\Delta G = -.3$ scenario is in black, the $\Delta G = .3$ scenario is in red, and the DSSV scenario is in blue.*

between the data and the fit normalized to the data, clearly shows structure in the data-fit difference. In particular, the fit consistently comes in below the data on the shoulders of the pion peak across all the p_T bins. The failure is most pronounced in the general region from .15-.25 GeV.

As the data and fit match each other well at higher invariant mass where the combinatorial background dominates, and the split-photon background is too small a contributor to be responsible for this effect, the fit problems must be due to our failure to correctly simulate the π^0 peak. We can test this hypothesis by comparing the data directly to simulation, as the full simulation should have the same backgrounds as data present. And to ensure the best possible modeling of the backgrounds, we use “embedding” simulations. Embedding works by taking low- p_T background from data and then adding to that our simulations of whatever high- p_T process we’re interested in (here, jet production, which should produce plenty of π^0 ’s, is used). Thus the interesting process is embedded in a

data-driven background, which should give us a better grasp of the backgrounds. Figure 4.13 plots the data-embedding difference normalized to the data: as in the previous section, JP1 & BHT3 triggers have been reweighted so that the embedding trigger ratios are the same as those in data. Additionally, since there is no fit to determine the normalization of the embedding, it is scaled by a factor which equalizes the data and embedding integrals over the invariant mass range .4-.5 GeV (i.e., a range where the combinatorial background is dominant). The error bars are large due to the relatively smaller size of the embedding sample (aggravated by the reweighting process, which, by giving half the sample the weight of one-tenth of the sample, effectively cuts our statistics in half), but structures are clearly visible around the pion mass peak, strongly suggesting that the discrepancy is due to a fundamental issue with the simulation.

This data-simulation mismatch prevents us from obtaining a complete understanding of the data, but what does this mean for the actual A_{LL} calculation? The answer lies in how we determine the number of pions in a given p_T bin, or more accurately how we don't determine that number. As our backgrounds are irreducible, there is no way to decide whether a given pion candidate is a genuine pion or merely background; instead, we count the total number of pions and then use the fit to determine what fraction of them likely arise from our two backgrounds (and then determine the extent to which our backgrounds dilute the final A_{LL} , as will be covered in detail later). On the other hand, a pion candidate with an invariant mass of 1 GeV is very unlikely to be an actual pion, so there is no point in counting all the pion candidates in a given p_T bin. Instead, we pick an invariant mass window and count all the candidates in that window. In this analysis, we use a mass window of .08-.25 GeV.

In the absolute worst-case scenario, the data-simulation mismatch indicates the presence of a previously unsuspected and strongly asymmetric background (or backgrounds), concentrated on the shoulders of the pion peak. But in that case, our A_{LL} should be strongly dependent on what mass window we use. A mass window more tightly focused

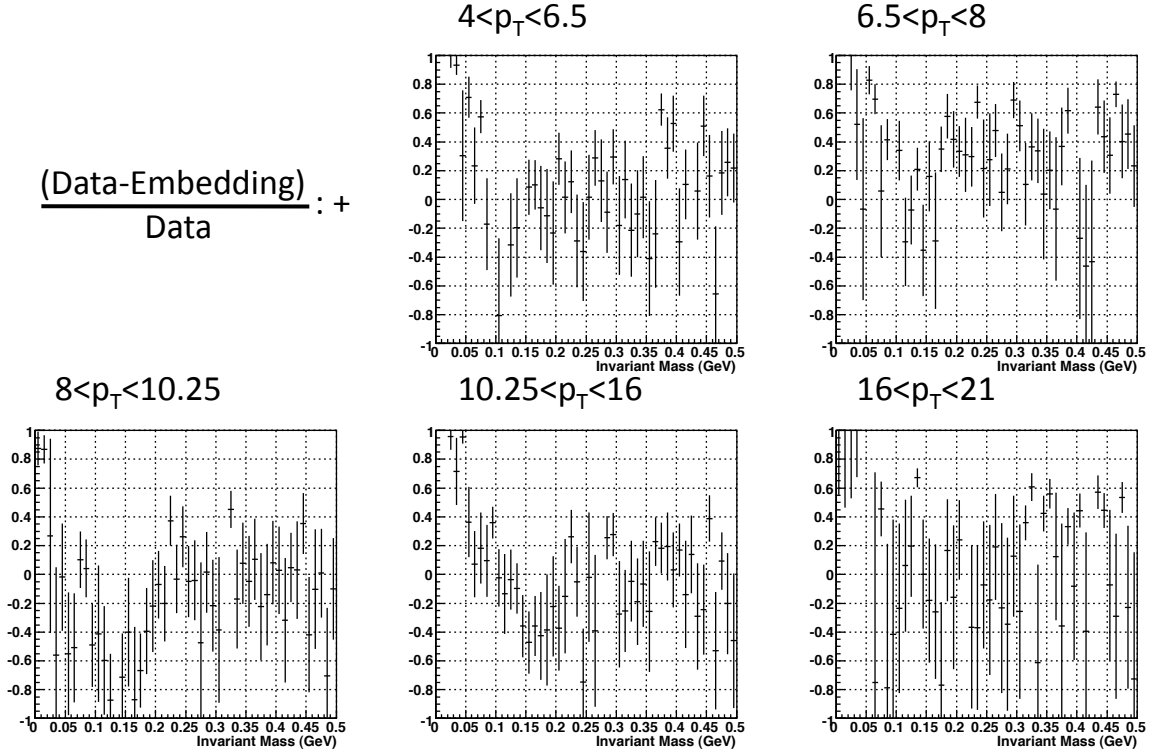


Figure 4-13: $(\text{Data-Embedding})/\text{Data}$ for each p_T bin. The embedding is scaled separately in each p_T bin such that the integrals of the data and the embedding over the range $.4 < m_{inv} < .5$ are the same.

about the pion mass, excluding most or all of the shoulder area, would produce an A_{LL} that was different from that we get from the mass window we actually use. This suggests a method for calculating a systematic uncertainty to describe this mismatch (which is most likely simply due to the BSMD simulator not matching reality: as the BSMD is the key detector for only a few analyses, less attention has been paid to its simulation than to other detectors). By calculating A_{LL} with various mass windows, we can see to what extent it depends on the mass window used and so to what extent the failure of the simulation to describe the data limits our understanding. The mass windows we use are one that is considerably tighter around the pion peak, running from .1 to .18, and the two made by taking one end from either the original or new window and pairing it with the other end of the other window. These are shown on Figure 4.14, which reproduces Figure 4.7 but adds lines to indicate the original and new mass windows (and the π^0 mass itself) so that what each window includes and excludes is clear. The results of the A_{LL} calculation for all four mass windows are shown in Figure 4.15. The systematic uncertainty is taken to be the largest difference between the A_{LL} using the original window and the A_{LL} 's from the new windows. The value of the uncertainty for each p_T bin is shown in Table 4.2.

Transverse Component

It's possible that there may be some net transverse polarization remaining in the beam even after the spin rotators have given us a longitudinally polarized beam. Such a transverse component to the beam might induce a transverse double spin asymmetry, A_Σ , which could affect our A_{LL} measurement. Due to a lack of statistics, determining A_Σ from neutral pion production is not feasible. Instead, we use the value of A_Σ measured from the 2009 pp200 inclusive jet production analysis. The transverse beam component turns out not to contribute very much to the overall uncertainty: its values in each p_T bin are shown in Table 4.2 below.

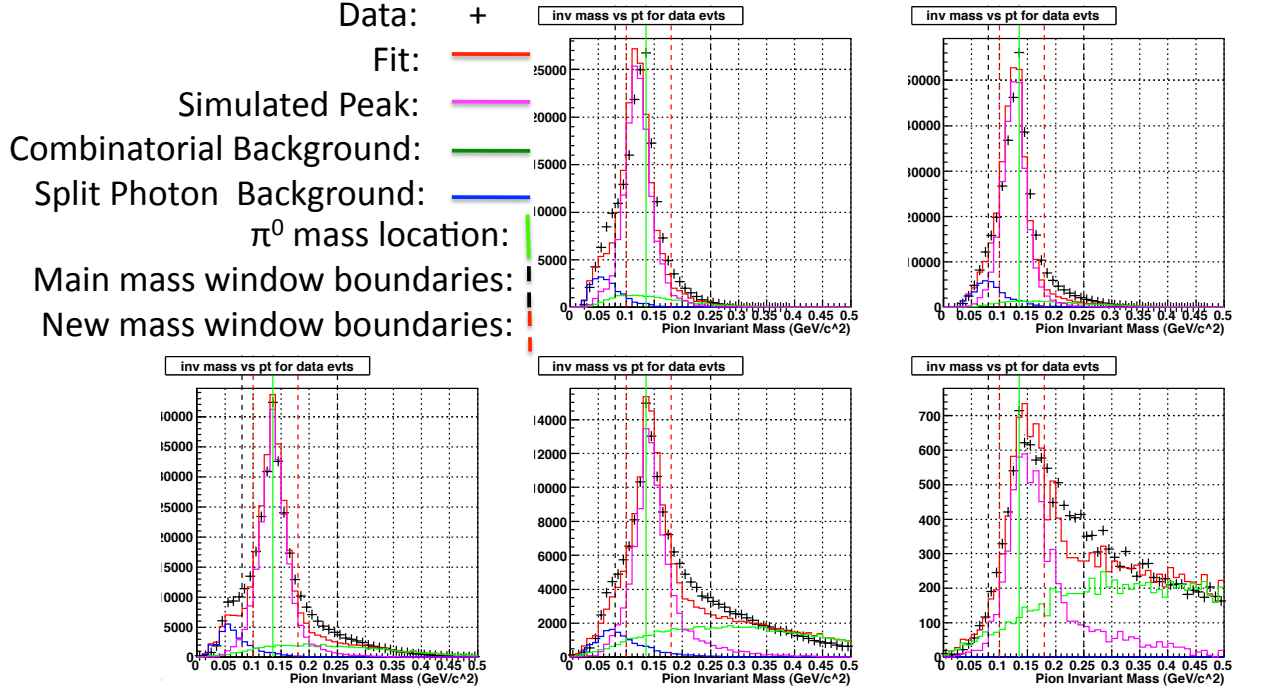


Figure 4-14: Mass windows, shown with data and simulation/background distributions for each p_T bin. The black dotted lines are the original mass window boundaries, the red dotted lines show the new mass window boundaries, and the green line indicates the location of the π^0 mass peak.




Figure 4-15: A_{LL} vs p_T for each mass window. Black is the original mass window, $.08 < m_{inv} < .25$; red is $.1 < m_{inv} < .25$; green is $.08 < m_{inv} < .18$; and blue is $.1 < m_{inv} < .18$.




Figure 4-16: A_{LL} with relative luminosity from the BBC (in black) and ZDC (in red).

Relative Luminosity

The final systematic uncertainty we have to deal with arises from our measurement of the relative luminosity. As previously explained, the main relative luminosity measurement comes from the BBC. However, we do have another detector which can also be used for relative luminosity measurements, the ZDC. Because the ZDC has lower rates than the BBC, the BBC relative luminosity measurement is preferred for the actual A_{LL} measurement, but we can use the ZDC relative luminosity to determine a systematic uncertainty, simply by recalculating the A_{LL} with the entirely independent value of R measured from the ZDC instead of that from the BBC. Due to problems with the west ZDC in 2009, the R used here is actually just from hits in the east ZDC, rather than from ZDC coincidences. Figure 4.16 shows the BBC and ZDC A_{LL} 's: the differences are too small to be seen on the scale of this plot, but are given with the rest of the uncertainties in Table 4.2.

4.7 Results

The formula for A_{LL} was given at the beginning of the chapter. However, this formula assumes that every pion we reconstruct is a true pion, which we know is not the case. Since we can't actually separate out the background from the signal, we must account for it in

Table 4.1: *Background Fractions*

p_T bin	Split-Photon Background Fraction	Combinatorial Background Fraction
4-6.5	.038	.078
6.5-8	.061	.06
8-10.25	.034	.11
10.25-16	.061	.2
16-21	0	.3

some other fashion. This is done by calculating an observed A_{LL} and then determining the true A_{LL} by taking into account the extent to which it is diluted by the backgrounds [21]. The first step is to calculate an A_{LL} for the backgrounds: this is done by using mass windows in which the backgrounds should dominate. For the combinatorial background, we use the range $1.2 \leq m_{inv} \leq 2$ GeV; for the split-photon background, $0 \leq m_{inv} \leq .06$ GeV. In calculating these background A_{LL} 's, we integrate over p_T : since both backgrounds arise from our algorithm or the detector, rather than physics, their asymmetry (should there be one) should not be p_T dependent. This also allows us to use two categories of pion candidates that were previously discarded, those of decay type 0 and $p_T < 6$ GeV and those of decay type 3 and $p_T > 4$ GeV. Both categories are dominated by background, the former by split-photon background and the latter by combinatorial background, and including them in the calculation of the background A_{LL} allows for increased statistics and so lower statistical uncertainties, helping to drive down our final uncertainties.

We must also calculate a background fraction for each background type and each p_T bin. This is done using the overall fit to the data, as discussed previously: each background distribution is integrated over the mass window using the normalization from the fit, and then this integral is divided by the integral of the data over the mass window. The background fractions for the primary mass window are shown in Table 4.1. Background fractions are also calculated for each of the three mass windows used in determining the data-simulation mismatch systematic.

Table 4.2: A_{LL} Uncertainties

p_T bin	Systematic Uncertainties				Statistical
	Trigger Bias	Fit	Relative Luminosity	Transverse Component	
4-6.5	.00028	.0069	.00014	.00025	.015
6.5-8	.00059	.0024	.00017	.00025	.01
8-10.25	.00097	.0078	.00033	.00025	.012
10.25-16	.0017	.016	.00017	.00025	.022
16-21	.0068	.021	.00018	.00025	.11

With these numbers, we are now prepared to calculate the true A_{LL} from the formula

$$A_{LL}^{measured} = (1 - f_{sp} - f_{comb}) * A_{LL}^{true} + f_{sp} * A_{LL}^{sp} + f_{comb} * A_{LL}^{comb},$$

where the f's are the background fractions.

Table 4.2 shows the uncertainties, both statistical and systematic, in each p_T bin; Table 4.3 shows the values of A_{LL} in each p_T bin, along with the statistical uncertainty, the total systematic uncertainty (obtained by summing the systematics in quadrature), and the total uncertainty (obtained by summing the statistical and total systematic uncertainties in quadrature). Figure 4.17 shows the final result over the full p_T range: the points are located in the center of each bin. The black error bars shown are statistical, with the green bands indicating the extent of the systematic uncertainties.

Table 4.3: A_{LL} Results

p_T bin	A_{LL}	Statistical Uncertainty	Total Systematic Uncertainty	Total Uncertainty
4-6.5	.0094	.015	.0069	.017
6.5-8	.0084	.01	.0024	.01
8-10.25	.022	.012	.0078	.014
10.25-16	.0046	.022	.016	.027
16-21	-.073	.11	.021	.11




Figure 4-17: $\pi^0 A_{LL}$ results, with systematic errors shown as green bands.

Chapter 5

Interpretation and Conclusion

To determine how meaningful our A_{LL} result is, we need to compare it both to theoretical predictions, to see how well it can discriminate among them, and to other A_{LL} measurements, to see how consistent it is with what is already known. Figure 5.1 shows this result (triangles and red error bars, statistical errors only) compared to the 2006 STAR π^0 A_{LL} result and to several theory curves [21], while Figure 5.2 shows the 2005 STAR π^0 result [8]. Note that different trigger values in different years affect the error bars in the lowest p_T bin. In 2006, the BHT3 trigger was at 5.2 GeV, and the L2 γ trigger required 3.8 GeV in a tower and 5.2 GeV in the cluster containing it; in 2009, those numbers were increased to 6, 4.5, and 6 respectively. The result is that the 2009 data is unable to improve significantly on the 2006 result in the 4-6 GeV region, as it is largely dependent on jet triggers which are both less efficient at finding high- p_T π^0 s than the high-tower triggers and are less represented in the data.

First of all, we note that the 2009 result is consistent with both the 2005 and 2006 results: there are no sudden changes in the value of the A_{LL} which might suggest a problem with this analysis. This means that, because of our improved error bars, we can be even firmer in rejecting the maximal gluon polarization scenario (GRSV ΔG Max) (the new data point at high p_T also helps here) [21]. Additionally, we can compare our A_{LL} result




Figure 5-1: Comparison of the 2009 $A_{LL} \pi^0$ result to the 2006 result and several theoretical predictions and global fits.




Figure 5-2: The published STAR neutral pion A_{LL} result using data taken in 2005.

to the DSSV and GRSV-STD global fits. At low p_T , our result is consistently above DSSV and falls closer to GRSV-STD; at higher p_T , by contrast, we are closer to DSSV. While the error bars are too large to make firm conclusions here, there is a clear trend.

The trend becomes even more significant when compared with other A_{LL} results, both from other experiments and using other observables. To start with, we can compare to the π^0 A_{LL} result from the PHENIX experiment, the other large physics experiment at RHIC. As you can see in Figure 5.3 [7], the PHENIX and STAR results are complementary: PHENIX does best at low p_T values that STAR cannot reach, while STAR provides access to the high- p_T region which PHENIX has difficulty reaching. Note that the STAR and PHENIX results are consistent in the overlap region. And again, at low p_T , though not at the lowest values that PHENIX can probe, the PHENIX result is above the DSSV prediction (though DSSV is not shown on this plot, values that are near or above GRSV-STD are above DSSV).

Finally, we compare to the STAR single-inclusive jet A_{LL} measurement, shown in Figure 5.4 [4]. This result is plotted with DSSV's single-inclusive jet prediction in green: once




Figure 5-3: The 2005 and 2006 $\pi^0 A_{LL}$ results from the PHENIX experiments, compared to several theory curves.

again, the data are above the DSSV prediction. The persistence of this trend across numerous observables and multiple detectors convinced DSSV to recalculate their polarized parton distributions. The result is shown in magenta on Figure 5.4: the new DSSV is above the old DSSV at lower p_T values before coming back to join it at the higher values. The previous DSSV polarized gluon distributions corresponded to a value of $\Delta G = -0.1$ (compared to $\Delta G = 1.8$ for GRSV-MAX, $\Delta G = 0.4$ for GRSV-STD, and $\Delta G = 0.1$ for GRSV-ZERO). The new DSSV polarized gluon distribution corresponds to a value of $\Delta G = 0.1$ integrated from $0.05 < x < 0.2$ (over the full x range, the value is $\Delta G \sim 0.12$) [4]. This new value of ΔG is almost as large as the quark spin contribution $\frac{1}{2}\Delta\Sigma \sim 0.15$. Rather than ΔG being negligible, as many people thought, recent A_{LL} results, including this one, suggest that the gluon spin contribution to the proton spin is actually of a similar magnitude to the quark spin contribution.

The next step from here is to move on to correlated results, such as measuring A_{LL} for dijets rather than for single-inclusive jets. Figure 5.5 shows the preliminary 2009 STAR dijet A_{LL} measurement, compared to several theory curves [3]. Dijet and other correlation measurements allow us to start investigating the shape of the $\Delta g(x, Q^2)$ distribution, instead of just its first moment. Similarly to the π^0 and inclusive jet results, the dijet result




Figure 5-4: The 2006 and 2009 STAR single-inclusive jet A_{LL} results, compared with a number of theory curves, including old DSSV (in green) and new DSSV (in magenta).




Figure 5-5: The 2009 STAR preliminary dijet result.

shows an asymmetry that is larger than predicted by DSSV at lower invariant mass values. If the above result of the non-zero value of ΔG holds up under further investigations such as these, it will establish that the polarization of the quarks and gluons account for only about half of the proton spin. To account for the entire proton spin, a non-zero net orbital motion of the quarks and gluons is needed.

THIS PAGE INTENTIONALLY LEFT BLANK

Bibliography

- [1] M. Bai, L. Ahrens, J. Alessi, K. Brown, G. Bunce, P. Cameron, C. M. Chu, J. W. Glenn, H. Huang, A. E. Kponou, K. Krueger, W. Lamble, A. Luccio, Y. I. Makdisi, S. Y. Lee, M. Okamura, L. Ratner, K. Reece, T. Roser, H. Spinka, M. J. Syphers, N. Tsoupas, D. G. Underwood, W. van Asselt, N. Williams, and A. Yokosawa. Overcoming intrinsic spin resonances with an rf dipole. *Phys. Rev. Lett.*, 80:4673–4676, May 1998.
- [2] Gerry Bunce, Naohito Saito, Jacques Soffer, and Werner Vogelsang. Prospects for spin physics at rhic. *Annual Review of Nuclear and Particle Science*, 50 (1):525–575, 2000.
- [3] Matthew Walker for The STAR Collaboration. Gluon Polarization Measurements from Correlated Probes. 2011.
- [4] D. de Florian, R. Sassot, M. Stratmann, and W. Vogelsang. QCD Spin Physics: Partonic Spin Structure of the Nucleon. *Prog.Part.Nucl.Phys.*, 67:251–259, 2012.
- [5] Daniel de Florian, G. A. Navarro, and R. Sassot. Extraction of polarized parton densities from polarized dis and sidis. *AIP Conf. Proc.*, 792:921–924, 2005.
- [6] John R. Ellis and R. L. Jaffe. A sum rule for deep inelastic electroproduction from polarized protons. *Phys. Rev. D*, 9:1444, 1974.

- [7] A. Adare et al. Gluon-spin contribution to the proton spin from the double-helicity asymmetry in inclusive π^0 production in polarized p+p collisions at $\sqrt{s}=200$ gev. *Phys. Rev. Lett.*, 103(012003), 2009.
- [8] B. I. Abelev et al. Longitudinal double-spin asymmetry and cross section for inclusive neutral pion production at midrapidity in polarized proton collisions at $\sqrt{s}=200$ gev. *Phys. Rev. D*, 80(111108), 2009.
- [9] C. Adler et al. The rhic zero degree calorimeters. *Nucl. Inst. Meth. A*, 470:488–499, 2001.
- [10] Elliot D. Bloom et al. High-energy inelastic e p scattering at 6-degrees and 10-degrees. *Phys. Rev. Lett.*, 23:930–934, 1969.
- [11] H. Hahn et al. The rhic design overview. *Nucl. Inst. Meth. A*, 499:245–263, 2003.
- [12] I. Alekseev et al. Polarized proton collider at rhic. *Nucl. Inst. Meth.*, January 2003.
- [13] J. Ashman et al. An investigation of the spin structure of the proton in deep inelastic scattering of polarized muons on polarized protons. *Nucl. Phys. B*, 328:1, 1989.
- [14] K. H. Ackermann et al. Star detector overview. *Nucl. Inst. Meth. A*, 499:624–632, 2003.
- [15] M. Anderson et al. The star time projection chamber: A unique tool for studying high multiplicity events at rhic. *Nucl. Inst. Meth. A*, 499:659–678, 2003.
- [16] M. Beddo et al. The star barrel electromagnetic calorimeter. *Nucl. Inst. Meth. A*, 499:725–739, 2003.
- [17] Martin Breidenbach et al. Observed behavior of highly inelastic electron-proton scattering. *Phys. Rev. Lett.*, 23:935–939, 1969.

- [18] K. Nakamura et al. (Particle Data Group). Review of particle physics. *J. Phys. G*, 37(075021), 2010 and 2011 partial update for the 2012 edition.
- [19] Oleksandr Grebenyuk. *Neutral Meson Production in $d+Au$ and $p+p$ collisions at $\sqrt{s_{NN}}=200$ GeV in STAR*. PhD thesis, University of Utrecht, 2007.
- [20] D. J. Gross and Frank Wilczek. Ultraviolet behavior of non-abelian gauge theories. *Phys. Rev. Lett.*, 30:1343–1346, 1973.
- [21] Alan Hoffman. *Longitudinal Double-Spin Asymmetry and Cross Section for Inclusive Neutral Pion Production in Polarized Proton Collisions at $\sqrt{s}=200$ GeV*. PhD thesis, M.I.T., 2009.
- [22] R. L. Jaffe. The theory of the nucleon spin. *Phil. Trans. Roy. Soc. Lond. A*, 349:391–404, 2001.
- [23] R. L. Jaffe and Aneesh Manohar. The g1 problem: Deep inelastic electron scattering and the spin of the proton. *Nucl. Phys. B*, 337:509–546, 1990.
- [24] Xiangdong Ji. Gauge invariant decomposition of nucleon spin. *Phys. Rev. Lett.*, 78(4):610–613, 1996.
- [25] J. Kiryluk. Local polarimetry for proton beams with the star beam-beam counters. *hep-ex/0501072*, 2005.
- [26] Adam Kocoloski. *Measurement of Longitudinal Double Spin Asymmetries of $\pi^+/-$ Production in $p+p$ Collisions at $\sqrt{s}=200$ GeV at RHIC*. PhD thesis, M.I.T., 2010.
- [27] S.Y. Lee. *Spin Dynamics and Snakes in Synchrotrons*. World Scientific, Singapore, 1997.
- [28] Hiromi Okada. *Measurement of the Analyzing Power in pp Elastic Scattering in the CNI Region with a Polarized Atomic Hydrogen Gas Jet Target*. PhD thesis, Kyoto University, 2006.

- [29] H. David Politzer. Reliable perturbative results for strong interactions? *Phys. Rev. Lett.*, 30:1346–1349, 1973.
- [30] J. Sowinski. Star internal documentation. <http://www.star.bnl.gov/protected/spin/sowinski/analysis/derivations/multiParticle.pdf>.
- [31] Matthew Walker and Michael Betancourt. 2009 bemc tower calibration report, July 2010.
- [32] Jeff Wood. *Polarimetry at the Brookhaven AGS Using Proton-Carbon Coulomb-Nuclear Interference*. PhD thesis, UCLA, 2005.