Title

Probing Quark Dynamics in Hadronization of Quark—Gluon Plasma Droplets with Anisotropic Flow and Omega—hadron Correlations

ISBN

9798189211270

Author

Wu, Xiatong

Publication Date

2026-08-21

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
Los Angeles

Probing Quark Dynamics in Hadronization of Quark-Gluon Plasma Droplets with Anisotropic Flow and Omega-hadron Correlations

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

by

Xiatong Wu

2026

ABSTRACT OF THE DISSERTATION

Probing Quark Dynamics in Hadronization of Quark-Gluon Plasma Droplets with Anisotropic Flow and Omega-hadron Correlations

by

Xiatong Wu
Doctor of Philosophy in Physics
University of California, Los Angeles, 2026
Professor Huan Z. Huang, Chair

At lower RHIC energies, hadronization in the quark–gluon-plasma (QGP) fireball is significantly influenced by transported quarks carried in from the colliding nuclei. Within the quark coalescence picture, these transported quarks, which experience the full evolution of the system's pressure gradients, imprint distinct flavor and dynamical signatures on final-state hadrons. This thesis presents three complementary measurements from the STAR Beam Energy Scan Phase II (BES-II) program in Au+Au collisions.

The influence of transported quarks is directly reflected in the elliptic flow (v_2) of charged pions. In Au+Au collisions at $\sqrt{s_{NN}} = 7.7$–27 GeV, π^- carries systematically larger v_2 than π^+ across centralities and energies, consistent with the higher availability of transported d quarks relative to u quarks in the neutron-rich Au nucleus. Using a coalescence sum rule with the antiproton as a proxy for the produced-quark flow, the transported quark ratio $N_d^{\mathrm{tr}}/N_u^{\mathrm{tr}}$ is extracted and found to be consistent with the Au stoichiometric value $315/276 \approx 1.14$, providing a direct link between the quark content of the colliding nuclei, via nuclear geometry, and the final collective expansion of the QGP drop.

Electromagnetic fields generated by the spectator protons in peripheral collisions offer a complementary window into the quark-level response at hadronization. The directed flow (v_1) splitting between Λ^0 and $\bar{\Lambda}^0$ is measured over the same BES-II energy range; the slope difference $\Delta(dv_1/dy)$ becomes negative in peripheral collisions, mirroring the behavior previously observed for charged hadrons and attributed to the combined Faraday and Coulomb effects. A first-order coalescence test using p–$\bar{p}$, supplemented by a second-order correction based on K^+–K^-, supports the coalescence picture and indicates that the electromagnetic response acts at the level of the constituent quarks and is not set by electric charge alone.

These flow-based analyses probe the transported-quark coalescence picture through anisotropic momentum signatures. A more direct handle on the transported-baryon population is provided by the Ω^- hyperon: at $\sqrt{s_{NN}}$ = 7.7–27 GeV, the Ω^- yield carries a substantial excess over $\bar{\Omega}^+$ from baryon transport. We establish the Combinatorial Background Subtraction (CBS) correlation method, which uses associated kaons to decompose the Ω^- population into its transported and pair-produced components and to probe whether their hadronization signatures differ. At $\sqrt{s_{NN}}$ = 19.6 GeV, the transported Ω^- shows a modest excess of associated kaons over a pair-produced Ω, 1.10 ± 0.74 per Ω, in the direction expected if its strangeness is balanced by kaons rather than by co-produced anti-baryons.

The dissertation of Xiatong Wu is approved.

Michail Bachtis
Graciela B. Gelmini
Zhongbo Kang
Huan Z. Huang, Committee Chair

University of California, Los Angeles
2026

To my mother.

Contents

List of Tables

List of Figures

List of Acronyms

AGS	Alternating Gradient Synchrotron
AMPT	A Multi-Phase Transport model
AtR	AGS-to-RHIC transfer line
BES	Beam Energy Scan
BES-II	Beam Energy Scan Phase II
BtA	Booster-to-AGS transfer line
CBS	Combinatorial Background Subtraction
CM	Central Membrane
CMW	Chiral Magnetic Wave
DAQ	Data Acquisition
DCA	Distance of Closest Approach
EBIS	Electron Beam Ion Source
EM	Electromagnetic
EPD	Event Plane Detector
EVB	Event Builder
FEE	Front-End Electronics
HLT	High-Level Trigger
HPSS	High Performance Storage System
IBS	Intrabeam Scattering
IP	Interaction Point
L0	Level-0 trigger
LEReC	Low Energy RHIC electron Cooling

LHC	Large Hadron Collider
MRPC	Multi-gap Resistive Plate Chamber
MWPC	Multi-Wire Proportional Chamber
NBD	Negative Binomial Distribution
NCQ	Number of Constituent Quarks
PDG	Particle Data Group
PID	Particle Identification
PV	Primary Vertex
QA	Quality Assurance
QCD	Quantum Chromodynamics
QGP	Quark–Gluon Plasma
RF	Radio-Frequency
RFQ	Radiofrequency Quadrupole
RHIC	Relativistic Heavy Ion Collider
RefMult	Reference Multiplicity
RefMultCorr	Corrected Reference Multiplicity
SiPM	Silicon Photomultiplier
STAR	Solenoidal Tracker at RHIC
STP	Star Trigger Pusher
TCD	Trigger Clock Distribution
TCU	Trigger Control Unit
TOF	Time-of-Flight
TPC	Time Projection Chamber
VPD	Vertex Position Detector

Acknowledgments

I want to first thank my parents and my grandparents, without whom I would not have become a person curious about the world around me and willing to devote myself to the pursuit of scientific wisdom. I also thank my uncle and aunt, Juntao Gao and Shan Lu, for their unwavering support ever since my high school years.

I owe a great deal to Professors Huan Huang and Gang Wang, my two closest mentors through these years. They introduced me to the field of heavy-ion physics and have been the ones I could always turn to when an analysis stopped making sense. Beyond that, they helped shape me as a scientist with rigor, honesty, and scrutiny of my own results.

I have been blessed with a wonderful group of colleagues within the UCLA nuclear physics group and beyond. I want to thank Dylan Neff and Zhiwan Xu for the many thought-provoking exchanges on technical and physics questions, from which I learned a great deal. I also thank Brian Chan, Zhongling Ji, Thomas Marshall, Aditya Prasad Dash (who produced the proton and kaon directed-flow measurements), Yunshan Cheng, Xiaozhou Yu, and Siyuan Ping for many fruitful discussions on various topics. My thanks also go to Aihong Tang and Prithwish Tribedy for their support and guidance.

Finally, I want to thank my partner, Jiayu Lyu, who has been beside me for the whole of this journey. She endured its demands as fully as I did and made sacrifices of her own along

the way, and gave nothing but patience and love in return. I would not have reached this point without her.

Vita

2020 B.S. in Physics (Minor in Mathematics), *summa cum laude*, University of California, Los Angeles.

2020–2026 Graduate Student Researcher, Department of Physics and Astronomy, University of California, Los Angeles.

2020–2026 Teaching Assistant and Teaching Associate, Department of Physics and Astronomy, University of California, Los Angeles.

2026 Rudnick–Abelmann Endowed Fellowship, UCLA Department of Physics and Astronomy.

Selected Publications

- X. Wu *et al.*, "Baryon Number Transport, Strangeness Conservation and Ω–hadron Correlations," EPJ Web Conf. **276**, 03002 (2023).
- Z. Xu, X. Wu *et al.*, "Flow-plane decorrelations in heavy-ion collisions with multiple-plane cumulants," Phys. Rev. C **105**, 024902 (2022).
- W. Dong, X. Yu, S. Ping, X. Wu *et al.*, "Study of baryon number transport dynamics

and strangeness conservation effects using Ω–hadron correlations," Nucl. Sci. Tech. **35**, 120 (2024).

- Co-author on more than twenty STAR Collaboration publications (2022–2024), including works in *Nature*, Phys. Rev. X, and Phys. Rev. Lett.

Selected Presentations

- Poster, Quark Matter 2025, Frankfurt, Germany, April 2025.
- Poster, Quark Matter 2023, Houston, USA, September 2023.
- Talk, APS Division of Nuclear Physics Fall Meeting, New Orleans, USA, October 2022.
- Talk, Strangeness in Quark Matter 2022, Busan, Republic of Korea, June 2022.

Chapter 1

Introduction

Quantum Chromodynamics (QCD) is the gauge theory that underlies the strong nuclear force, governing the interactions of quarks and gluons, the building blocks of visible matter. At typical hadronic scales, quarks and gluons are confined into protons, neutrons, and other hadrons, but under extreme conditions of temperature or density, QCD predicts a transition to a deconfined phase known as the Quark–Gluon Plasma (QGP). The QGP is believed to have filled the universe within the first few microseconds after the Big Bang and may also exist in the cores of neutron stars at high baryon density. Creating and studying this phase in the laboratory provides a unique opportunity to test QCD in the strongly coupled regime and to understand the transition from quarks and gluons to the ordinary hadrons that form the visible matter of the universe today.

1.1 Fundamentals of Quantum Chromodynamics

As part of the Standard Model, QCD is the non-Abelian gauge theory describing the strong interaction, based on the $SU(3)_{\text{color}}$ gauge group. Quarks, spin-$\frac{1}{2}$ fermions carrying a color charge, transform in the fundamental representation $\mathbf{3}$ of SU(3). Gluons, the gauge bosons mediating the strong force, transform in the adjoint representation $\mathbf{8}$ and themselves carry color charge, allowing for gluon self-interactions.

The QCD Lagrangian is given by

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} G^a_{\mu\nu} G^{a\,\mu\nu} + \sum_{f=1}^{N_f} \bar{\psi}_f (i\gamma^\mu D_\mu - m_f)\psi_f, \tag{1.1}$$

where $G^a_{\mu\nu}$ is the gluon field-strength tensor,

$$G^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu + g f^{abc} A^b_\mu A^c_\nu, \tag{1.2}$$

and $D_\mu = \partial_\mu - igT^a A^a_\mu$ is the covariant derivative acting on quark fields ψ_f of flavor f. The T^a are the SU(3) generators in the fundamental representation, f^{abc} are the structure constants, and g is the strong coupling. The non-Abelian term $gf^{abc}A^b_\mu A^c_\nu$ in $G^a_{\mu\nu}$ generates cubic and quartic gluon self-interactions absent in Abelian gauge theories such as quantum electrodynamics. These self-interactions underlie the two defining features of QCD: asymptotic freedom and confinement. Observable hadrons are color-singlet composites of these degrees of freedom: quark–antiquark pairs (mesons) or three-quark states (baryons).

The strong coupling $\alpha_s \equiv g^2/4\pi$ depends sensitively on the energy scale Q^2. At large Q^2, α_s becomes small and quarks and gluons behave as quasi-free partons; this property, asymptotic freedom, makes perturbative QCD predictive at high energies. At small Q^2, α_s grows large and the theory is non-perturbative: quarks and gluons are confined inside hadrons, and the energy required to separate two color charges grows with their separation until it exceeds the threshold for vacuum pair production and hadronization.

This scale separation has a profound consequence for nuclear matter. Under ordinary conditions, confinement forbids the existence of free quarks and gluons as asymptotic states. At sufficiently high temperature or baryon density, however, the relevant momentum scales become large compared to Λ_{QCD}, the coupling weakens, and the medium can no longer support hadronic bound states. QCD then predicts a transition to a deconfined phase of quasi-free quarks and gluons, the Quark–Gluon Plasma. The conditions under which this transition occurs, and the properties of the matter it produces, provide the physics context for this thesis, which examines the collective dynamics and hadronization of the deconfined matter created in heavy-ion collisions.

1.2 The QCD Phase Diagram

The thermodynamic state of strongly interacting matter is governed by two intensive variables: the temperature T and the baryon chemical potential μ_B, which quantifies the net excess of quarks over antiquarks. The QCD phase diagram, plotted in the T–μ_B plane, encodes the boundaries between the confined hadronic phase and the deconfined QGP. Mapping this diagram experimentally is one of the principal goals of heavy-ion physics research.

At vanishing baryon chemical potential ($\mu_B \approx 0$), the phase structure can be computed from first principles using lattice QCD. These calculations establish that the transition from hadronic matter to the QGP is not a sharp phase transition but a smooth analytic crossover, centered on a pseudocritical temperature $T_{pc} \approx 155\,\mathrm{MeV}$, across which both deconfinement and the restoration of chiral symmetry set in [18]. The deconfined matter is not a weakly interacting gas of quarks and gluons: measurements of collective flow at the nearly baryon-free top energies of the Relativistic Heavy Ion Collider (RHIC) and the Large Hadron Collider (LHC) show that it behaves as a strongly coupled, nearly perfect liquid, with a shear-viscosity-to-entropy ratio approaching the conjectured quantum lower bound $\eta/s \approx 1/(4\pi)$ [37].

At large μ_B, lattice methods are obstructed by the fermion sign problem, and the structure of the phase diagram becomes correspondingly less certain. Effective models and functional approaches generally predict that the crossover at $\mu_B = 0$ steepens into a first-order phase transition at sufficiently large μ_B, with a discontinuity in the order parameter and a finite latent heat. If both predictions are correct, the crossover and first-order regions must connect at a second-order critical endpoint, the QCD critical point, whose existence and location remain among the most important open questions in the field.

The baryon chemical potential reached at midrapidity in a heavy-ion collision is controlled primarily by the collision energy. At top RHIC and LHC energies the matter produced at midrapidity is almost baryon-free, with $\mu_B \sim 20\,\mathrm{MeV}$; as $\sqrt{s_{NN}}$ is reduced, a larger fraction of the incoming baryon number is stopped at midrapidity and μ_B rises accordingly. The Beam Energy Scan Phase II (BES-II) program at RHIC, with Au+Au collisions in the range $\sqrt{s_{NN}} = 7.7$–$27\,\mathrm{GeV}$, populates chemical freeze-out conditions spanning $\mu_B \approx 150$–$420\,\mathrm{MeV}$. This reaches close to the region where any first-order transition or critical point is expected to lie, and where the transported-quark dynamics that motivate this thesis become quantitatively dominant.

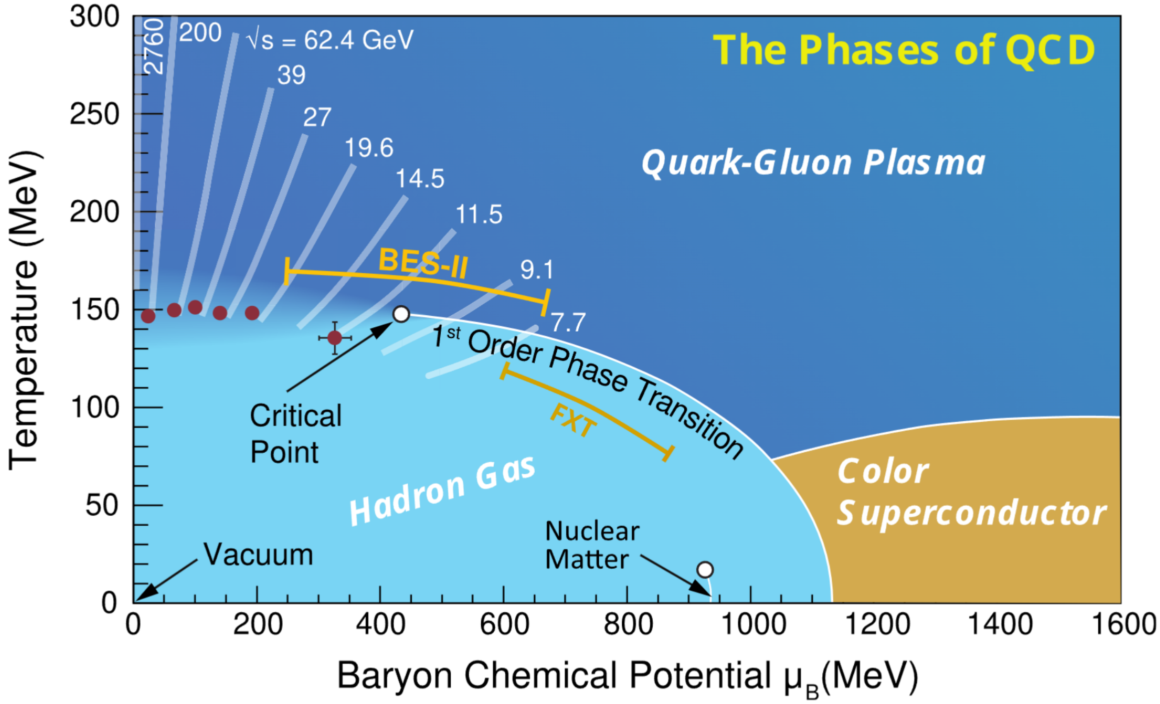

Figure 1.1: The QCD phase diagram in the T–μ_B plane, as conjectured from current understanding. Red circles mark chemical freeze-out points determined experimentally. The white lines indicate schematically the trajectories followed by the hydrodynamically expanding fireball, labeled by the initial collision energy $\sqrt{s_{NN}}$ in GeV; lowering the beam energy shifts the trajectory to higher μ_B, so that a beam-energy scan translates directly into a scan of the phase diagram. The coverage of the second phase of the RHIC Beam Energy Scan is marked "BES-II" (collider mode) and "FXT" (fixed-target mode). At small μ_B the transition between the hadron gas and the Quark–Gluon Plasma is a smooth crossover; at large μ_B it is conjectured to become first order, terminating in a critical point. Figure from [24].

1.3 Relativistic Heavy-Ion Collisions

The QGP is produced in the laboratory by colliding heavy nuclei at ultrarelativistic energies. At RHIC, gold nuclei (^{197}Au) are accelerated to center-of-mass energies of $\sqrt{s_{NN}} = 7.7$–200 GeV per nucleon pair, while the LHC extends the reach to 5.02 TeV with Pb+Pb collisions. The analyses presented in this thesis use Au+Au data from the BES-II program at RHIC, with $\sqrt{s_{NN}} = 7.7$–27 GeV. The present section introduces the geometric and dynamical concepts required to interpret the measurements.

1.3.1 Collision Geometry and Centrality

The geometry of a heavy-ion collision is characterized by the impact parameter b, defined as the transverse distance between the centers of the two colliding nuclei. In a head-on (*central*) collision, $b \approx 0$ and the full nuclear volume participates; in a grazing (*peripheral*) collision, b approaches the sum of the nuclear radii and only a small overlap region is formed. The plane spanned by the impact parameter vector and the beam direction defines the *reaction plane*, which provides the natural reference for collective observables (Sect. 1.4).

The impact parameter is not directly measurable but is strongly correlated with the final-state charged-particle multiplicity, which grows with the size of the overlap region. Experimentally, events are classified into *centrality* classes defined as percentiles of the total inelastic cross section, ordered from most central (0–5%) to most peripheral (70–80%). The mapping from measured multiplicity to centrality is constructed using a Monte Carlo Glauber model [42], which simulates collisions of nucleons sampled from a Woods–Saxon density profile and provides the average number of participating nucleons N_{part} and binary nucleon–nucleon collisions N_{coll} in each class.

1.3.2 Space-Time Evolution

The evolution of a heavy-ion collision proceeds through several distinct stages:

1. **Initial state** ($\tau \lesssim 1\,\mathrm{fm}/c$). The two Lorentz-contracted nuclei pass through one another, depositing energy and baryon number into the overlap region. At top RHIC and LHC energies the nucleons are almost transparent and the midrapidity region is nearly baryon-free; at BES-II energies, baryon stopping is substantial and a significant fraction of the incoming baryon number is transported to midrapidity as *transported quarks*. The initial energy density distribution reflects the positions of the participating nucleons together with event-by-event fluctuations.

2. **Thermalization and QGP formation** ($\tau \sim 1\,\mathrm{fm}/c$). The deposited energy density is well above the crossover value [18], and the system thermalizes into a deconfined QGP. In non-central collisions the overlap region has an almond-shaped spatial anisotropy, conventionally characterized by the eccentricity ε_n. Pressure gradients along the minor axis of the almond drive the subsequent expansion preferentially in that direction, converting the initial spatial anisotropy into momentum-space anisotropy.

3. **Hydrodynamic expansion**. The QGP evolves as a relativistic fluid, cooling as it expands. Pressure gradients translate the initial geometric anisotropies into azimuthal modulations of the particle momentum distribution that survive to the final state.

4. **Hadronization** ($T \approx T_{pc}$). As the local temperature drops below the pseudocritical value, quarks and gluons recombine into color-singlet hadrons. The two relevant mechanisms, string fragmentation and quark coalescence, are discussed in Sect. 1.5; coalescence dominates at the intermediate transverse momenta accessed by the analyses in this thesis.

5. **Hadronic rescattering and freeze-out**. After hadronization the system enters a hadronic phase in which inelastic and elastic collisions continue. *Chemical freeze-out* occurs when inelastic processes cease, fixing the particle yields; *kinetic freeze-out* occurs when elastic scattering also stops, fixing the momentum spectra. The detected particles are emitted from the kinetic freeze-out surface.

The evolution from a thermalized, nearly perfect-fluid QGP to the final hadronic state imprints the early-time geometry on the momenta of emitted particles. Quantifying this imprint through the Fourier harmonics of the azimuthal distribution, and interpreting it in terms of the underlying quark dynamics, is the subject of the next two sections.

1.4 Collective Flow and Azimuthal Anisotropy

Collective flow refers to the correlations between the momenta of emitted particles and the geometry of the collision. These correlations arise because the initial spatial anisotropy of the overlap region is converted, by pressure gradients in the expanding fireball, into a momentum-space anisotropy that survives to the final state [45, 51]. Because the conversion is most efficient during the earliest, densest phase of the evolution, flow observables are among the most sensitive probes of the equation of state and transport properties of the QGP.

1.4.1 Fourier Decomposition of the Azimuthal Distribution

The azimuthal distribution of particles emitted in a heavy-ion collision can be expanded in a Fourier series with respect to the reaction-plane angle Ψ_{RP}:

$$E\frac{d^3N}{d^3p} = \frac{1}{2\pi}\frac{d^2N}{p_T\,dp_T\,dy}\left[1 + 2\sum_{n=1}^{\infty} v_n \cos\big(n(\phi - \Psi_{\mathrm{RP}})\big)\right], \tag{1.3}$$

where ϕ is the azimuthal angle of the particle, p_T its transverse momentum, and y its rapidity. The Fourier coefficients v_n quantify the strength of each harmonic and are obtained as

$$v_n = \langle \cos\big(n(\phi - \Psi_{\mathrm{RP}})\big)\rangle, \tag{1.4}$$

where the average runs over all particles in all events within a given kinematic class. The two harmonics relevant to this thesis are the directed flow (v_1) and the elliptic flow (v_2); higher harmonics (v_3, v_4, ...) carry complementary information about initial-state fluctuations but are not used here [47, 51].

1.4.2 Elliptic Flow (v_2)

In non-central collisions, the overlap region of the two nuclei has an almond shape with the minor axis aligned along the impact-parameter direction. The resulting pressure gradient is stronger along the minor axis, preferentially accelerating particles in that direction and producing a $\cos(2\phi)$ modulation of the azimuthal distribution. The coefficient of this modulation, v_2, is the most extensively studied flow observable.

Elliptic flow is self-quenching: as the system expands, the spatial eccentricity is consumed and the driving pressure gradient diminishes. The bulk of v_2 is therefore generated during the earliest, densest phase of the evolution, the QGP phase, making it a sensitive probe of the QGP shear viscosity. Viscous hydrodynamic calculations describe the measured v_2 at RHIC and LHC with $\eta/s \approx 1/(4\pi)$, providing the empirical foundation for the perfect-fluid characterization of the QGP discussed in Sect. 1.2.

A further structural feature of the v_2 data established the partonic origin of collective flow. When the elliptic flow of identified hadrons is scaled as v_2/n_q versus p_T/n_q, with $n_q = 2$ for mesons and $n_q = 3$ for baryons, the data for all species collapse onto a common curve [2, 26, 36]. This number-of-constituent-quark (NCQ) scaling is the canonical signature of *quark coalescence*: it implies that flow develops at the level of the constituent quarks and is inherited additively by the formed hadrons. The implications of this picture for the analyses in this thesis are developed in Sect. 1.5.

1.4.3 Directed Flow (v_1)

The first Fourier coefficient, v_1, describes the collective sideward deflection of particles in the reaction plane. The rapidity-odd component, which carries the physics of the collision dynamics, is approximately linear near midrapidity:

$$v_1(y) \approx \left.\frac{dv_1}{dy}\right|_{y=0} \cdot y, \tag{1.5}$$

and the slope $dv_1/dy|_{y=0}$ characterizes the tilt of the emitting source. Directed flow is generated at the earliest moments of the collision, before the eccentricity-driven elliptic flow develops, and is sensitive to the initial longitudinal tilt of the fireball, the passage of the spectator nucleons, and, at BES energies, the equation of state near the phase transition.

In addition to the hydrodynamic response, v_1 at BES energies carries a distinct sensitivity to the transient electromagnetic fields generated by the spectator protons. Because these fields act differentially on positively and negatively charged quarks before hadronization, they produce a charge-dependent splitting of the directed flow between particles and antiparticles, $\Delta(dv_1/dy)$. The underlying mechanisms (the Coulomb deflection, the Faraday induction from the decaying magnetic field, and the Lorentz/Hall force) and their detailed phenomenology are treated in Sect. 3.4, where the Λ–$\bar{\Lambda}$ splitting is used as a probe of the partonic response to these fields.

1.4.4 The Event Plane Method

The reaction plane Ψ_{RP} is not directly observable. In practice, it is estimated event by event from the azimuthal distribution of the detected particles themselves, yielding the *event plane* Ψ_n associated with the n-th harmonic. The event plane is constructed from the flow vector

$$\vec{Q}_n = \Big(\textstyle\sum_i w_i \cos(n\phi_i),\ \sum_i w_i \sin(n\phi_i)\Big), \tag{1.6}$$

where the sum runs over the tracks or detector hits used and the weights w_i are chosen to optimize the resolution. The event-plane angle is

$$\Psi_n = \frac{1}{n}\arctan\left(\frac{Q_{n,y}}{Q_{n,x}}\right). \tag{1.7}$$

Because Ψ_n is estimated from a finite number of particles, it fluctuates around the true reaction plane and dilutes the measured v_n. The dilution is corrected by the event-plane resolution,

$$v_n = \frac{\langle\cos\big(n(\phi - \Psi_n)\big)\rangle}{\mathrm{Res}(\Psi_n)}, \qquad \mathrm{Res}(\Psi_n) = \langle\cos\big(n(\Psi_n - \Psi_{\mathrm{RP}})\big)\rangle, \tag{1.8}$$

which is determined from the data using the sub-event method: the detector is split into two independent regions (typically east and west in pseudorapidity), and the correlation between their event planes is measured. For sub-events of equal resolution,

$$\mathrm{Res}(\Psi_n) = \sqrt{\left\langle \cos\left(n\Psi_n^A - n\Psi_n^B\right)\right\rangle}; \tag{1.9}$$

when the resolutions differ, additional correction factors are required. The specific event-plane detectors and weighting choices used in this thesis are described in Sect. 3.3.4 (pion v_2) and Sect. 3.4.3 (Λ v_1) [47].

1.5 Hadronization and Quark Coalescence

As the QGP cools below the pseudocritical temperature T_{pc}, quarks and gluons must re-assemble into color-singlet hadrons. Two mechanisms operate in parallel and dominate in different kinematic regimes.

1.5.1 String Fragmentation

At high transverse momentum ($p_T \gtrsim 6\,\mathrm{GeV}/c$), hadron production is dominated by the fragmentation of hard-scattered partons. In the Lund string model, a color-connected $q\bar{q}$ pair is joined by a chromoelectric flux tube of approximately uniform tension $\kappa \approx 1\,\mathrm{GeV/fm}$. As the quarks separate, the string stretches and breaks by Schwinger pair production, $P \propto \exp(-\pi m_\perp^2/\kappa)$, generating new $q\bar{q}$ pairs and ultimately yielding a collimated jet of hadrons. Baryons are produced when the string breaks into a diquark–antidiquark pair instead of a quark–antiquark pair, and the suppression of strange relative to light quarks ($m_s \approx 0.45\,\mathrm{GeV}$ vs. $m_{u,d} \approx 0.33\,\mathrm{GeV}$) carried by the Schwinger exponential explains the observed hierarchy of yields among hadron species in jet fragmentation [14]. The string-fragmentation picture, with its specific implications for multi-strange baryon production and baryon-number transport at finite baryon density, is revisited in Sect. 1.6.

1.5.2 Quark Coalescence

At intermediate and low transverse momentum, a different mechanism takes over. In the coalescence (or recombination) picture, hadrons are formed when co-moving constituent quarks find themselves close in phase space at the moment of hadronization. Because the quark distribution in the QGP is approximately thermal, low-momentum quarks are abundant and the probability of finding two (three) quarks with appropriate quantum numbers to form a meson (baryon) is large; this is the bulk hadronization mechanism relevant to the analyses in this thesis [26, 27, 31].

The defining experimental signature of coalescence is the NCQ scaling introduced in Sect. 1.4.2: hadronic flow is, to a good approximation, the sum of the flows of the constituent quarks evaluated at fixed transverse momentum per constituent,

$$v_n^{\mathrm{hadron}}(p_T) \;\approx\; \sum_{i=1}^{n_q} v_n^{q_i}\left(\frac{p_T}{n_q}\right). \tag{1.10}$$

Equation 1.10 is the *coalescence sum rule*, the impulse approximation to the general coalescence dynamics, in which hadron formation is sudden and leaves the constituent momenta unchanged. It encodes the assertion that collective flow develops at the partonic level, before confinement assembles the final-state hadrons, and that hadron formation preserves the information about the underlying quark dynamics.

1.5.3 Transported vs. Produced Quarks

The coalescence sum rule (Eq. 1.10) treats all constituent quarks of a given flavor as equivalently equilibrated, but at BES energies this equilibration is incomplete. Two populations of quarks coexist at midrapidity, and they behave differently:

- **Transported quarks** originate from the valence content of the incoming nucleons and are carried into the fireball at the moment of impact. They are present from $\tau = 0$ and experience the full development of the pressure gradients, acquiring strong collective flow.

- **Produced quarks** are pair-created from the color field as the QGP evolves and therefore acquire systematically weaker flow than the transported population.

Antibaryons such as $\bar{p}$ and $\bar{\Lambda}$, composed exclusively of pair-produced antiquarks, serve as experimental proxies for the flow of the produced-quark population, while their particle partners receive contributions from both populations. This asymmetry imprints a particle–antiparticle splitting on the flow observables that is a generic feature of BES-II measurements and provides the basic handle exploited by the analyses in this thesis.

1.5.4 Isospin Asymmetry of the Gold Nucleus

A further consequence of the transported–produced distinction follows from the isospin composition of the colliding nuclei. The ^{197}Au nucleus contains $Z = 79$ protons and $N = 118$ neutrons, with a neutron-to-proton ratio of $N/Z \approx 1.49$. Counting valence quarks under the constituent-quark assignment (two u and one d per proton, one u and two d per neutron):

$$N_u = 2Z + N = 276, \qquad N_d = Z + 2N = 315, \tag{1.11}$$

giving a transported d/u ratio of

$$\frac{N_d^{\mathrm{tr}}}{N_u^{\mathrm{tr}}} = \frac{315}{276} \approx 1.14. \tag{1.12}$$

The neutron excess implies that d transported quarks are more abundant than u transported quarks at midrapidity. Under coalescence, this differential availability is inherited by the formed hadrons, producing a flow splitting between species that carry different valence content, most simply between π^- ($\bar{u}d$, sensitive to the transported d) and π^+ ($u\bar{d}$, sensitive to the transported u). The recovery of the value in Eq. 1.12 from the pion v_2 data, with the antiproton serving as the produced-quark reference, is the central test of the coalescence picture pursued in Sect. 3.3.

Two refinements modify this simple counting. The neutron distribution in ^{197}Au has a slightly larger root-mean-square radius than the proton distribution (the *neutron skin*, $\Delta R_{np} \approx 0.15$–$0.25\,\mathrm{fm}$ [50]), so that the surface of the nucleus is more neutron-rich than the bulk. Peripheral collisions, in which only the outer layers overlap, are therefore expected to sample a d/u ratio above the bulk value of Eq. 1.12. The ^{197}Au nucleus is also modestly deformed ($\beta_2 \approx -0.13$ [43]), introducing additional event-by-event fluctuations in the geometry of the overlap region. Both effects are expected to introduce centrality-dependent modifications to the bulk prediction, and are discussed where relevant in Chapter 3.

1.6 Baryon Stopping, Gluon Junctions, and the Ω Hyperon

The transported quarks introduced in Sect. 1.5.3 are the partonic manifestation of *baryon stopping*: the deposition of a fraction of the incoming baryon number at midrapidity. The mechanism by which baryon number is transported over the large rapidity interval separating the beam fragments from the central region is a long-standing question in QCD, and one to which the multi-strange Ω hyperon offers an unusually clean experimental handle.

1.6.1 Baryon Stopping at BES Energies

At top RHIC and LHC energies the colliding nucleons are nearly transparent to one another: the midrapidity region is populated almost entirely by newly produced particles, with a net-baryon density (baryons minus antibaryons) close to zero. As the collision energy is lowered, the nucleons are increasingly decelerated in the collision, and a growing share of their baryon number is carried into the central region. By $\sqrt{s_{NN}} = 7.7\,\mathrm{GeV}$ the net-proton $(p - \bar{p})$ rapidity density at midrapidity has become significant, in contrast to the near-zero net-baryon density at top RHIC energy, signaling substantial stopping [6]. This trend is the same one that, expressed thermodynamically, raises the baryon chemical potential along the chemical freeze-out curve (Sect. 1.2). The valence quarks of the stopped nucleons are precisely the transported quarks whose enhanced collective flow underlies the pion v_2 and Λ v_1 analyses of Chapter 3.

1.6.2 Mechanisms of Baryon-Number Transport

How baryon number is conveyed from the beam to midrapidity is not uniquely prescribed by QCD. Two pictures are commonly contrasted.

- **Diquark transport.** A nucleon's valence content can be organized as a quark and a compact (ud) *diquark*, a color-antitriplet object that behaves, to first approximation, like a single slow-moving constituent. In the standard Lund implementation the diquark sits at a string endpoint and forms the leading baryon by combining with one quark from the nearest string break. Because the diquark retains its original light-quark content, the leading baryon naturally carries u and d valence quarks: it emerges as a proton, Λ, or Σ, not as a multi-strange state [15].

- **Gluon-junction transport.** Alternatively, the baryon number of a nucleon can be associated not with any specific quark but with a Y-shaped topological configuration of the gluon field, the *junction* at which three color flux tubes meet [35]. In this picture the junction is flavor-blind: it dresses with whichever quarks are present in its vicinity at hadronization, and Regge phenomenology allows it to be transported over several units of rapidity. A junction stopped at midrapidity in a strangeness-rich QGP can therefore hadronize with three strange quarks essentially as readily as with light quarks.

The two pictures are not mutually exclusive, and intermediate scenarios (a diquark that gradually exchanges flavor through successive string breaks, $ud \to us \to ss$, or a diquark that dissolves entirely) blur the distinction. The qualitative contrast that survives, and that motivates the measurement, is that the diquark picture has a strong intrinsic bias toward light-quark leading baryons, whereas the junction picture has none.

1.6.3 The Ω Hyperon as a Probe

Strange-quark production is enhanced in a deconfined medium relative to hadronic matter, because the thermal QGP supplies $s\bar{s}$ pairs copiously while strangeness production in a hadron gas is suppressed by large reaction thresholds [48]. Multi-strange baryons sit at the extreme of this hierarchy: the Ω^- (sss), with no valence light quarks at all, is among the cleanest hadronic probes of the strange sector of the QGP.

This same quark content makes the Ω^- a sharp discriminator between the transport mechanisms above. A leading baryon carrying transported baryon number reaches the Ω^- trivially in the junction picture, where it simply picks up three medium strange quarks, but only through a contrived chain of flavor exchanges in the intact-diquark picture. An Ω^- population that carries net transported baryon number therefore weighs against the diquark picture and in favor of a flavor-blind, junction-like mechanism, without claiming that any individual Ω^- originates from a junction.

Experimentally, the signature of transported baryon number in the Ω^- sample is an *excess* over the pair-produced baseline. At $\sqrt{s_{NN}} = 7.7$–$27\,\mathrm{GeV}$ the measured Ω^- yield exceeds the $\bar{\Omega}^+$ yield; the difference is the population of Ω^- that cannot be accounted for by $B\bar{B}$ pair production from the vacuum and must instead carry baryon number transported from the colliding nuclei [3]. The two components differ in their associated-particle environment: a pair-produced Ω^- is balanced in strangeness by a co-produced anti-hyperon ($\bar{\Xi}$, $\bar{\Sigma}$), whereas an Ω^- carrying transported baryon number balances its strangeness with associated kaons (K^+). Correlations between Ω^- and associated kaons thus provide a statistical handle on the transported component of the Ω^- sample. The correlation-based decomposition method (Combinatorial Background Subtraction) developed for this purpose is presented in Chapter 5.

1.7 Thesis Outline

This thesis presents three measurements from the BES-II program of the Solenoidal Tracker at RHIC (STAR) in Au+Au collisions. The first two test the quark-coalescence picture in the presence of quarks of different origin and flavor; the third studies the interplay between strangeness conservation and baryon number transport in the multi-strange sector.

Chapter 2 describes the RHIC accelerator complex and the STAR detector subsystems used in these analyses, in particular the Time Projection Chamber, the Time-of-Flight system, and the Event Plane Detector, together with the event and track reconstruction procedures, including the KFParticle topological fit used for weakly decaying hyperons.

Chapter 3 presents two collective-flow analyses. The first measures the elliptic-flow splitting between π^- and π^+ and applies the coalescence sum rule, with the antiproton as the produced-quark reference, to extract the transported d/u ratio at midrapidity and compare it to the ^{197}Au stoichiometric value of Eq. 1.12. The second measures the directed-flow splitting between Λ^0 and $\bar{\Lambda}^0$ as a function of centrality and tests whether it follows the coalescence prediction constructed from proton and kaon measurements, probing the response of the constituent quarks to the early-time electromagnetic fields. Chapter 4 reports the results of both analyses and their physics interpretation.

Chapter 5 turns to the multi-strange sector. It establishes the Combinatorial Background Subtraction method, which uses Ω–kaon correlations to characterize the transported component of the Ω^- sample, thereby testing the baryon-number-transport picture in a regime where the diquark and gluon-junction mechanisms make qualitatively different predictions. Chapter 6 applies the method at $\sqrt{s_{NN}} = 19.6$ GeV, isolating the correlation of the excess kaons with the excess Ω^- from the measured same- and mixed-event correlations.

Chapter 7 summarizes what each of the three analyses establishes about hadronization at finite baryon density.

Chapter 2

RHIC and the STAR Detector

Launched in 2000, the RHIC complex at Brookhaven National Laboratory is a versatile accelerator facility designed to explore the properties of strongly interacting matter at high temperature and baryon density. By accelerating and colliding heavy ions at nearly the speed of light, RHIC creates extremely hot and dense systems whose evolution provides a snapshot of conditions similar to those in the early universe microseconds after the Big Bang. RHIC is capable of colliding a wide range of beam species, from polarized protons to heavy nuclei such as gold and uranium, at center-of-mass energies from 7.7 to 200 GeV. This operational flexibility enables systematic studies across collision energies, system sizes, and beam species, offering a controlled environment for mapping the QCD phase structure and investigating the formation and evolution of the quark–gluon plasma.

Two independent rings ("blue" and "yellow") allow counter-circulating beams, while superconducting magnets guide and focus them for collisions at six interaction points [46]. STAR, located at the 6 o'clock interaction point, is one of the four original RHIC experiments and is one of the two currently operating experiments alongside the newly constructed sPHENIX. As a general-purpose detector with large and uniform acceptance, STAR enables precision

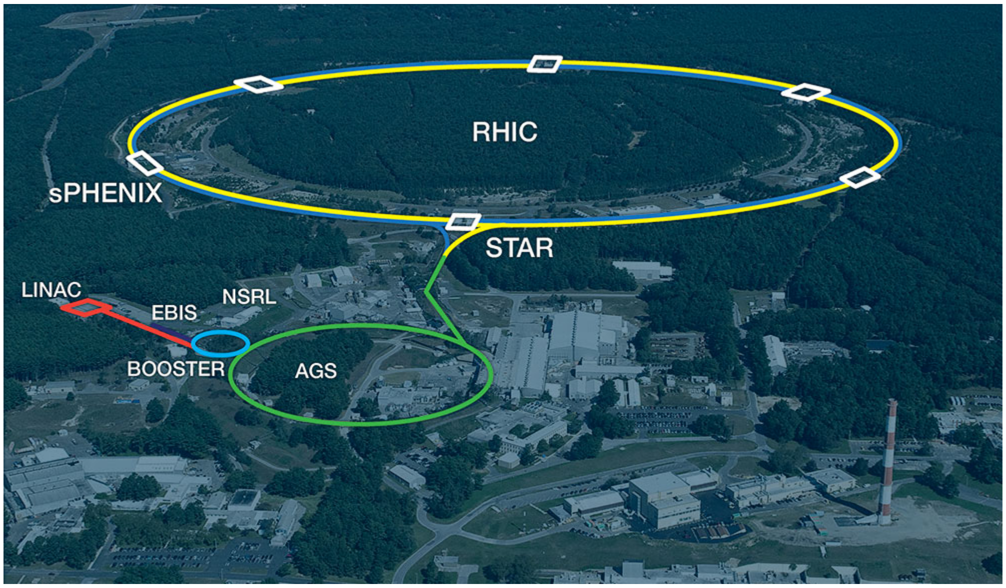

Figure 2.1: An overview of the RHIC complex, including the two active experiments, STAR and sPHENIX, located at the 6 o'clock and 8 o'clock positions, respectively. Figure from [21].

measurements of key QGP phenomena such as jet quenching, anisotropic flow and quark coalescence, as well as a broad array of discoveries in other QCD and cold nuclear physics.

2.1 RHIC Accelerator Complex

RHIC is an integrated accelerator complex consisting of several sequential stages for beam preparation and acceleration, beginning with ion production and culminating in final beam energization within the superconducting collider rings.

2.1.1 Electron Beam Ion Source

Ion species are first generated in dedicated sources, such as the Electron Beam Ion Source (EBIS) for heavy nuclei and a polarized H^- source for proton operations. The EBIS is a modern, high-performance injector designed to provide the high-charge-state heavy ions required for RHIC operations. Unlike traditional Tandem injectors, which rely on multiple stages of foil stripping and long transport lines, the EBIS produces fully stripped or highly charged ions in short, intense pulses suitable for single- or few-turn injection into the Booster synchrotron [11]. This allows greater control over beam quality, simplifies injection, and supports a wide range of ion species from helium to uranium.

In EBIS, low-charge ions are first injected from external ion sources into a long, superconducting solenoid trap, where they are confined by a high-density electron beam. The electron beam, emitted from a convex LaB_6 or IrCe cathode, can reach currents up to 10 A at energies around 20 keV. Compression of the electron beam by the solenoid creates the high-density environment necessary for sequential electron impact ionization, producing the desired high-charge-state ions. Auxiliary warm solenoids near the electron gun and collector allow precise control of beam transport and compression while keeping these regions separated from the central trap to maintain ultra-high vacuum ($\sim 10^{-9}$ Torr) [11].

Once ions reach the desired charge state, they are extracted using a carefully designed sequence of electrostatic and magnetic elements, including an ion extractor, accelerating tube, gridded lens, and solenoid lens, which together match the beam into the downstream accelerators. The electron collector, water-cooled and optimized for uniform power deposition, dissipates up to 230 kW from the electron beam while minimizing perturbations to the extracted ions [11].

After extraction from the EBIS trap, ions are transported to a radiofrequency quadrupole (RFQ) accelerator, which provides strong transverse focusing and initial bunching while increasing the beam energy to approximately 300 keV per nucleon. The RFQ also performs charge-state selection, efficiently rejecting unwanted ions and ensuring that only the desired charge state continues downstream. A short Linac then accelerates the beam further to about 2 MeV per nucleon, preparing it for synchrotron injection. For gold beams, EBIS routinely delivers pulses of Au^{32+} with widths of 10–40 μs at repetition rates up to 10 Hz, with intensities of around 1.2×10^9 ions per pulse. The short, intense pulses enable single- or few-turn injection into the Booster synchrotron within a few minutes for a physics store [28].

2.1.2 Synchrotrons

The prepared ion or proton beam then undergoes three stages of synchrotron acceleration at the RHIC complex, including the Booster synchrotron, the Alternating Gradient Synchrotron (AGS), and the main RHIC rings. These circular accelerators utilize two major components, magnets and RF cavities, to accelerate particles while maintaining longitudinal and transverse focus in circular motion around the ring.

Magnets in synchrotrons are primarily used to guide the beam around a circular trajectory and to focus the beam transversely. Realistic magnets are characterized by a multi-pole expansion, but they can be identified by the leading component in the series expansion. Dipole magnets, where the dipole term is the dominant term, bend the beam in a certain direction and are usually placed in the curved section along the ring. Quadrupole magnets ("quads"), where the quadrupole term dominates, focus the beam transversely in the straight sections. They are further divided into "focusing" quads that focus the beam horizontally but disperse it vertically, and "defocusing" ones that perform the opposite. A lattice setup consisting of a series of alternating focusing and defocusing quads is needed to focus the beam in both the vertical and horizontal transverse directions. Higher-order magnets and smaller trims are used to finely adjust the beam trajectory or profile, such as modifying the chromaticity and the tune of the betatron oscillation due to the lattice quad setup.

The acceleration of the beam in synchrotrons is achieved by the oscillating electric field generated in radio-frequency (RF) cavities. These cavities are designed to efficiently store

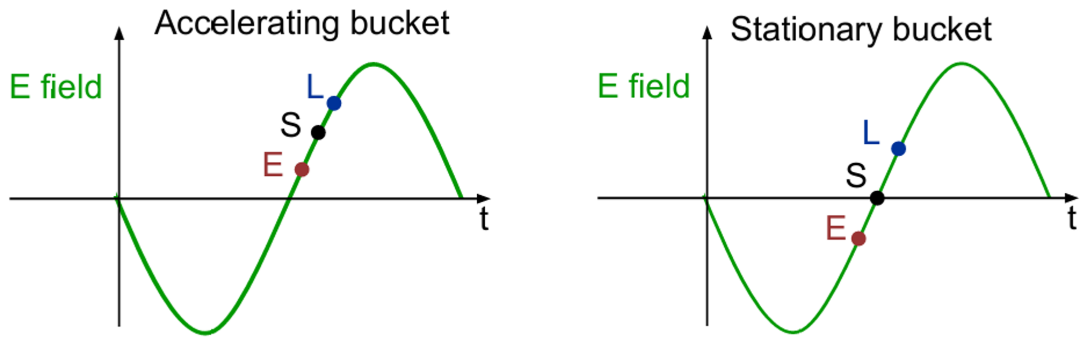

Figure 2.2: An accelerating bucket versus a stationary (storage) bucket. The three dots refer to particles that arrive at the cavity early ("E"), at the synchronous phase ("S"), and late ("L"). Adapted from [17].

electromagnetic energy at a resonant frequency. This frequency determines the maximum number of accelerating phases in the cavity during the time it takes for a particle to complete a full revolution through

$$h = t_{\mathrm{rev}} f_{\mathrm{rf}}, \tag{2.1}$$

where h is also known as the harmonic number. For the main RHIC rings with a circumference of around 3834 m, for example, a gold ion with $E/A \approx 9.8$ GeV is matched with a RF frequency of around 28.02 MHz, leading to a harmonic number of [29]

$$h = \frac{f_{\mathrm{rf}}\, C}{\beta c} = \frac{28.02 \times 10^6\,\mathrm{Hz} \times 3833.845\,\mathrm{m}}{0.9955 \times 2.998 \times 10^8\,\mathrm{m/s}} \approx 360. \tag{2.2}$$

In the accelerating stage, the phase of the RF EM wave is selected so that all beam particles experience an accelerating electric field when entering the cavities. However, due to the sinusoidal shape of the wave profile, particles that arrive earlier tend to experience a smaller accelerating field and vice versa, as shown in the left panel of Figure 2.2. As a result, particles with slightly different initial momenta will eventually be grouped into synchronized beam packets, a mechanism known as phase focusing. These synchronized bunches each occupy an accelerating bucket, where the maximum possible number of buckets within a ring is defined by the harmonic number h.

Longitudinal acceleration by the RF cavities must be coordinated with the ramping of the magnetic fields to maintain a stable circular orbit and minimize beam loss. At RHIC, superconducting magnets provide strong magnetic fields with minimal energy dissipation. Typically, arc dipole magnets at RHIC provide magnetic field strength of around 3.46 T [46]. Once the beam reaches the target energy, the RF phase can be adjusted to form stationary buckets. In this storage stage, particles arriving early or late relative to the synchronous phase experience restoring forces that confine them longitudinally, resulting in stable oscillations inside the stationary bucket, as illustrated in the right panel of Figure 2.2.

A critical parameter in synchrotron operation is the transition energy, denoted by γ_t. When the momentum of a particle increases within the synchrotron, two competing effects modify the revolution time: 1) the increase in velocity tends to reduce the revolution time, and 2) the increased magnetic rigidity forces the particle onto a longer orbital path, which increases the revolution time. Due to relativistic effects, it becomes increasingly energy-demanding for the RF cavities to increase the particles' velocities as they approach the speed of light. Therefore, the transition energy represents the critical threshold where the velocity effect and the path-length effect exactly balance. Below this threshold, the velocity effect dominates, and particles with higher momentum have a shorter revolution time. Above it, the path-length effect dominates, and higher-momentum particles take longer to complete a full cycle.

This phenomenon has significant consequences for the operation of RF cavities and the maintenance of phase stability. While an accelerating bucket like the one in the left panel of Figure 2.2 is stable below γ_t, as more energetic particles arrive at the cavities earlier and receive a smaller acceleration kick, the RF synchronous phase must be flipped to the opposite slope of the waveform to maintain longitudinal focus once γ_t is crossed. Failure to perform this phase jump would result in the rapid debunching and eventual loss of the beam.

The Booster Synchrotron

Following the pre-injector stage, gold ions are injected into the Booster, the first of the three synchrotrons in the RHIC injector complex. The EBIS provides high-intensity pulses with a width typically ranging from 10 to 40 μs [11]; given the Booster's revolution period, this corresponds to a highly efficient few-turn (1–4) injection. The ring's circumference of 201.78 m is specifically designed to be exactly one-quarter that of the AGS [8]. This geometric ratio facilitates a seamless bucket-to-bucket transfer, as the Booster accelerates ion bunches that are then mapped into one-quarter of the AGS's longitudinal phase space. To optimize beam intensity for the experimental program, the RF system performs an adiabatic bunch merging sequence: the beam is initially captured at harmonic $h = 4$, merged into $h = 2$, and finally into $h = 1$, resulting in a single high-intensity bunch prior to extraction [28].

The acceleration cycle increases the kinetic energy of the Au^{32+} ions to approximately 105 MeV/u. At this energy, the single bunch is extracted and directed through a thin stripping foil in the Booster-to-AGS (BtA) transfer line. The foil system consists of aluminum and carbon foils with optimized thickness to produce the highest yield of Au^{77+}. The downstream BtA dipoles and collimators act as a magnetic spectrometer to ensure only the Au^{77+} state is transported. This produces a beam with the requisite charge-to-mass ratio for efficient acceleration and stable passage through the transition energy in the AGS [28].

The Alternating Gradient Synchrotron

Initially commissioned in 1960 as the world's highest-energy particle accelerator, the AGS was designed to accelerate protons to 33 GeV for fixed-target nuclear and particle physics experiments. After RHIC commissioning in 2000, the AGS transitioned to serve primarily as the final stage of the RHIC injector chain, receiving Au^{77+} bunches from the Booster. During each AGS cycle, 8 Booster bunches are injected into $h = 16$ buckets and eventually merged to 2 bunches with $h = 4$ [28].

The primary challenge during the AGS acceleration cycle is the crossing of the transition energy at $\gamma_t \approx 8.5$. To mitigate emittance growth and beam loss, the AGS employs a specialized γ_t-jump system. This system utilizes high-speed quadrupole magnets to rapidly alter the machine's magnetic lattice, effectively shifting the transition energy of the ring itself. By "jumping" γ_t across the beam's current energy faster than the normal acceleration rate, the time the ions spend in the unstable longitudinal phase is minimized [16].

Following acceleration to an extraction energy of approximately 8.9 GeV/u (for $\sqrt{s_{NN}} =$ 200 GeV operation), the ions are extracted and transported via the AGS-to-RHIC (AtR) transfer line. Within this line, the ions pass through a final stripping foil. At these relativistic speeds, the remaining two K-shell electrons are removed with high efficiency, resulting in fully stripped $^{197}Au^{79+}$ nuclei. A final magnetic spectrometer in the AtR line selects this bare charge state, ensuring that only the highest quality, fully ionized gold beam is injected into the Blue and Yellow rings of RHIC for collisions.

RHIC Main Rings

The final acceleration and storage of the beam occur in the RHIC main rings, consisting of two independent, superconducting storage rings, the "Blue" ring for clockwise and the

"Yellow" ring for counter-clockwise circulating beams. Each ring has a circumference of approximately 3834 m and utilizes 1,740 superconducting magnets. Unlike the conventional magnets in the Booster and AGS, RHIC's dipoles and quadrupoles are cooled by liquid helium to 4.5 K, allowing for a peak magnetic field of approximately 3.45 T for 100 GeV/u operation [46]. The two rings intersect at six interaction points (IPs), with the STAR experiment currently situated at IP6 and sPHENIX at IP8.

Ions are injected from the AGS via the AtR line and captured into RF buckets at harmonic number $h = 360$ using a 28 MHz RF system. Typically only ∼111 buckets are filled with circulating bunches while the remaining gaps are reserved as an "abort gap" to allow the extraction kickers to ramp up without damaging machine components. During the energy ramp, the ions must cross a transition energy of $\gamma_t \approx 22.8$. Similar to the AGS, the RF system performs a rapid synchronous phase jump at this threshold to maintain longitudinal stability [46].

Once the beam reaches the target collision energy, a "rebucketing" process occurs where the bunches are transferred from the 28 MHz acceleration system into a 197 MHz storage RF system ($h = 2520$) [46]. These higher-frequency buckets provide significantly shorter bunch lengths, which are critical for localizing the collision vertex within the narrow longitudinal acceptance of the experiments' silicon vertex detectors.

At storage energies, heavy ions are particularly susceptible to Intrabeam Scattering (IBS), which causes a rapid increase in emittance and a subsequent decay in luminosity. To counteract this, RHIC employs two world-first cooling technologies:

- **Stochastic Cooling:** High-bandwidth pickups detect stochastic fluctuations in bunch density and apply corrective kicks to compress the beam, used primarily at high energies (>100 GeV/u) [41].

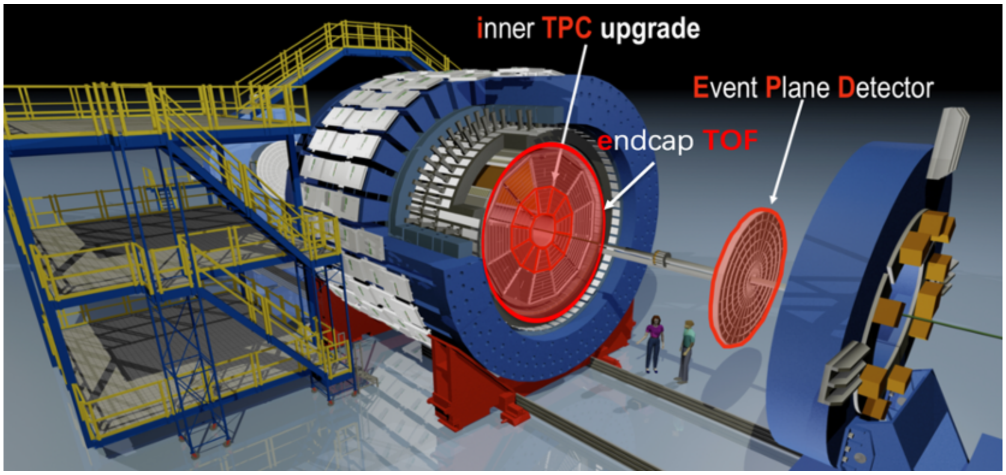

Figure 2.3: A side view of the STAR detector after the BES-II upgrades. Figure from [54].

- **Low Energy RHIC electron Cooling (LEReC):** Specifically implemented for the BES-II program, LEReC co-propagates a "cold" electron beam with the gold ions at low energies ($\sqrt{s_{NN}} = 7.7$–19.6 GeV) to absorb their thermal energy and maintain high bunch intensities [25].

2.2 The STAR Detector

STAR is a large-acceptance, multi-purpose detector system optimized for the detection of the thousands of particles produced in high-energy heavy-ion collisions. The design philosophy of STAR emphasizes tracking, momentum analysis, and particle identification (PID) at midrapidity with full azimuthal symmetry ($0 < \phi < 2\pi$). The detector is built within a large-volume solenoidal magnet that provides a uniform 0.5 T magnetic field, enabling the measurement of charged particle momenta ranging from 100 MeV/c to over 10 GeV/c [10].

2.2.1 The Main Solenoidal Magnet

The STAR Magnet is a room-temperature, iron-dominated solenoidal system that provides the necessary magnetic environment for high-precision momentum analysis and particle identification. Housed within a 1100 ton steel yoke, the magnet is designed to produce a highly uniform longitudinal field of 0.25–0.50 T. A critical physics requirement is maintaining a field uniformity of better than 1000 ppm over the active volume of the Time Projection Chamber (TPC) [19]. This level of precision is fundamental to minimizing systematic track distortions, ensuring that the spatial reconstruction of ionization clusters is not compromised by non-parallel electric and magnetic fields within the tracking volume.

The structural framework and magnetic return path are provided by a massive steel yoke consisting of 30 trapezoidal flux return bars (backlegs), four end rings, and two 73 ton pole pieces. Each 18 ton backleg is 6.85 m long and forms the outer cylindrical wall of the magnet. This iron-dominated design ensures that the flux density remains below 1.4 T throughout the yoke, resulting in a magnetic environment where field uniformity is largely independent of minor excitation fluctuations. To preserve tracking integrity, the structure is engineered to keep mechanical deflections below 1 mm under the intense magnetic pressures of operation [19].

The primary field is generated by ten Main coils and two Space Trim coils wound from large-scale aluminum conductors. These coils are built from two-layer pancakes wound in a "two-in-hand" (bifilar) fashion. To cancel undesired field components arising from layer-to-layer winding transitions, the pancakes are wound in alternating "right-hand" and "left-hand" threads [19]. In standard operation, these coils carry a current of approximately 4500 A. The Space Trim coils are utilized to flatten the magnetic field at the axial ends of the tracking volume, ensuring a more uniform drift environment for ionization electrons.

Active correction of the field line straightness is achieved via specialized Poletip Trim coils recessed into the inner face of each pole piece. These coils consist of six layers wound in a “three-in-hand” (trifilar) fashion. By utilizing these active trim coils instead of physically shaping the iron pole tips, the system can dynamically compensate for minor mechanical misalignments and achieve superior field line straightness. Thermal stability of the entire system is maintained via 88 parallel water-cooling circuits which dissipate the heat generated by the high-current operation [19]. This cooling is critical for protecting the gas-filled detectors housed within the magnet, as it prevents thermal gradients that would otherwise fluctuate the TPC gas density and negatively impact electron drift velocity (v_{drift}) and energy loss (dE/dx) calibrations.

2.2.2 The Time Projection Chamber

The TPC is the primary tracking system for the STAR experiment. It is responsible for the three-dimensional reconstruction of charged particle trajectories, momentum measurement via track curvature in the solenoidal magnetic field, and particle identification through the measurement of specific ionization energy loss (dE/dx). The TPC consists of a 4.2 m long gas-filled cylinder with an inner diameter of 1 m and an outer diameter of 4 m, providing full azimuthal coverage ($0 < \phi < 2\pi$) and a pseudorapidity range of $|\eta| < 1.0$ [10, 12].

The active volume is bisected by a Central Membrane (CM), a thin conductive cathode held at a high negative potential (typically 28 kV). A series of field-shaping rings along the inner and outer cylinder walls, in conjunction with the CM and grounded endcaps, establishes a highly uniform longitudinal electric field of approximately 135 V/cm. The volume is filled with P10 gas (90% Argon, 10% Methane) regulated at 2 mbar above atmospheric pressure. Argon acts as the primary ionization medium, while methane serves as a quencher to stabilize the electron drift velocity ($v_{drift} \approx 5.45$ cm/μs) [12].

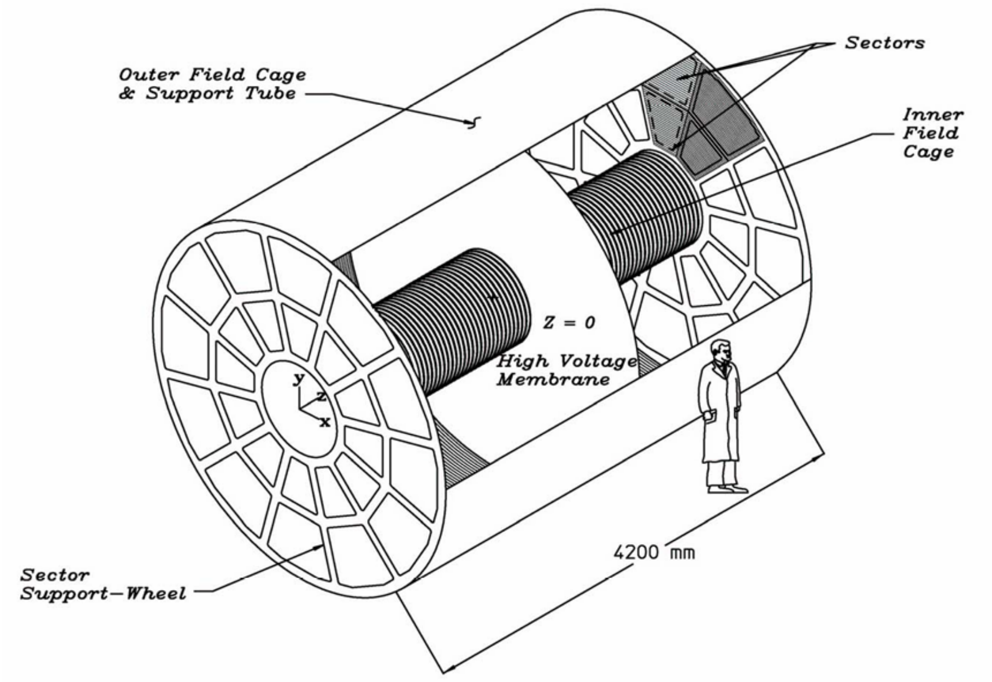

Figure 2.4: A cutaway sketch of the STAR Time Projection Chamber (TPC), illustrating the central high-voltage membrane, the inner and outer field cages, and the support wheels that house the 24 readout sectors. Figure from [12].

As a charged particle traverses the gas volume, it creates primary ionization clusters along its path. These electrons drift at a constant velocity toward the endcaps, which are divided into 24 azimuthal sectors (12 per side) [12]. Each sector contains a Multi-Wire Proportional Chamber (MWPC) readout system consisting of three wire planes: a gated grid, a ground plane, and an anode wire plane. The gated grid acts as an electronic shutter, opening to allow electrons from triggered events to pass while closing to prevent ions from the readout region from drifting back into the main volume, thereby protecting the field uniformity.

The signal is captured by a segmented pad plane behind the anode wires. The electron avalanche at the anode wires induces a charge on several adjacent pads, allowing for precise transverse (x, y) reconstruction through barycentric clusters. The longitudinal (z) coordinate is determined by the arrival time of the signal relative to the trigger. By combining these coordinates, the TPC achieves a single-point position resolution of approximately 0.5–1 mm [12]. Particle identification is performed by calculating the truncated mean of the energy loss measured by the TPC pad rows, which is then compared to the Bichsel function expectations for various particle species as a function of their reconstructed momentum to determine the $n\sigma$ deviation [20].

The inner TPC (iTPC) upgrade was implemented to meet the high-precision requirements of BES-II. While the outer sectors remained unchanged, the original 13 sparse pad rows in the inner sectors were replaced with 40 rows of smaller, more densely packed pads, providing continuous radial coverage. This increased granularity allowed the experiment to maintain a full 0.50 T magnetic field while lowering the transverse momentum threshold from 125 MeV/c to approximately 60 MeV/c. Furthermore, the iTPC extended the active pseudorapidity coverage from $|\eta| < 1.0$ to $|\eta| < 1.5$ and increased the maximum number of sampled hits along a track from 45 to 72. This improvement in sampling density enhanced the dE/dx resolution by nearly 20%, significantly bolstering particle identification capabilities for the

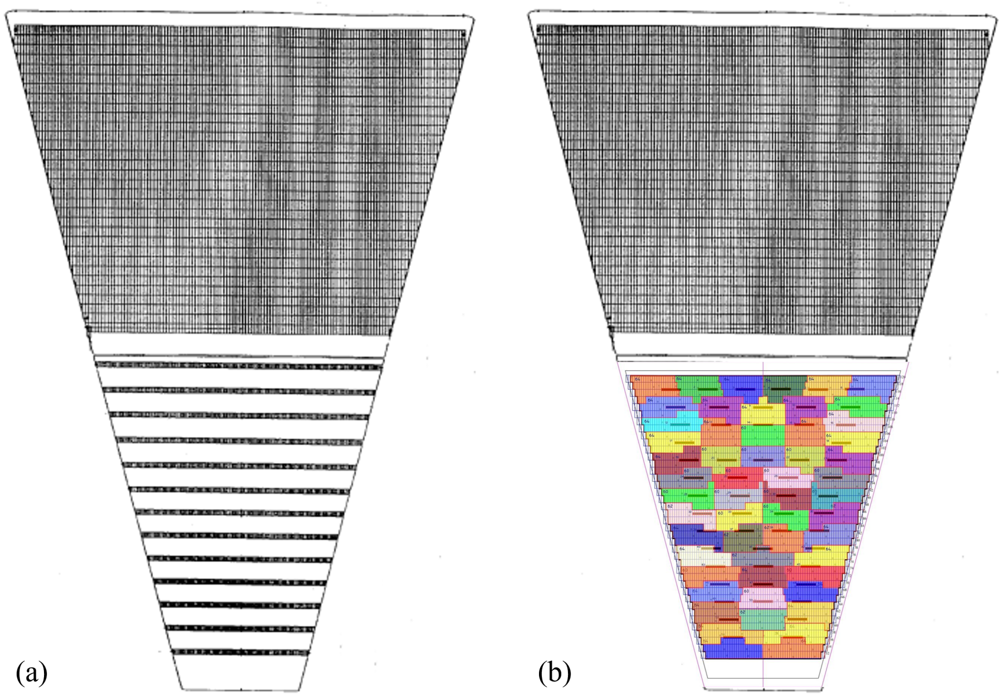

Figure 2.5: A comparison of the inner TPC sector before (a) and after (b) the iTPC upgrade. The number of pad rows was increased from 13 to 40, with the color-coding in the upgraded panel (b) illustrating the mapping to the iFEE cards. Adapted from [49]

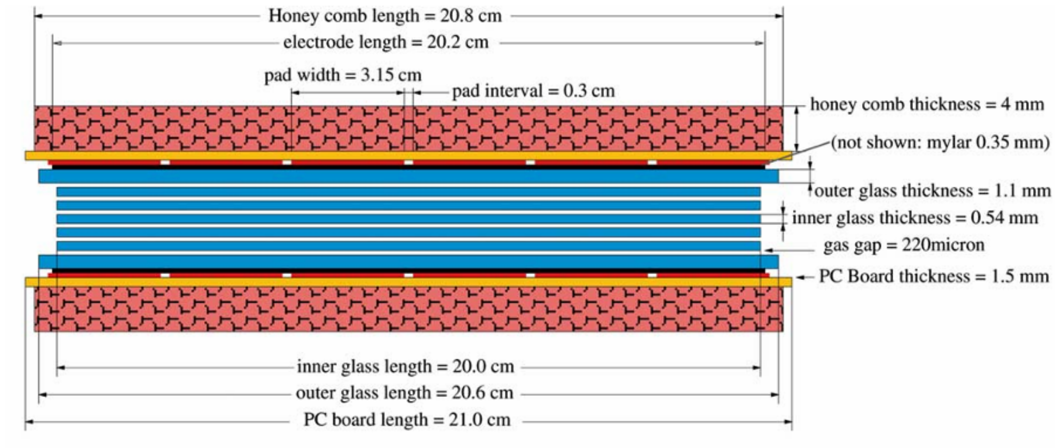

Figure 2.6: A side view of one MRPC module of the STAR bTOF. Adapted from [40].

high-baryon-density physics of the BES program [55, 49].

2.2.3 The Time-of-Flight System

The barrel Time-of-Flight (bTOF) system is a high-resolution timing detector that surrounds the TPC, providing enhanced particle identification at intermediate and high momenta. It covers the full azimuth ($0 < \phi < 2\pi$) and a pseudorapidity range of $|\eta| < 0.9$. The system is based on Multi-gap Resistive Plate Chamber (MRPC) technology, which offers a cost-effective solution for large-area coverage while maintaining an intrinsic timing resolution of better than 80 ps [39, 40].

The bTOF barrel consists of 120 trays, each containing 32 MRPC modules. Each module is constructed from a stack of six glass plates separated by thin spacers to create narrow (220 μm) gas gaps. The outer surfaces of the stack are coated with resistive graphite and held at a high differential voltage (typically ± 7 kV), creating a high electric field across the gaps. When a charged particle traverses the gas (a mixture of 95% Freon-134a and 5% Isobutane), it initiates a Townsend avalanche. The high resistivity of the glass limits the lateral spread of the discharge, while the induced signal is captured by copper readout pads [40].

The primary function of the Time-of-Flight (TOF) system is to measure the time interval (Δt) between the initial collision at the vertex and the particle's impact on the MRPC. The start time (t_0) is provided by the Vertex Position Detector (VPD), located near the beam pipe. By combining the measured time-of-flight with the track length (L) and momentum (p) determined by the TPC, the particle's mass-squared (m^2) is calculated using the relativistic relation:

$$m^2 = p^2 \left(\frac{c^2 \Delta t^2}{L^2} - 1 \right) \tag{2.3}$$

This measurement allows for clear π/K separation up to $p \approx 1.6$ GeV/c and K/p separation up to $p \approx 3.0$ GeV/c [40]. In conjunction with the TPC, the TOF system ensures robust PID over a wide kinematic range, which is critical for reconstructing the chemical and thermal properties of the medium.

2.2.4 The Event Plane Detector

The Event Plane Detector (EPD) is a high-granularity plastic scintillator detector located at forward rapidity, covering the range $2.1 < |\eta| < 5.1$. Positioned on both the East and West sides of the TPC endcaps, the EPD consists of two wheels, each divided into 12 azimuthal super-sectors. Its primary function is to provide a precise and independent measurement of the event plane angle (Ψ_n) and the collision centrality, while simultaneously serving as a high-performance trigger for the experiment [7].

Each EPD wheel is composed of 372 individual scintillator tiles (31 tiles per super-sector) with a thickness of 1.2 cm. The tiles are optically isolated and fiber-coupled to silicon photomultipliers (SiPMs) [7]. The radial segmentation of the EPD allows for a detailed measurement of the charged particle multiplicity distribution at forward rapidity. Unlike

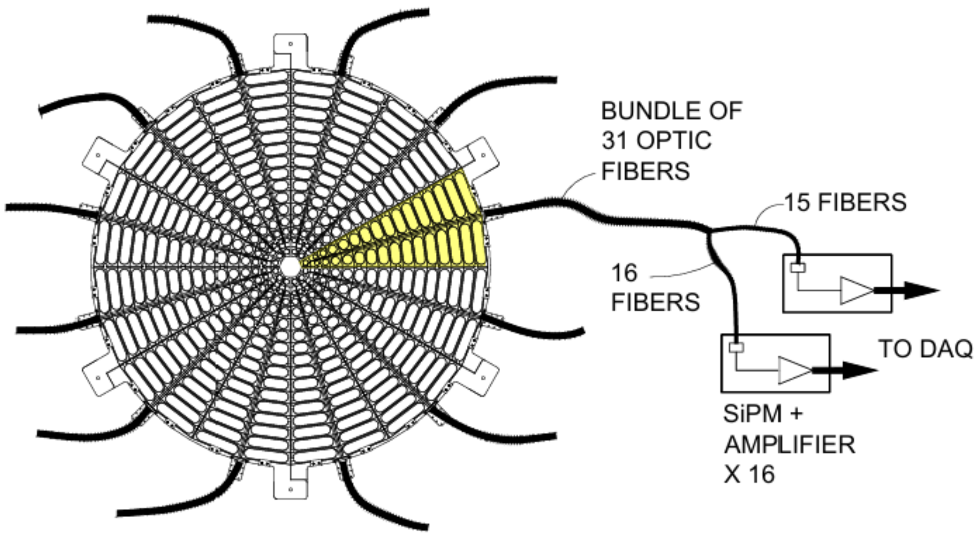

Figure 2.7: A complete Event Plane Detector (EPD) wheel. The layout consists of 12 super-sectors per wheel; one super-sector is highlighted to show the radial segmentation into 31 tiles, as well as its readout chain. Figure from [7].

the legacy detectors it replaced, the EPD's high granularity significantly reduces the "non-flow" correlations (such as di-jets or resonance decays) that can bias flow measurements when the event plane is determined using the same particles measured in the TPC.

In the context of anisotropic flow studies, the EPD provides a superior determination of the first-order event plane (Ψ_1), which is crucial for measurements of directed flow (v_1). By providing a large rapidity gap between the forward-moving remnants and the midrapidity particles tracked by the TPC, the EPD enables a cleaner separation of initial-state fluctuations from final-state collective effects. Furthermore, the EPD's fast timing and high efficiency make it an integral part of the STAR trigger system, ensuring that central and semi-central events are accurately identified and recorded at high rates.

2.2.5 The Trigger and Data Acquisition System

The STAR experiment utilizes a hierarchical, pipelined architecture to manage the disparate timescales between the RHIC bunch crossing rate (107 ns) and the varying readout speeds of its sub-detectors. While fast detectors can recover within a single clock cycle, the TPC is far slower: the full drift of the ionization electrons to the readout plane takes $\sim$40 μs, and reading the event out of the TPC takes a further $\sim$300 μs. To meet the increased demand for high-luminosity data collection, the Trigger and Data Acquisition (DAQ) systems act in coordination to identify physically significant interactions and manage the resulting data flow through a strictly timed hardware handshake [9].

Event Selection and Trigger Decision

The process begins with the Level-0 (L0) trigger. Fast trigger detector data are digitized on QT boards and passed, along with summary data from slow detectors, to DSM boards for further calculations. The results from the DSM tree are passed to the Trigger Control Unit (TCU), which makes the final trigger decision for a given bunch crossing. The TCU supports up to 64 independent triggers running simultaneously. If a collision satisfies the logic conditions and the detector is not busy, the TCU assigns a unique 12-bit Token to the event. This token acts as a global identifier, ensuring that data fragments from various subsystems can be correctly reassembled by the Event Builders.

Readout Initialization and Data Flow

Once a decision is made, the TCU notifies the slow detector sub-systems (such as the TPC and TOF) through a Trigger Clock Distribution (TCD) board. The TCD translates the L0 accept into a hardware signal that initializes the digitization and readout cycles in the

Front-End Electronics (FEE). Simultaneously, the TCU notifies the Level 0 CPU (L0), which informs the digitizer and DSM crates of the trigger through the Star Trigger Pusher (STP) network. This notification initiates the readout of the raw trigger data itself, which is passed to Level 2 using the STP network.

Data Processing and Event Building

The event data then enters the DAQ pipeline, which is highly parallelized to handle the large data volume. The TPC FEEs utilize ALTRO chips to perform pedestal subtraction and zero suppression, discarding empty channels to reduce the event size from ~ 50 MB to approximately 2–12 MB.

Data fragments are transmitted via optical fibers to sector-based readout PCs (DETs), where further data reduction, such as 2D cluster finding, is performed. These fragments are then pushed through a high-speed Ethernet network to Event Builders (EVBs). Using the 12-bit Token, the EVBs gather data from all subsystems to assemble a coherent, complete event. The assembled events are passed to the High-Level Trigger (HLT) for real-time tracking and vertex reconstruction, providing online quality assurance. Successfully processed events are finally shipped to the High Performance Storage System (HPSS) at aggregate bandwidths of up to 2000 MB/s for archival on magnetic tape. Once the event is safely buffered, the Token is released back to the TCU for reuse in a subsequent collision.

2.3 Event and Track Reconstruction

Once the raw electronic signals from a collision are recorded by the DAQ system, the data must be processed through offline reconstruction algorithms to extract physical observables.

This transformation from raw “hits” to physics-ready variables involves determining the global event geometry and identifying the properties of individual particle tracks.

2.3.1 Primary Vertex Finding

The reconstruction process begins with the identification of the primary vertex (PV), the precise location where the initial nucleon-nucleon interaction occurred. Individual tracks are first reconstructed using a Kalman filter algorithm [34] that follows the ionization clusters left in the TPC. These tracks are then extrapolated back toward the beam axis.

The PV is defined as the common point of origin for the majority of these extrapolated tracks. To ensure that the collision took place within the region of uniform detector acceptance, we require the longitudinal position to be within $|V_z| < 145$ cm for most analyzed energies; the exact per-energy vertex selections used in each analysis are listed in the corresponding chapters. Furthermore, a radial cut ($V_r < 2$ cm) is applied to exclude interactions with the beam pipe.

2.3.2 Pile-up Rejection and Centrality Determination

After vertex reconstruction, the “centrality” of the event is determined to quantify the geometric overlap of the two colliding nuclei. Since the impact parameter b is not a direct observable, STAR utilizes the charged particle multiplicity as a proxy for the collision’s “violence.”

The primary variable used is the Reference Multiplicity (RefMult), defined as the number of primary tracks in the TPC within $|\eta| < 0.5$ and a Distance of Closest Approach (DCA) to the primary vertex of less than 3 cm. Initial processing involves “pile-up” rejection to account for

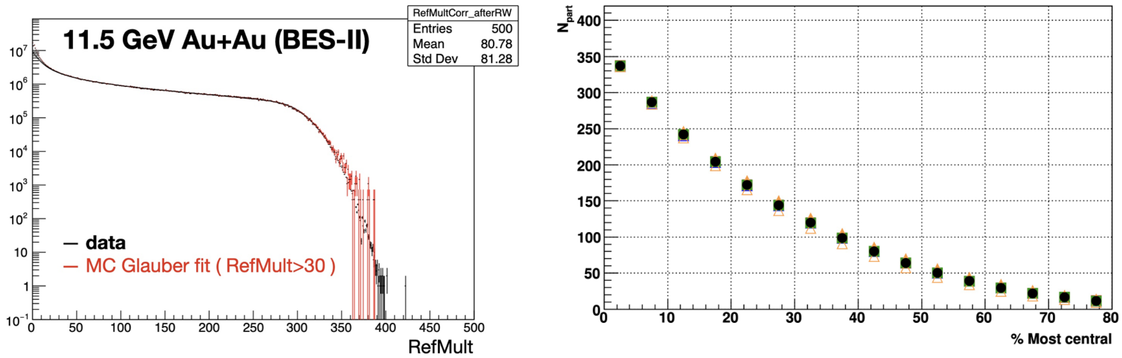

Figure 2.8: (Left) Measured corrected reference multiplicity distribution (RefMultCorr) compared with a Monte-Carlo Glauber simulation with best-fit parameters. (Right) Mapping of centrality classes to the calculated number of participating nucleons (N_{part}) based on the Glauber model thresholds. Adapted from [53].

the TPC's relatively long drift time, which can cause multiple collision events to be recorded in a single readout window. This is mitigated by correlating the TPC multiplicity with "fast" detectors such as the TOF system. Events where the TPC track count and TOF hit count are significantly mismatched are identified as pile-up and excluded from the sample.

Once the sample is cleaned, the RefMult is corrected for detector-related effects, including the V_z dependence of the TPC acceptance and run-by-run luminosity variations. This yields a corrected Reference Multiplicity (RefMultCorr). To relate RefMultCorr to the underlying geometry, a Monte Carlo Glauber model is employed to simulate the collision based on a Woods–Saxon density distribution of the nucleons in the ^{197}Au nuclei [42].

The simulated multiplicity is generated using a two-component model, which assumes particle production scales with both the number of participating nucleons (N_{part}) and the number of binary collisions (N_{coll}). This result is convolved with a Negative Binomial Distribution (NBD) to account for statistical fluctuations. A multi-dimensional parameter scan is performed to find the best fit between the simulation and the experimental RefMultCorr distribution, as shown in Fig. 2.8. The final centrality bins are then determined by integrating the RefMultCorr distribution to specific percentage thresholds.

2.3.3 Particle Identification

Following event reconstruction, the identification of individual particle species is performed using a hybrid PID strategy. This approach combines the ionization energy loss measured in the TPC with the velocity measurements from the TOF.

Energy Loss (dE/dx)

The TPC identifies particles by measuring the ionization energy loss (dE/dx) of tracks as they traverse the gas. While the standard Bethe–Bloch formula describes the *mean* rate of energy loss, it does not accurately reflect the signals in thin-absorber detectors like the P10-filled STAR TPC. Consequently, STAR utilizes the Bichsel function as the standard reference for PID, which accounts for the stochastic nature of individual collisions and the detector-specific most probable energy loss (Δ_p) for a given track segment.

This method provides excellent species separation at low momentum ($p < 0.6$ GeV/c for K/π and $p < 1.1$ GeV/c for p/π), as shown in Fig. 2.9. The deviation of the measured signal from the Bichsel expectation for a mass hypothesis X is quantified by the $n\sigma_X$ variable:

$$n\sigma_X = \frac{1}{\sigma_R} \ln\left(\frac{(dE/dx)_{meas}}{\langle dE/dx\rangle_{X,\mathrm{Bichsel}}}\right) \tag{2.4}$$

where σ_R represents the dE/dx resolution. Ideally, the $n\sigma$ distribution should follow a standard normal distribution ($N(0,1)$). To account for residual dependencies on kinematics or detector occupancy, a Gaussian fit is performed for each p_T bin. If the observed width deviates from unity, the $n\sigma$ values are scaled by the fitted standard deviation to ensure a consistent selection efficiency across the entire kinematic range.

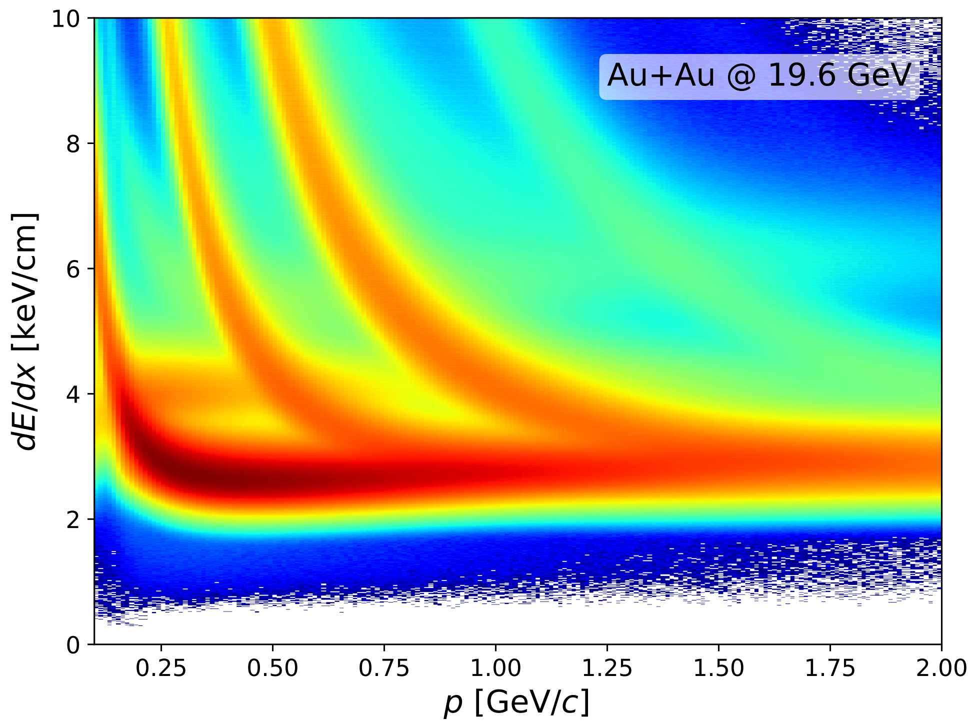

Figure 2.9: Specific energy loss (dE/dx) as a function of momentum p for charged particles in Au+Au collisions at $\sqrt{s_{NN}} = 19.6$ GeV, with the color scale logarithmic in track yield. At a given momentum the bands correspond, in order of increasing ionization, to π, K, p, and d, and each follows the Bichsel expectation for its mass.

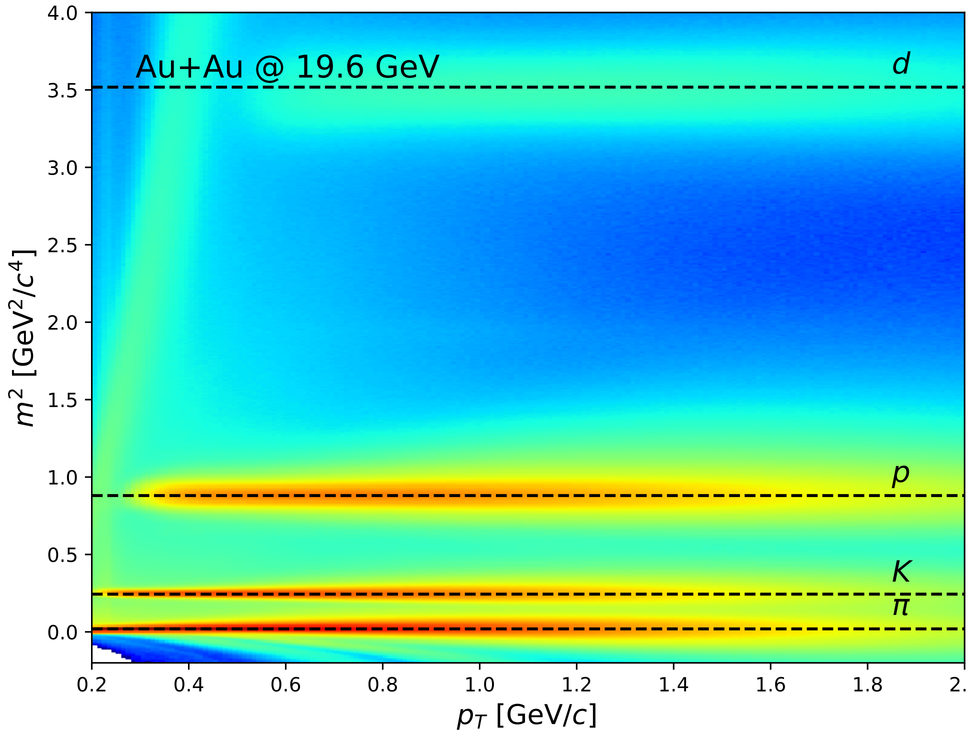

Figure 2.10: Mass squared (m^2) distribution as a function of transverse momentum p_T for charged tracks at $\sqrt{s_{NN}} = 19.6$ GeV. The horizontal dashed lines indicate the Particle Data Group (PDG) mass squared values for pions, kaons, protons, and deuterons. The Time-of-Flight (TOF) information provides significant separation between particle species, particularly in the region where dE/dx bands overlap.

Time-of-Flight

As momentum increases into the "relativistic rise" region, the dE/dx bands for different species begin to overlap. To maintain high purity, the TOF system is used to measure the particle velocity β. By combining β with the momentum p, the mass-squared (m^2) is calculated:

$$m^2 = p^2 \left(\frac{1}{\beta^2} - 1 \right) \tag{2.5}$$

The measured m^2 clusters about the Particle Data Group (PDG) value for each species, and the TOF extends the PID reach significantly, allowing for the unambiguous separation of pions and kaons up to $p \approx 1.6$ GeV/c and protons up to $p \approx 3.0$ GeV/c. However, the requirement of a valid TOF hit introduces a matching inefficiency, as not every track reconstructed in the TPC can be successfully projected and matched to a hit in the TOF trays. This results in a non-negligible loss of statistics compared to the TPC-only sample.

To maximize the available dataset while maintaining high purity, the TOF is utilized only above specific momentum thresholds tailored to each particle species. In the low-momentum region where the TPC provides sufficient dE/dx separation, TPC-only identification is prioritized to preserve statistics. The TOF is then integrated into the selection criteria at higher momenta, specifically where the Bichsel bands for different species begin to overlap, ensuring robust identification across the entire analyzed transverse momentum range.

2.3.4 The Kalman Filter

The Kalman filter is a recursive algorithm for estimating the state of a dynamic system from a sequence of noisy measurements [34]. It is employed extensively in high-energy physics for track fitting, vertex reconstruction, and short-lived particle reconstruction. This section summarizes the general formalism; its specific application to particle reconstruction is described in Sect. 2.3.5.

Consider a system whose state at step k is described by a vector $\mathbf{r}_k$ with covariance $\mathbf{C}_k$. A measurement $\mathbf{m}_k$ with covariance $\mathbf{V}_k$ is related to the state by a measurement model $\mathbf{H}_k$:

$$\mathbf{m}_k = \mathbf{H}_k\,\mathbf{r}_k + \boldsymbol{\epsilon}_k, \qquad \boldsymbol{\epsilon}_k \sim \mathcal{N}(\mathbf{0},\, \mathbf{V}_k). \tag{2.6}$$

Given a prior estimate $\mathbf{r}_{k-1}$ with covariance $\mathbf{C}_{k-1}$, the filter incorporates the new measurement through the following update. The residual (innovation) between the measurement and the predicted observation is

$$\boldsymbol{\zeta}_k = \mathbf{m}_k - \mathbf{H}_k\, \mathbf{r}_{k-1}. \tag{2.7}$$

The innovation covariance is

$$\mathbf{S}_k = \mathbf{H}_k\, \mathbf{C}_{k-1}\, \mathbf{H}_k^\top + \mathbf{V}_k, \tag{2.8}$$

and the Kalman gain, which optimally weights the new information against the prior, is

$$\mathbf{K}_k = \mathbf{C}_{k-1}\, \mathbf{H}_k^\top\, \mathbf{S}_k^{-1}. \tag{2.9}$$

The updated state and covariance are then

$$\mathbf{r}_k = \mathbf{r}_{k-1} + \mathbf{K}_k\, \boldsymbol{\zeta}_k, \tag{2.10}$$

$$\mathbf{C}_k = (\mathbf{I} - \mathbf{K}_k\, \mathbf{H}_k)\, \mathbf{C}_{k-1}, \tag{2.11}$$

and the χ^2 contribution of the measurement is $\Delta\chi^2_k = \boldsymbol{\zeta}_k^\top\, \mathbf{S}_k^{-1}\, \boldsymbol{\zeta}_k$.

In STAR, the Kalman filter is applied at several stages of the reconstruction chain. During track fitting, the filter follows ionization clusters left in the TPC, sequentially updating the helix parameters and their covariance as each hit layer is incorporated [12]. For vertex reconstruction, tracks are extrapolated to a common point and iteratively fitted using a Kalman-based algorithm that rejects outliers via χ^2 cuts (Sect. 2.3). For short-lived particles that decay before reaching the detector, the same mathematical framework is adapted to reconstruct the decay vertex and the mother particle's full kinematics from its daughter tracks, as described next.

2.3.5 Short-Lived Particle Reconstruction with KFParticle

Many particles of physics interest, such as the Λ hyperon, K^0_S, Ξ, and Ω, decay via the weak or electromagnetic interaction before reaching the active detector volume. These short-lived particles must be reconstructed from their charged decay products. The KFParticle package [30] adapts the Kalman filter formalism of Sect. 2.3.4 to this problem, with two key modifications: (i) the measurement model is chosen so that the Kalman update constrains only the *spatial* vertex, while daughter momenta are accumulated additively by four-momentum conservation; and (ii) the unknown transport distance of each daughter to the decay vertex is absorbed into the measurement covariance, allowing the vertex position to be determined purely from the intersection of multiple tracks. This framework is used in this thesis for the reconstruction of Λ hyperons and Ω baryons.

State Vector and Covariance

Each particle is described by an eight-parameter state vector:

$$\mathbf{r} = (x,\, y,\, z,\, p_x,\, p_y,\, p_z,\, E,\, s)^\top \equiv \begin{pmatrix} \mathbf{v} \\ \underline{\mathbf{p}} \\ s \end{pmatrix}, \tag{2.12}$$

where $\mathbf{v} = (x, y, z)^\top$ is the spatial position, $\underline{\mathbf{p}} = (p_x, p_y, p_z, E)^\top$ the four-momentum with $E = \sqrt{M^2 + |\mathbf{p}|^2}$ for an assumed mass hypothesis M, and $s = l/|\mathbf{p}|$ is the normalized decay length (flight path l divided by the momentum magnitude). This detector-independent parametrization allows derived quantities (invariant mass, total momentum, flight length, and proper lifetime) to be obtained from the fitted state via Jacobian projections [30]. The full uncertainty is carried in the symmetric 8×8 covariance matrix $\mathbf{C}$. When a reconstructed

TPC track is promoted to a KFParticle object, its six helix parameters and 6×6 covariance are extended to the eight-dimensional representation using $\partial E/\partial p_i = p_i/E$; the decay-length parameter is initialized to $s = 0$ with $\sigma_s^2 \to \infty$.

Daughter Transport

Before two daughter tracks can be combined at a common vertex, each must be propagated from its point of measurement to the estimated decay point. In the uniform solenoidal field of the STAR magnet ($B_z = 0.5$ T), this transport follows a helical trajectory and is performed analytically, with the covariance propagated via

$$\mathbf{C}' = \mathbf{J}\,\mathbf{C}\,\mathbf{J}^{\top}, \tag{2.13}$$

where $\mathbf{J}$ is the Jacobian of the transport equations with respect to the state parameters.

Because the arc-length s_k^d from each daughter's parametrization point to the unknown decay vertex is itself unknown ($\langle s_k^d \rangle = 0$, $\sigma^2_{s_k^d} \to \infty$), the daughter's measurement covariance at the vertex acquires an additional term:

$$\mathbf{V}_k = \mathbf{C}_k^d + \begin{pmatrix} \mathbf{p}_k^d \\ \mathbf{p}_k^d \times \mathbf{B}\,q_k \\ 0 \end{pmatrix} \begin{pmatrix} \mathbf{p}_k^d \\ \mathbf{p}_k^d \times \mathbf{B}\,q_k \\ 0 \end{pmatrix}^{\top} \cdot\ \sigma^2_{s_k^d}, \tag{2.14}$$

where $\mathbf{B}$ is the local magnetic field and q_k the daughter's charge. This inflation renders the daughter's position uncertainty effectively infinite, so that no single track can determine the vertex; only the *intersection* of multiple daughters constrains it.

Vertex Fit

The vertex fit operates on the seven-parameter subvector $\mathbf{r}_k = (\mathbf{v}_k,\, \underline{\mathbf{p}}_k)^\top$. The measurement model $\mathbf{H} = (\mathbf{I}_{3\times3},\, \mathbf{0}_{3\times4})$ projects the state onto the position block. Because $\mathbf{H}$ is trivially the identity on $\mathbf{v}$, the innovation covariance (Eq. 2.8) reduces to a 3×3 matrix:

$$\mathbf{S}_k = \left(\mathbf{C}^v_{k-1} + \mathbf{V}^v_k\right)^{-1}, \tag{2.15}$$

where $\mathbf{C}^v_{k-1}$ and $\mathbf{V}^v_k$ are the position blocks of the mother and daughter covariances, respectively, and the inversion is performed via Cholesky decomposition.

Since no prior estimate of the mother exists, the first daughter ($k = 1$) is simply copied: $\mathbf{r}_1 = \mathbf{m}_1$, $\mathbf{C}_1 = \mathbf{V}_1$, with $\chi^2_1 = 0$ and $\mathrm{ndf}_1 = -1$ (a single track cannot define a vertex). For each subsequent daughter $k \geq 2$, the position residual $\boldsymbol{\zeta}_k = \mathbf{v}^d_k - \mathbf{v}_{k-1}$ drives the update:

$$\mathbf{v}_k = \mathbf{v}_{k-1} + \mathbf{C}^v_{k-1}\,\mathbf{S}_k\,\boldsymbol{\zeta}_k, \tag{2.16}$$

$$\underline{\mathbf{p}}_k = \underline{\mathbf{p}}_{k-1} + \underline{\mathbf{p}}^d_k + \left(\mathbf{C}^{vp}_{k-1} - \mathbf{V}^{vp}_k\right)\mathbf{S}_k\,\boldsymbol{\zeta}_k. \tag{2.17}$$

The daughter's four-momentum $\underline{\mathbf{p}}^d_k$ enters *additively* in Eq. 2.17: the Kalman filter constrains only the decay vertex geometry, while the mother's momentum is the sum of its daughters' momenta. Each daughter contributes $\Delta\mathrm{ndf} = 2$ (three spatial measurements minus one for the unknown arc-length s^d_k) and $\Delta\chi^2_k = \boldsymbol{\zeta}_k^\top\,\mathbf{S}_k\,\boldsymbol{\zeta}_k$.

Mass Constraint and Production Vertex Fit

After all daughters are added, the precision of the mother's parameters can be improved if the invariant mass is known. The constraint $E^2 - |\mathbf{p}|^2 = M^2_{\mathrm{PDG}}$ is imposed as a one-dimensional

Kalman measurement (Eq. 2.10) with measurement model

$$\mathbf{H}_{M^2} = \big(0,\, 0,\, 0,\, -p_x,\, -p_y,\, -p_z,\, E,\, 0\big) \tag{2.18}$$

and zero measurement error.

Optionally, the reconstructed mother is transported from the decay vertex back toward its production vertex (typically the primary vertex $\mathbf{v}^p$) using the arc-length parameter s. The primary vertex position and its covariance then serve as a further Kalman measurement with $\mathbf{H} = (\mathbf{I},\, \mathbf{0})$, adding $\Delta\text{ndf} = 2$. The resulting topological χ^2 quantifies how well the mother's flight direction points back to the collision point.

Topological Selection of Decay Candidates

The vertex-fitting procedure is applied to all combinatorial pairs (or triplets, for cascades) of daughter candidates. Neutral particles that decay into two oppositely charged daughters, such as $\Lambda^0 \to p\pi^-$ and $K_S^0 \to \pi^+\pi^-$, produce a characteristic "V"-shaped topology displaced from the primary vertex and are known as V0 decays. Charged cascades such as $\Omega^- \to \Lambda K^-$ involve an intermediate V0 combined with a bachelor track at a secondary cascade vertex. In both cases, the following sequence of topological cuts is applied to suppress the combinatorial background:

1. **Distance of closest approach:** The DCA between the two daughter helices must satisfy $d_{\text{DCA}} < d_{\max}$, ensuring geometric consistency with a common decay point.

2. **Vertex quality:** The χ^2/NDF of the reconstructed decay vertex must be below a threshold, testing the compatibility of the daughter tracks with a single spatial origin.

3. **Daughter displacement from the primary vertex:** Each daughter must have a

large χ^2 with respect to the primary vertex,

$$\chi^2_{\mathrm{PV}} = \Delta\mathbf{r}^\top \left(\mathbf{C}^v + \mathbf{C}^v_{\mathrm{PV}}\right)^{-1} \Delta\mathbf{r} > \chi^2_{\min}, \tag{2.19}$$

confirming that it is inconsistent with being a prompt track.

4. **Decay length and significance:** The flight distance L from the primary vertex to the decay vertex must exceed a minimum value, and its significance L/σ_L must be large.

Candidates passing all cuts are stored both with and without the mass constraint: the unconstrained version serves as input for subsequent cascade reconstruction, while the mass-constrained version is used directly for physics analyses. The specific numerical values of these topological cuts are analysis-dependent and are listed in the relevant methodology chapters.

Chapter 3

Constituent Quark Dynamics and Collective Flow in BES-II

3.1 Introduction

The primary objective of the BES program at RHIC is to explore the QCD phase diagram and characterize the properties of the QGP. While previous studies have firmly established *Quark Coalescence* as the dominant mechanism for hadron production at intermediate p_T via the observation of NCQ scaling [2, 26, 27, 31, 36], a more nuanced question remains: how do specific flavor species and their respective origins (transported vs. pair-produced) dictate the collective expansion?

This chapter seeks to extend the coalescence sum rule to distinguish between the dynamics of transported valence quarks from the initial nuclei and pair-produced "sea" quarks from the vacuum. Empirically, the coalescence model describes the combination of quarks moving in the same direction with similar momenta into final-state hadrons [26]. Under this constituent quark hypothesis, the collective flow of a hadron is essentially the sum of the flow of its valence constituents:

$$v_n^{hadron}(p_T) \approx \sum_{i=1}^{n_q} v_n^{q_i}(p_T/n_q) \tag{3.1}$$

where n_q is the number of constituent quarks (2 for mesons, 3 for baryons) and q_i represents both the species and the origin of the quark [36].

We apply this framework to isolate the initial-state isospin asymmetry. Specifically, we use the elliptic flow (v_2) of π^+ ($u\bar{d}$) and π^- ($\bar{u}d$) to isolate the average number of transported u and d quarks. By extracting the ratio of these transported constituents and comparing it to the initial neutron-to-proton stoichiometry of the ^{197}Au nuclei:

$$\frac{\langle N_d^{trans} \rangle}{\langle N_u^{trans} \rangle} = \frac{2N+Z}{N+2Z} = \frac{315}{276} \approx 1.14 \tag{3.2}$$

we illustrate that the collective expansion preserves a clear signature of the initial-state isospin. This consistency highlights the capability of the coalescence model to account for the distinct partonic history of transported versus produced quarks.

Furthermore, we investigate the interplay between this coalescence mechanism and the intense, transient electromagnetic (EM) fields generated by spectator protons. The neutral Λ hyperon (uds) provides a critical test case; despite its overall neutrality, its constituent u, d, and s quarks carry electric charges and are susceptible to EM-induced forces before they coalesce. By testing if the Λ directed flow splitting (Δv_1) is consistent with the summed EM-response of its constituent quarks, proxied by the combination of proton and kaon flow, we probe whether these quarks act as independent degrees of freedom that interact with the EM field during the pre-hadronization phase.

Together, these analyses aim to provide a unified picture of coalescence in which hadron flow is governed not only by the number of constituent quarks, but by their flavor and partonic origin.

3.2 Data Sample and Event Selection

Both analyses presented in this chapter draw on Au+Au collision data collected by the STAR detector at RHIC during the BES-II program, spanning center-of-mass energies $\sqrt{s_{NN}} = 7.7$–27 GeV. The data were recorded with a minimum-bias trigger. To ensure the integrity of the physics observables and a stable detector acceptance, event-level selection criteria are applied, with the trigger IDs and the vertex constraints for each collision energy summarized in Table 3.1.

Table 3.1: Summary of the trigger IDs and event-level vertex cuts for different Au + Au collision energies.

$\sqrt{s_{NN}}$ **(GeV)**	**Trigger ID**	$\lvert V_z \rvert$ **(cm)**	V_r **(cm)**
7.7	810010, 810020, 810030, 810040	< 145	< 2
9.2	780010, 780020	< 145	< 2
11.5	710000, 710010, 710020	< 145	< 2
14.6	650000	< 145	< 2
17.3	870010	< 145	< 2
19.6	640001, 640011, 640021, 640031, 640041, 640051	< 70	< 2
27	610001, 610011, 610021, 610031, 610041, 610051	< 70	< 2

In addition to vertex positioning, several quality assurance (QA) procedures were implemented to remove non-hadronic backgrounds and instrumental artifacts:

- **Pile-up Rejection:** Out-of-time pile-up events, resulting from collisions occurring within the TPC drift window but outside the trigger time, were rejected. This was achieved by analyzing the correlation between the number of tracks matched to the TOF trays and the total TPC track multiplicity, following standard STAR criteria.

- **Bad Run Removal:** Data periods characterized by unstable detector conditions, hardware trips, or abnormal gain calibrations were identified through a "bad run" list and excluded from the analysis to ensure the stability of the flow observables. For this analysis, the official STAR bad run lists are used.

Together, these criteria define the common event sample on which both the pion elliptic flow isospin analysis (Sect. 3.3) and the Λ directed flow analysis (Sect. 3.4) are based. Track-level selection criteria specific to each analysis are described in their respective sections.

3.3 Isospin Asymmetry in Pion Elliptic Flows

3.3.1 Theoretical Framework

In the study of collective flow in heavy-ion collisions, a common observation is that particles carry larger elliptic flow (v_2) than their antiparticles [4]. This phenomenon is rooted in the "transported quark effect": quarks originating from the colliding nuclei are present from the onset of the collision, experiencing the full development of the pressure gradients. In contrast, pair-produced quarks from the vacuum appear later in the evolution and thus develop weaker flow.

For the proton and antiproton system, this distinction is absolute, as the antiproton contains no transported quarks. However, the situation for pions is more nuanced. Both the π^- ($\bar{u}d$) and the π^+ ($u\bar{d}$) contain one valence quark that can be a transported quark from the initial nuclei. Because both species should benefit from this transport effect, any difference in their flow is often assumed to be driven by their opposite electric charges. Specifically, the Chiral Magnetic Wave (CMW) is typically invoked to explain the pion asymmetry [5]; it posits that the interplay between intense magnetic fields and the chiral anomaly induces a finite electric

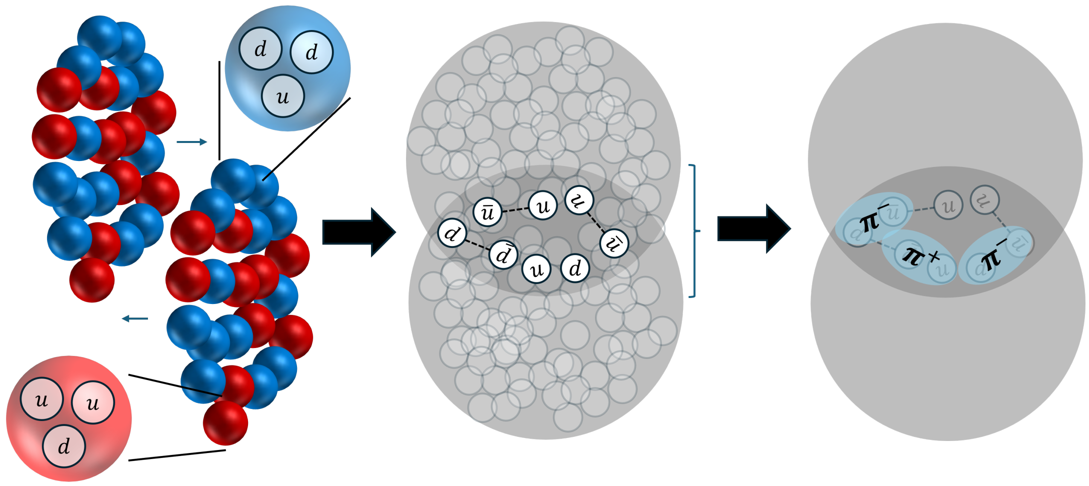

Figure 3.1: A schematic illustration of pion formation via quark coalescence in non-central collisions. Linked quarks in the middle panel represent pair-produced pairs while separate quarks denote transported quarks originating from the colliding nuclei. Although pions may be formed from various combinations of these sources, the initial dynamical properties and flavor abundance of the transported quarks are, on average, imprinted onto the final state hadrons.

quadrupole moment, leading to a charge-dependent splitting of elliptic flow [22].

We propose that this splitting has a natural origin in flavor abundance. While both π^- and π^+ can contain up to one transported valence quark, the probability of them doing so is not equal. Because the Gold nucleus is neutron-rich ($N/Z \approx 1.5$), the initial state contains a significant surplus of transported d quarks compared to u quarks. Within the coalescence framework, this higher availability means that a larger fraction of the π^- population is formed using "high-flow" transported d quarks. Conversely, the smaller reservoir of u quarks means a higher proportion of π^+ must be formed from "low-flow" produced quarks. Consequently, the π^- flow is naturally enhanced relative to the π^+ flow. In this view, the observed splitting is a direct translation of the initial isospin asymmetry, the d/u ratio of the colliding nuclei, into the final-state momentum observables.

Additionally, the spatial distribution of these nucleons plays a crucial role in the centrality dependence of the observed flow splitting. If a "neutron skin" effect is present [50], where neutrons are distributed with a larger radius than protons, the initial overlap region in peripheral collisions is expected to be even more neutron-rich than in central collisions. This would result in an even stronger isospin asymmetry in the peripheral region. Furthermore, other factors such as nuclear deformation [43] may contribute to the variation of this effect across different centralities, as the specific orientation and shape of the nuclei alter the local d/u ratio in the partonic medium.

3.3.2 Analysis Method

To quantify how this initial flavor imbalance translates to the final state, we utilize the quark coalescence sum rule. We decompose the measured elliptic flow of π^- and π^+ into the contributions from transported quarks (q_{tr}) and produced sea quarks (q_{pr}):

$$v_2(\pi^-) = N_d^{tr} v_2(q_{tr}) + (2 - N_d^{tr}) v_2(q_{pr}) \tag{3.3}$$

$$v_2(\pi^+) = N_u^{tr} v_2(q_{tr}) + (2 - N_u^{tr}) v_2(q_{pr}) \tag{3.4}$$

where N_q^{tr} represents the average number of transported quarks of flavor q. In this formulation, we assume that the flow of all transported quarks is identical ($v_2(d_{tr}) = v_2(u_{tr}) = v_2(q_{tr})$) and that all produced quarks share a common flow $v_2(q_{pr})$. To eliminate the unknown $v_2(q_{pr})$, we utilize the antiproton ($\bar{p}$) as a proxy for the produced quark flow. Since the antiproton is composed of three produced antiquarks, its flow is given by:

$$v_2(\bar{p}) = 3v_2(q_{pr}) \implies v_2(q_{pr}) = \frac{1}{3} v_2(\bar{p}) \tag{3.5}$$

Substituting this into equations 3.3 and 3.4, we obtain:

$$v_2(\pi^-) = N_d^{tr} v_2(q_{tr}) + \frac{2 - N_d^{tr}}{3} v_2(\bar{p}) \tag{3.6}$$

$$v_2(\pi^+) = N_u^{tr} v_2(q_{tr}) + \frac{2 - N_u^{tr}}{3} v_2(\bar{p}) \tag{3.7}$$

Then,

$$\frac{\dfrac{v_2(\pi^-)}{2} - \dfrac{v_2(\bar{p})}{3}}{\dfrac{v_2(\pi^+)}{2} - \dfrac{v_2(\bar{p})}{3}} = \frac{N_d^{tr}}{N_u^{tr}} \tag{3.8}$$

If the coalescence model holds, this ratio should reflect the initial stoichiometry of the Gold nuclei (≈ 1.14), providing a direct link between the initial-state flavor geometry and the final-state collective expansion.

3.3.3 Track Selection and Particle Identification

Table 3.2: Track quality selection criteria applied to all primary tracks for the pion coalescence analysis.

Cut	Value
Hits used in track fit, N_{fit}	≥ 15
Hits contributing to dE/dx, $N_{dE/dx}$	≥ 15
Hit ratio, $N_{\text{fit}}/N_{\text{max}}$	> 0.52
DCA to primary vertex	< 3.0 cm

Primary tracks used in the pion coalescence analysis must satisfy the quality criteria listed in Table 3.2. The requirements on hits used in track fitting and energy loss measurements ensure sufficient track quality for accurate momentum and particle identification. The hit-ratio cut suppresses track splitting, in which a single particle trajectory is reconstructed as two shorter segments. The DCA cut suppresses secondary tracks from weak decays and interactions with detector material.

Pions, protons, and antiprotons are then identified using the TPC ionization energy loss (dE/dx) and, above certain threshold transverse momenta, the TOF detector information. To avoid possible deviation from a normally distributed $n\sigma$ distribution, the raw $n\sigma$ values are further fitted with a Gaussian function, and the mean and sigma from the fit are used to recalculate the $n\sigma$ values for the final PID selection. The m^2 window cut is applied to the mass-squared distribution calculated from TOF information, which provides additional separation between species at higher momenta where dE/dx bands overlap.

To ensure even acceptance in the TOF detector, identified particles are restricted to the pseudorapidity range $|\eta| < 0.6$. The transverse momentum range for each particle is selected so that the momentum range per constituent quark is the same across mesons and baryons, in accordance with the coalescence model. The detailed PID and kinematic cuts are listed in Table 3.3.

Table 3.3: Particle identification cuts for the pion v_2 analysis. The $n\sigma$ cut applies to the relevant species hypothesis in all cases.

Particle	p_T range (GeV/c)	η range	TOF required above (GeV/c)	m^2 window (GeV^2/c^4)	$\lvert n\sigma\rvert$ (normalized)
$\pi^\pm$	0.16–1.2	$\lvert\eta\rvert < 0.6$	0.4	$[-0.01,\ 0.10]$	< 2
$p/\bar{p}$	0.24–1.8	$\lvert\eta\rvert < 0.6$	0.6	$[0.75,\ 1.10]$	< 2

3.3.4 Event Plane Determination

Elliptic flow is measured relative to the second-order event plane Ψ_2. The primary event plane is reconstructed from TPC tracks at all collision energies; an independent EPD-based measurement is performed at $\sqrt{s_{NN}} \geq 14.6\,\text{GeV}$ for comparison.

TPC event plane. The second-order event plane is constructed from TPC tracks in the pseudorapidity range $|\eta| > 0.5$, split into east ($\eta < -0.5$) and west ($\eta > 0.5$) sub-events. The optimal weight for event plane reconstruction is the flow signal v_n itself; for the second-order TPC plane we use $w = p_T$ as a practical proxy, since v_2 scales approximately linearly with p_T at low momenta. To suppress non-flow effects between the event plane and the particles of interest ($|\eta| < 0.6$), each particle's v_2 is evaluated using the event plane reconstructed from the opposite sub-event: particles in the east hemisphere ($\eta < 0$) use the west plane and particles in the west hemisphere ($\eta > 0$) use the east plane.

Phi-weight and Fourier shift corrections are applied day-by-day and centrality-by-centrality to flatten the azimuthal distribution of the reconstructed event plane.

The event plane resolution is determined via the east–west sub-event method:

$$\mathrm{Res}(\Psi_{2,\mathrm{TPC}}) = \sqrt{\left\langle \cos\left(2\Psi_{2,\mathrm{east}} - 2\Psi_{2,\mathrm{west}}\right)\right\rangle}. \tag{3.9}$$

EPD event plane ($\sqrt{s_{NN}} \geq 14.6\,$GeV). At collision energies of 14.6 GeV and above, the second-order event plane is also reconstructed from EPD participant tiles. Following the same principle of using v_n as the optimal weight, each EPD tile is weighted by its measured v_2 signal with respect to a TPC reference event plane; this produces a data-driven, η-dependent weight that reflects the actual elliptic flow modulation across the EPD acceptance. Tiles with $|\eta|$ above an energy-dependent threshold (2.7, 2.9, 3.0, and 3.4 at $\sqrt{s_{NN}} = 14.6$, 17.3, 19.6, and 27 GeV) are assigned zero weight, restricting the plane to the participant region. The same phi-weight and shift corrections are applied day-by-day and centrality-by-centrality. The EPD-based v_2 results are presented alongside the TPC results.

The elliptic flow is then:

$$v_2 = \frac{\big\langle \cos\big(2(\phi - \Psi_2)\big)\big\rangle}{\mathrm{Res}(\Psi_2)}. \tag{3.10}$$

3.3.5 Systematic Uncertainties

The systematic uncertainty on the pion v_2 and the extracted coalescence ratio $N_d^{\mathrm{tr}}/N_u^{\mathrm{tr}}$ is assessed by repeating the analysis with varied selection criteria. The sources and their variations are listed in Table 3.4.

Table 3.4: Systematic uncertainty sources and their variations for the pion v_2 analysis.

#	Source	Default	Variation
1	Vertex z position	$\lvert V_z\rvert < 145$ (70) cm	$\lvert V_z\rvert < 70$ (35) cm
2	N_{fit}	≥ 15	≥ 20
3	Primary-track PID $\lvert n\sigma\rvert$	< 2	< 3

For each variation the full analysis chain is repeated and the resulting v_2 and coalescence ratio are compared with the default values. The significance of each deviation is evaluated using the Barlow test: a variation is considered a significant source of systematic uncertainty if

$$|\Delta| > \sigma_\Delta, \qquad \sigma_\Delta = \sqrt{|\sigma_{\mathrm{var}}^2 - \sigma_{\mathrm{def}}^2|}, \tag{3.11}$$

where Δ is the difference between the varied and default results and σ_{var}, σ_{def} are the corresponding statistical uncertainties. The total systematic uncertainty is taken as the quadrature sum of all significant deviations, divided by $\sqrt{3}$ under the assumption that each variation samples a uniform distribution.

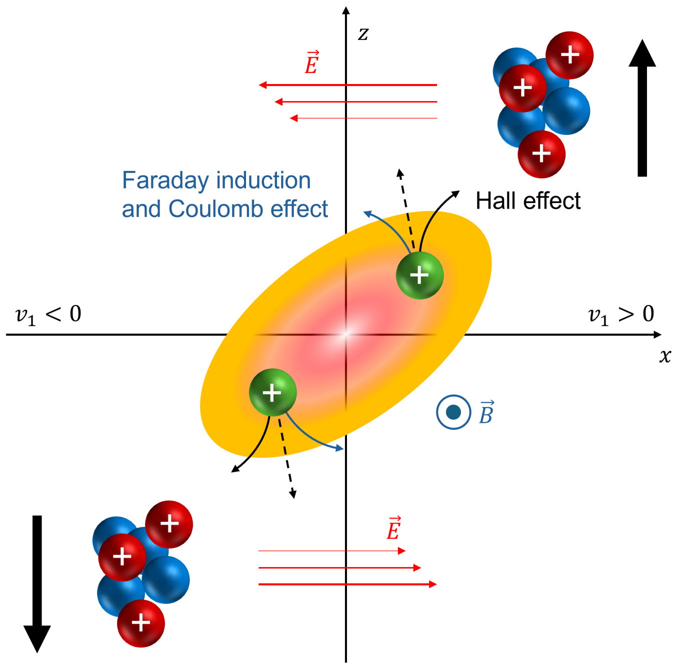

Figure 3.2: A schematic illustration of a non-central heavy-ion collision viewed along the y-axis (perpendicular to the reaction plane). While the dashed lines represent the trajectory of a positively charged quark driven by QGP expansion, the solid curves denote the deviations induced by specific electromagnetic phenomena.

3.4 Directed Flow Splitting of Λ Hyperons

3.4.1 Electromagnetic Effects

In non-central heavy-ion collisions, the spectator protons passing by the collision zone generate an intense, transient magnetic field [23]. The interaction between this electromagnetic field and the expanding QGP induces several distinct effects on the momentum of charged particles. These effects are fundamentally driven by the electric charge of the constituent quarks and are often categorized into the following three mechanisms [32, 33]:

- **The Coulomb Effect:** The net positive charge of the spectator remnants creates a repulsive electric field. In the geometry of Fig. 3.2, where the forward spectators

($z > 0$) are located at positive x, this field pushes positively charged quarks toward the negative x-direction at forward longitudinal positions. Conversely, at backward positions ($z < 0$), positive quarks are pushed toward the positive x-direction.

- **The Faraday Effect:** As the magnetic field strength rapidly decreases, the induced electric field opposes the change in magnetic flux. This induced field drives a current that, for the described setup, pushes positive charges toward the $-x$ direction at $z > 0$ and toward the $+x$ direction at $z < 0$. Thus, the Faraday effect acts in the same direction as the Coulomb deflection.

- **The Hall Effect:** Quarks moving with longitudinal velocity (v_z) through the transverse magnetic field (B_y) experience a Lorentz force ($F_x = -qv_zB_y$). With the magnetic field pointing in the $-y$ direction, positive quarks at forward positions ($v_z > 0$) are deflected toward the $+x$ direction, while those at backward positions ($v_z < 0$) are deflected toward the $-x$ direction. The Hall effect works against the Coulomb and Faraday effects.

Building on these spatial deflections, the collective motion of the quarks can be well quantified by the first-order harmonic coefficient of the Fourier expansion of the azimuthal distribution, known as directed flow (v_1). In this study, we focus strictly on the rapidity-odd component of v_1, which describes the collective "tilt" of the expanding medium. At midrapidity, this distribution is typically linear, and we characterize the strength of this motion by the slope at $y = 0$, denoted as dv_1/dy.

In the established geometry, the Hall effect deflects positive quarks toward the $+x$ direction at forward rapidity ($y > 0$), contributing a positive dv_1/dy. Conversely, the Coulomb and Faraday effects deflect them toward the $-x$ direction, contributing a negative dv_1/dy. To isolate these electromagnetic contributions from the charge-independent hydrodynamic flow of the bulk medium, we define the directed flow splitting, $\Delta(dv_1/dy)$. This is the difference between the slopes of positively and negatively charged species (or, in the case of neutral Λ, between Λ^0 and $\bar{\Lambda}^0$):

$$\Delta(dv_1/dy) = (dv_1/dy)_{q>0} - (dv_1/dy)_{q<0} \tag{3.12}$$

By taking this difference, the symmetric hydrodynamic background effects, which affect $q > 0$ and $q < 0$ particles identically, cancel out, leaving a signature that is sensitive to the early-stage electromagnetic field. A positive $\Delta(dv_1/dy)$ would indicate that the Hall effect (Lorentz force) dominates the interaction, while a negative value suggests that the combined influence of the Coulomb and Faraday fields has a greater impact on the final-state quark distribution. Regardless of which effect dominates, the electromagnetically induced v_1 splitting is expected to increase as the collision becomes more peripheral due to an increased quantity of spectators.

In addition to electromagnetic effects, the measured directed flow splitting is significantly influenced by the transported quark effect. Quarks transported from the initial-state nuclei experience the entire system evolution and thus acquire stronger flows compared to quarks produced later in the collision. This effect is particularly pronounced for protons, which may contain up to three transported quarks, whereas the $\bar{p}$ is composed entirely of pair-produced antiquarks. This asymmetry in the initial conditions contributes a positive background to the splitting across all collision centralities.

A recent STAR study [1] provided direct evidence of these EM-field effects, reporting measurements of $\Delta(dv_1/dy)$ for $\pi^{\pm}$, $K^{\pm}$, and $p(\bar{p})$ in Au+Au and isobar collisions at $\sqrt{s_{NN}} = 200$ GeV, as well as Au+Au at 27 GeV. As shown in Figure 3.3, a pivotal finding was that the $p(\bar{p})$ v_1 splitting exhibits a distinct sign change, transitioning from positive to negative values as the collisions become more peripheral; the $K^{\pm}$ splitting follows the same trend, though more weakly and most clearly at $\sqrt{s_{NN}} = 27$ GeV. Since baryon transport contributes a persistent positive background across all centralities, the observation of a negative total splitting in peripheral events provides a crucial benchmark. It indicates that the electromagnetic field effects, specifically Faraday induction and Coulomb forces, become sufficiently dominant in peripheral collisions to overcome the transport-driven background and flip the sign of the observed splitting.

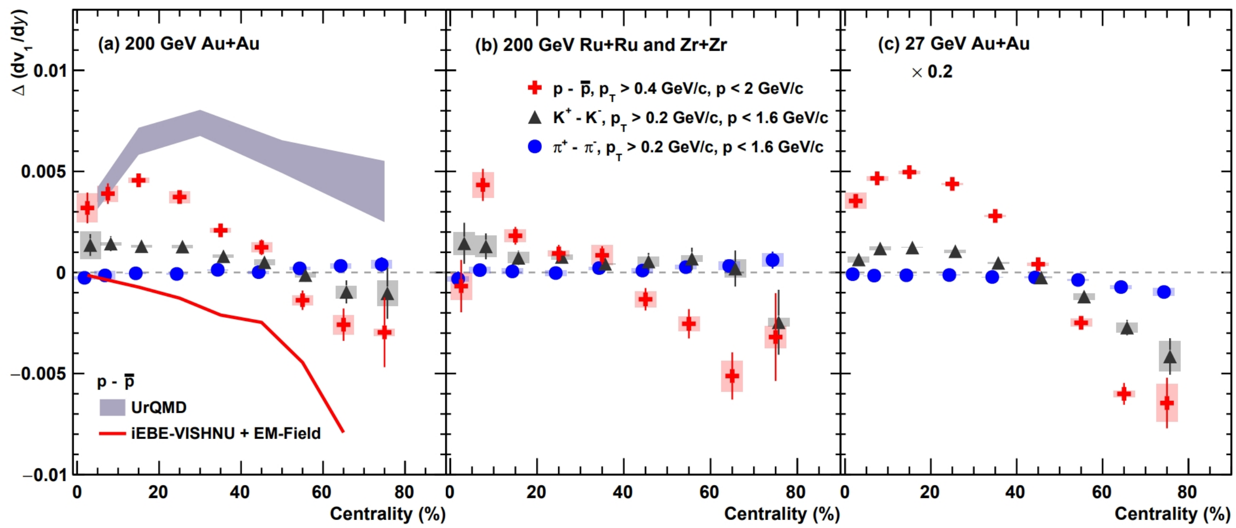

Figure 3.3: $\Delta(dv_1/dy)$ between positively and negatively charged pions, kaons, and protons as a function of centrality in (a) Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV, (b) isobar collisions at $\sqrt{s_{NN}} = 200$ GeV, and (c) Au+Au collisions at $\sqrt{s_{NN}} = 27$ GeV. The sign change from positive to negative values with increasing centrality percentile, clearest for $p(\bar{p})$, indicates that Coulomb and Faraday effects dominate over the baryon transport-driven background in peripheral events [1].

The neutral Λ^0 hyperon (uds) occupies a unique position in this landscape. Because it carries no net electric charge, the Λ^0 itself is immune to any direct electromagnetic deflection at the hadronic level. However, its three constituent quarks individually carry charges ($q_u = +\frac{2}{3}e$, $q_d = q_s = -\frac{1}{3}e$) and are therefore each subject to the Coulomb, Faraday, and Hall forces described above during the pre-hadronization phase. If the constituent quarks act as independent degrees of freedom in the electromagnetic field *before* hadronization, the coalescence sum rule yields an exact prediction, following from the quark content identity $(uud) - (\bar{u}\bar{u}\bar{d}) - [(u\bar{s}) - (\bar{u}s)] \equiv (uds) - (\bar{u}\bar{d}\bar{s})$:

$$\Delta\left(\frac{dv_1}{dy}\right)_\Lambda = \Delta\left(\frac{dv_1}{dy}\right)_p - \Delta\left(\frac{dv_1}{dy}\right)_K, \tag{3.13}$$

where $\Delta X \equiv X_{\text{particle}} - X_{\text{antiparticle}}$ (i.e. $\Delta v_1(K) = v_1(K^+) - v_1(K^-)$). A measurement consistent with this prediction would constitute direct evidence that the electromagnetic field influences hadron formation at the partonic level, before color confinement assembles the final-state hadron.

In the following sections we describe the analysis procedure applied to BES-II Au+Au collision data at $\sqrt{s_{NN}} = 7.7$–$27\,\text{GeV}$ to measure $\Delta(dv_1/dy)$ for Λ^0 and $\bar{\Lambda}^0$ as a function of collision centrality, and compare the results to the coalescence prediction of Eq. 3.13.

3.4.2 Λ Reconstruction

Λ hyperons are reconstructed via the weak decay channel $\Lambda^0 \to p + \pi^-$ (and $\bar{\Lambda}^0 \to \bar{p} + \pi^+$), which has a branching fraction of 63.9% [44]. Candidate daughter tracks are first selected from the reconstructed global tracks and identified as protons or pions; oppositely charged $p\pi$ pairs are then combined into Λ candidates using the KFParticle package [30, 34], a Kalman-filter-based topology fitter that constrains the daughter tracks to a common secondary vertex while propagating the full covariance matrix of each track.

Daughter track quality. Daughter candidates are taken from the global track collection and must satisfy the quality criteria in Table 3.5. The requirement on the number of fit hits ensures a well-measured trajectory, while the bound on the dE/dx resolution,

$\sigma_{dE/dx} \in [0.04, 0.12]$, removes tracks with poorly determined energy loss that would degrade the particle identification. Unlike the primary tracks, no DCA cut to the primary vertex is imposed on the daughters; the displacement is instead required through the χ^2 of each daughter with respect to the primary vertex (Table 3.6), which separates the secondary decay daughters from prompt primary tracks.

Table 3.5: Track quality criteria applied to Λ daughter candidates (p, π).

Cut	Value
Charge	$\neq 0$
Hits used in track fit, N_{fit}	≥ 15
dE/dx resolution, $\sigma_{dE/dx}$	[0.04, 0.12]

Daughter particle identification. Each daughter is identified by combining the TPC ionization energy loss (dE/dx) with the TOF mass-squared m^2, when a matched TOF hit is available. For each track the dE/dx pull $n\sigma_i^{dE/dx}$ (the deviation from the expected ionization for mass hypothesis i, in units of the resolution) and, when available, the TOF mass-squared pull $n\sigma_i^{m^2}$ (the deviation from the momentum-dependent expected m^2) are evaluated for the pion and proton hypotheses. The selection proceeds in two stages:

- **TPC** dE/dx**:** a hypothesis is retained if its pull satisfies the hard cut $|n\sigma^{dE/dx}| < 3$ for pions and protons. Because the dE/dx bands overlap at low momentum, this stage is non-exclusive: a track may be compatible with more than one hypothesis.

- **TOF** m^2**:** the m^2 pulls of the competing hypotheses are compared, and only the single best-matching species (smallest $|n\sigma^{m^2}|$) is retained, and only if it satisfies the hard cut $|n\sigma^{m^2}| < 3$. This stage is therefore exclusive.

The final identity is the intersection of the two stages: when a TOF hit is present, a daughter must pass the dE/dx cut *and* be the best TOF hypothesis with $|n\sigma^{m^2}| < 3$; in the absence of a TOF hit, the dE/dx selection alone is used. This momentum-dependent $n\sigma_{m^2}$ requirement replaces the fixed m^2 window employed in the pion v_2 analysis (Sect. 3.3.3).

The pairs surviving daughter selection are combined into Λ candidates with the topological cuts listed in Table 3.6.

Table 3.6: KFParticle configuration for Λ reconstruction.

Parameter	Value
CleanLowPVTrackEvents	True
SetSoftKaonPIDMode	True
SetSoftTofPidMode	True
SetChiPrimaryCut2D (χ^2 of each daughter to the primary vertex)	> 3
SetChi2Cut2D (χ^2/NDF of the Λ vertex fit)	< 10
SetLCut (minimum decay length)	> 1.0 cm

The Λ ($\bar{\Lambda}$) signal is identified from the invariant mass distribution of $p\pi^-$ ($\bar{p}\pi^+$) pairs around the PDG mass $M_\Lambda = 1.11568$ GeV/c^2, retained over $M_{\mathrm{inv}} \in [1.0, 1.2]$ GeV/c^2. No mass window selection is applied; the signal and background contributions are separated from the shape of the invariant mass distribution itself. Λ candidates are accepted in the kinematic range $0.4 < p_T < 1.8$ GeV/c and $|y| < 0.6$.

3.4.3 Event Plane Determination

The first-order event plane Ψ_1 is reconstructed using the EPD. To maximize sensitivity to the early-time electromagnetic fields generated by spectators, only tiles in the *spectator* region ($|\eta| > 3.2$) contribute to the event plane calculation. An odd η-weight,

$$w(\eta) = c_1\,\eta + c_3\,\eta^3, \tag{3.14}$$

is applied, with nine centrality-dependent coefficient sets tuned to optimize the v_1 resolution. A phi-weight flattening correction, derived from the EPD tile geometry, is applied first, followed by an iterative Fourier shift correction. Both corrections are stored and applied on a day-by-day, centrality-by-centrality basis.

The event plane resolution is measured with the east–west sub-event method:

$$\mathrm{Res}(\Psi_{1,\mathrm{EPD}}) = \sqrt{\big\langle \cos\big(\Psi_{1,\mathrm{east}} - \Psi_{1,\mathrm{west}}\big)\big\rangle}. \tag{3.15}$$

The first-order event-plane resolution peaks in mid-central collisions and decreases toward central and peripheral events; the mid-central peak falls from approximately 0.65 at $\sqrt{s_{NN}} =$ 7.7 GeV to 0.32 at 27 GeV, with values across all energies and centralities spanning roughly 0.1 to 0.65.

The directed flow of each Λ ($\bar{\Lambda}$) candidate is then:

$$v_1 = \frac{\big\langle \cos(\phi - \Psi_1)\big\rangle}{\mathrm{Res}(\Psi_1)}, \tag{3.16}$$

where the average is taken over all candidates across all events in a given centrality and rapidity bin.

3.4.4 v_1 Signal Extraction

The v_1 of Λ ($\bar{\Lambda}$) as a function of rapidity and centrality is extracted using the *v_1-vs-invariant-mass* method, which avoids the need to separate signal and background events event-by-event. The raw observable stored per (y, p_T) bin is the profile $\langle \cos(\phi_\Lambda - \Psi_1)\rangle$ as a function of M_{inv}, together with the invariant mass spectrum dN/dM_{inv}. These two distributions are fitted simultaneously in two sequential steps.

Step 1: Invariant Mass Fit

The invariant mass spectrum of $p\pi^-$ ($\bar{p}\pi^+$) candidates is described by

$$\frac{dN}{dM_{\text{inv}}} = S(M_{\text{inv}}) + B(M_{\text{inv}}), \tag{3.17}$$

where the signal S is a double truncated-Gaussian with a shared peak position μ,

$$S(M_{\text{inv}}) = s\Big[z\,\mathcal{N}_{\text{tr}}(M_{\text{inv}}\,|\,\mu,\sigma_1) + (1-z)\,\mathcal{N}_{\text{tr}}(M_{\text{inv}}\,|\,\mu,\sigma_2)\Big], \tag{3.18}$$

and the combinatorial background B is a second-order polynomial normalized over the fit range,

$$B(M_{\text{inv}}) = b\,\frac{p_0 + p_1 M_{\text{inv}} + p_2 M_{\text{inv}}^2}{\int_{M_{\text{lb}}}^{M_{\text{rb}}}(p_0 + p_1 m + p_2 m^2)\,dm}. \tag{3.19}$$

Here s and b are the total signal and background counts; $z \in (0,1)$ is the fraction carried by the first Gaussian; σ_1 and σ_2 are the two widths; $\mathcal{N}_{\text{tr}}$ denotes a Gaussian truncated to the fit range $[M_{\text{lb}},\, M_{\text{rb}}] = [M_\Lambda - 6\sigma_0,\, M_\Lambda + 10\sigma_0]$, with $\sigma_0 = 2.3$ MeV/c^2. If the post-fit width ratio $|\sigma_1/\sigma_2| > 5$ (or its inverse), the fit falls back to a single Gaussian ($z = 1$) to avoid a degenerate solution.

The fit is performed using an extended binned negative log-likelihood cost function (implemented with `iminuit`/`Minuit2`) in three steps:

1. Mask the signal region $|M_{\text{inv}} - M_\Lambda| < 3 \times 3\,\text{MeV}/c^2$, fix $s = 0$, and fit the background polynomial to the sidebands alone.

2. Fix the background parameters (b, p_0, p_1, p_2), release s, and fit the signal amplitude.

3. Release all parameters and perform a global unconstrained fit.

The resulting $S(M_{\text{inv}})$ and $B(M_{\text{inv}})$ are used to compute the signal fraction $f_S(M_{\text{inv}}) \equiv S/(S+B)$, which is passed directly to the v_1 fit.

Step 2: v_1 Profile Fit

The measured profile $\langle\cos(\phi_\Lambda - \Psi_1)\rangle(M_{\text{inv}})$ contains contributions from both signal and background:

$$v_{1,\text{total}}(M_{\text{inv}}) = f_S(M_{\text{inv}}) \cdot v_{1,\text{sig}} + \big[1 - f_S(M_{\text{inv}})\big] \cdot v_{1,\text{bkg}}(M_{\text{inv}}), \tag{3.20}$$

where $v_{1,\text{sig}}$ is a single free parameter (the signal v_1 for the given centrality and rapidity bin), and the background v_1 is parameterized as a first-order polynomial centered on the Λ PDG mass,

$$v_{1,\text{bkg}}(M_{\text{inv}}) = a_0 + a_1\,(M_{\text{inv}} - M_\Lambda). \tag{3.21}$$

The profile is fitted with a least-squares cost function using a robust `soft-`ℓ_1 loss to reduce sensitivity to statistical outliers, following the same three-step masking procedure as the mass fit:

1. Mask the signal region $|M_{\text{inv}} - \mu| < 3\sigma_{\text{eff}}$, fix $v_{1,\text{sig}} = 0$, and fit (a_0, a_1) to the sideband profile. Only bins with at least one entry are included.

2. Fix (a_0, a_1), release $v_{1,\text{sig}}$, and fit the signal.

3. Release all parameters and perform a global fit.

Final parameter values and symmetric uncertainties are obtained from the HESSE algorithm; asymmetric uncertainties from the profile-likelihood (MINOS) method are computed when the fit converges to a valid and accurate minimum.

Figure 3.4 shows a representative example of the invariant mass spectrum and $v_1(M_{\mathrm{inv}})$ profile fit for centrality 40–50%, rapidity $0.7 < y < 0.8$, at $\sqrt{s_{NN}} = 19.6$ GeV.

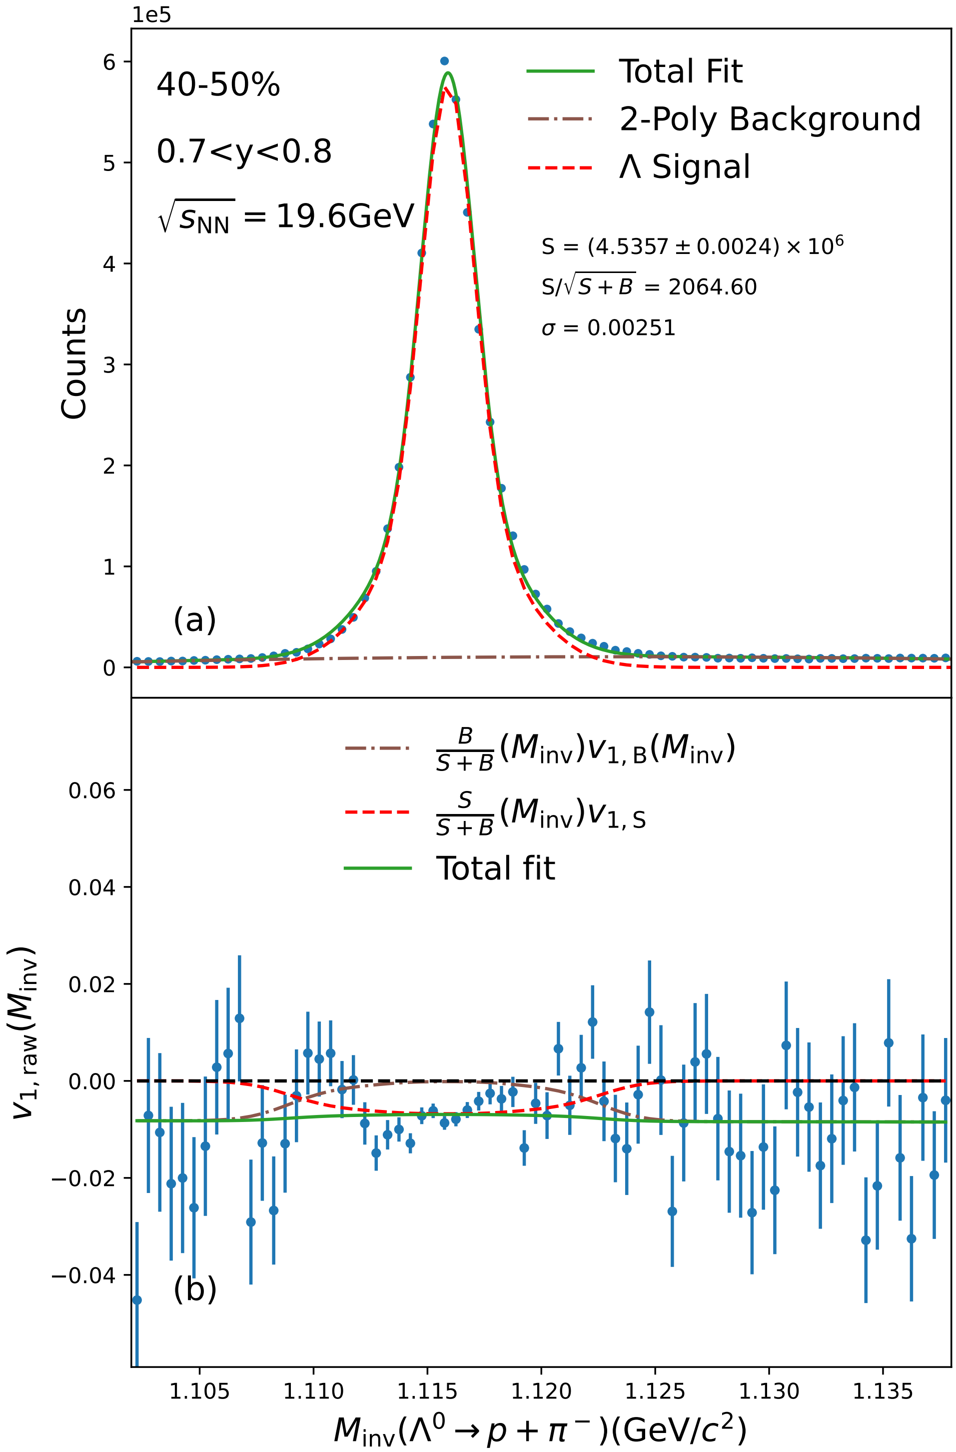

Figure 3.4: (a) Invariant mass spectrum of Λ^0 (points) with the double-Gaussian signal (red dashed), the polynomial background (brown dash-dotted), and their sum (green solid), for centrality 40–50%, $0.7 < y < 0.8$ at $\sqrt{s_{NN}} = 19.6$ GeV. (b) Corresponding $v_1(M_{\mathrm{inv}})$ profile with the signal and background terms of Eq. 3.20 shown separately in the same line styles, together with their sum.

3.4.5 Systematic Uncertainties

The systematic uncertainties on the Λ^0 ($\bar{\Lambda}^0$) directed flow are estimated from the variations listed in Table 3.7, each assessed with the significance-gated Barlow procedure of Sect. 3.3.5. The cut-based sources (vertex z, n_{HitsFit}, the KFParticle topology, and the efficiency-correction scheme) are treated as half-width uniform variations ($/\sqrt{3}$); the $v_1(y)$ fit-range source, evaluated as separate one-sided fits over the positive-y and negative-y halves, is a full-width variation and is divided by $\sqrt{12}$.

Table 3.7: Variations of cuts for estimating systematic uncertainties and the corresponding default cuts.

#	Source		Default	Variation
1	Vertex z position		$\lvert V_z\rvert < 145(70)$ cm[1]	$\lvert V_z\rvert < 70(35)$ cm
2	nHitsFit		> 15	> 20
3	KFParticle Parameters	Max Distance[2]	< 1	< 0.8
		L Cut[3]	> 1.0	> 3.0
4	$v_1(y)$ fitting range		$\lvert y\rvert < 0.6$	$0 < y < 0.6$ and $-0.6 < y < 0$
5	Efficiency correction		per-y-bin	y-integrated

[1] Values in parentheses apply to $\sqrt{s_{NN}} = 19.6$ and 27 GeV.
[2] SetMaxDistanceBetweenParticleCut(). Set the maximum DCA between two daughters.
[3] SetLCut(). Set the minimum decay length of the reconstructed particles.

Chapter 4

Collective Flow Measurements in BES-II

This chapter presents the results of the two collective flow analyses developed in Chapter 3. Both are interpreted through a single framework, quark coalescence, which relates the measured flow of identified hadrons to that of their constituent quarks and thereby turns each measurement into a probe of a distinct piece of the underlying physics. Sect. 4.1 reports the elliptic flow splitting between π^- and π^+, which under coalescence measures the transported quark ratio $N_d^{\mathrm{tr}}/N_u^{\mathrm{tr}}$ and hence the isospin composition of the colliding nuclei. Sect. 4.2 reports the directed flow splitting between Λ^0 and $\bar{\Lambda}^0$ and its centrality and energy dependence, which through the coalescence sum rule built from proton and kaon flow reveals the response of the constituent quarks to the early-time electromagnetic field. Sect. 4.3 draws the two results together.

4.1 Isospin Asymmetry in Pion Elliptic Flows

The elliptic flow of π^+ and π^- is measured as a function of centrality at each of the seven collision energies. The p_T-integrated v_2 is computed over the constituent quark momentum range $p_T/n_q \in [0.08, 0.6]$ GeV/c (corresponding to $p_T \in [0.16, 1.2]$ GeV/c for pions and $p_T \in [0.24, 1.8]$ GeV/c for protons), using the second-order TPC event plane reconstructed with the sub-event method described in Sect. 3.3.4. At $\sqrt{s_{NN}} \geq 14.6$ GeV, an independent EPD-based event plane is also used and the two results are compared.

4.1.1 Pion v_2 as a Function of Centrality

Figure 4.1 shows v_2 for π^- and π^+ as a function of centrality at all seven collision energies. Several features are common across energies:

- v_2 rises from central to mid-peripheral collisions, reflecting the increasing geometric anisotropy of the initial overlap region. The elliptic flow peaks near 30–40% centrality and decreases again in the most peripheral collisions where both the eccentricity and the system size become small.

- π^- carries systematically larger v_2 than π^+ at all centralities. This splitting is the key observable of this analysis: it reflects the higher availability of transported d quarks relative to u quarks in the neutron-rich Au nucleus.

- The splitting magnitude is small in the most central collisions and remains approximately constant or increases mildly toward peripheral collisions, consistent with the expectation from the neutron-skin effect, though the statistical precision at very peripheral centralities is insufficient to resolve a clear trend.

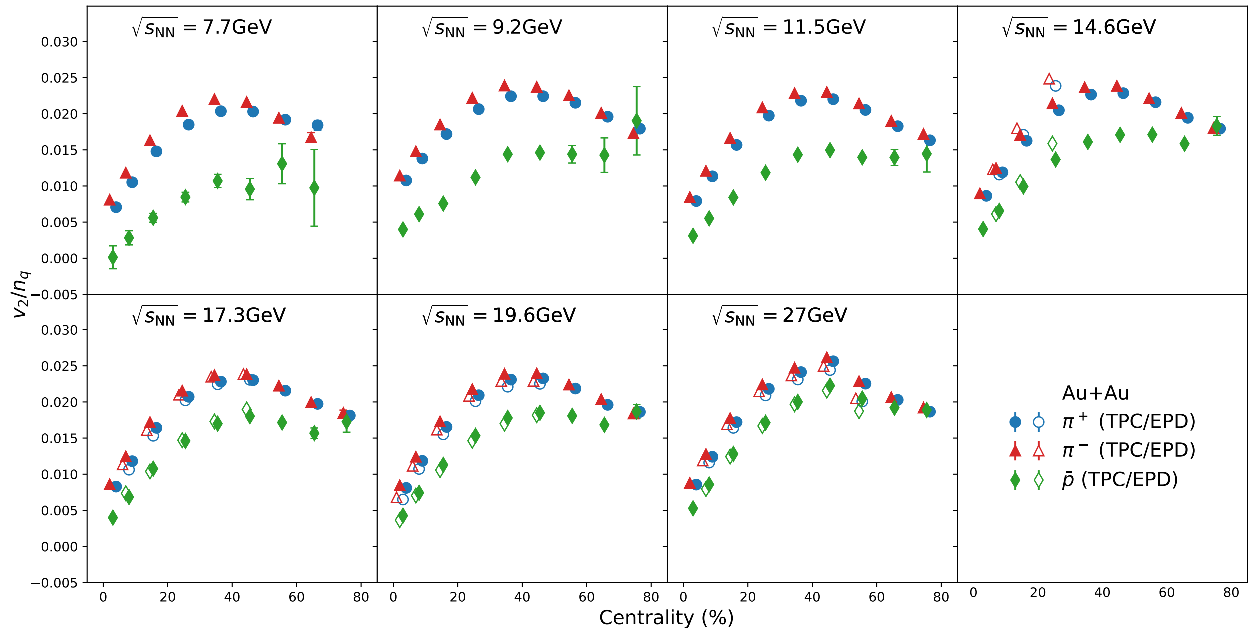

Figure 4.1: p_T-integrated v_2 for π^- (filled markers) and π^+ (open markers) as a function of centrality at $\sqrt{s_{NN}}$ = 7.7, 9.2, 11.5, 14.6, 17.3, 19.6, and 27 GeV. The integration range is $p_T/n_q \in [0.08, 0.6]$ GeV/c. Statistical uncertainties are shown as error bars; systematic uncertainties are shown as boxes.

4.1.2 Extracted Transported Quark Ratio

Figure 4.2 shows the transported quark ratio $N_d^{\rm tr}/N_u^{\rm tr}$, extracted from the measured v_2 values using the coalescence relation of Sect. 3.3:

$$\frac{N_d^{\rm tr}}{N_u^{\rm tr}} = \frac{\dfrac{v_2(\pi^-)}{2} - \dfrac{v_2(\bar{p})}{3}}{\dfrac{v_2(\pi^+)}{2} - \dfrac{v_2(\bar{p})}{3}}.$$

The ratio is consistent with the theoretical prediction from the Au nuclear stoichiometry, $N_d^{\rm tr}/N_u^{\rm tr} = (2N+Z)/(N+2Z) = 315/276 \approx 1.14$, across most centralities and energies. In mid-central and semi-peripheral collisions (10–60%), where the statistical precision is highest, the measured ratio clusters tightly around 1.14 at every energy, with no clear centrality dependence within the current uncertainties. Only in the most central collisions (0–10%)

at 27 GeV does the ratio fall toward unity, possibly reflecting a larger contribution from resonance-decay pions that dilutes the π^-–π^+ splitting. The systematic uncertainty on the 10–40% ratio ranges from 0.002 at 19.6 GeV to 0.038 at 9.2 GeV, dominated by the vertex-z and PID variations.

Glauber calculations incorporating the neutron-skin density profile of ^{197}Au predict that the ratio should rise above 315/276 toward peripheral collisions, where the dilute nuclear surface contributes more strongly. The present data are consistent with such an enhancement but lack the statistics in the most peripheral bins ($> 60\%$) to establish it. At the lowest energies the systems formed are small and survive only briefly, so the assumption that the flow develops in the partonic stage is least justified there; and because the ratio is a quotient of two small, nearly canceling differences, $\frac{1}{2}v_2(\pi^{\mp}) - \frac{1}{3}v_2(\bar{p})$, the statistical fluctuations grow largest just where the elliptic flow is smallest. At 7.7 GeV, the ratio departs from the prediction more strongly, lying above it in mid-central collisions and dropping well below it toward peripheral. Such deviations may reflect a breakdown of coalescence, a statistical fluctuation, or both; the present data cannot distinguish between them.

4.1.3 Energy Dependence

Figure 4.3 shows $N_d^{\mathrm{tr}}/N_u^{\mathrm{tr}}$ integrated over the 10–40% centrality range as a function of collision energy, for both TPC and EPD event planes where available. The measured ratio at each energy is summarized in Table 4.1.

With statistical and systematic uncertainties combined, six of the seven measurements agree with the Au stoichiometric prediction within 2σ; the largest deviation is at $\sqrt{s_{NN}} = 19.6$ GeV ($+2.5\sigma$), where the systematic uncertainty is smallest, and 7.7 GeV lies $+1.8\sigma$ above. No systematic energy trend is observed, consistent with the initial nuclear geometry, and therefore

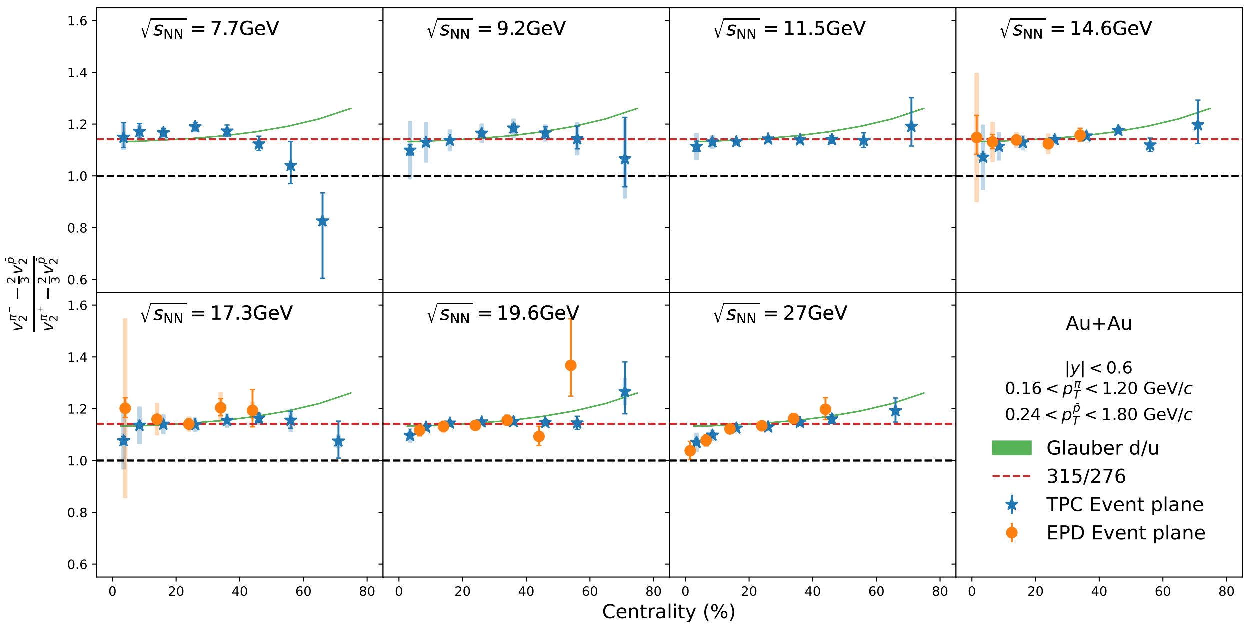

Figure 4.2: Extracted transported quark ratio $N_d^{\rm tr}/N_u^{\rm tr}$ as a function of centrality at $\sqrt{s_{NN}}$ = 7.7, 9.2, 11.5, 14.6, 17.3, 19.6, and 27 GeV. The dashed horizontal line indicates the predicted value $315/276 \approx 1.14$ from Au nuclear stoichiometry. Statistical uncertainties are shown as error bars; systematic uncertainties are shown as boxes.

Table 4.1: Transported quark ratio $N_d^{\rm tr}/N_u^{\rm tr}$ at 10–40% centrality for each collision energy (TPC event plane). Uncertainties are statistical and systematic, respectively. The theoretical prediction is $315/276 \approx 1.141$.

$\sqrt{s_{NN}}$ (GeV)	$N_d^{\rm tr}/N_u^{\rm tr}$
7.7	$1.176 \pm 0.010 \pm 0.017$
9.2	$1.157 \pm 0.005 \pm 0.038$
11.5	$1.138 \pm 0.004 \pm 0.011$
14.6	$1.139 \pm 0.003 \pm 0.021$
17.3	$1.144 \pm 0.004 \pm 0.031$
19.6	$1.150 \pm 0.003 \pm 0.002$
27	$1.134 \pm 0.003 \pm 0.009$

the transported d/u ratio, being independent of $\sqrt{s_{NN}}$.

At energies $\sqrt{s_{NN}} \geq 14.6$ GeV, the TPC and EPD event plane results are in good agreement, providing an important cross-check and ruling out event-plane-related systematic biases as the source of the observed splitting.

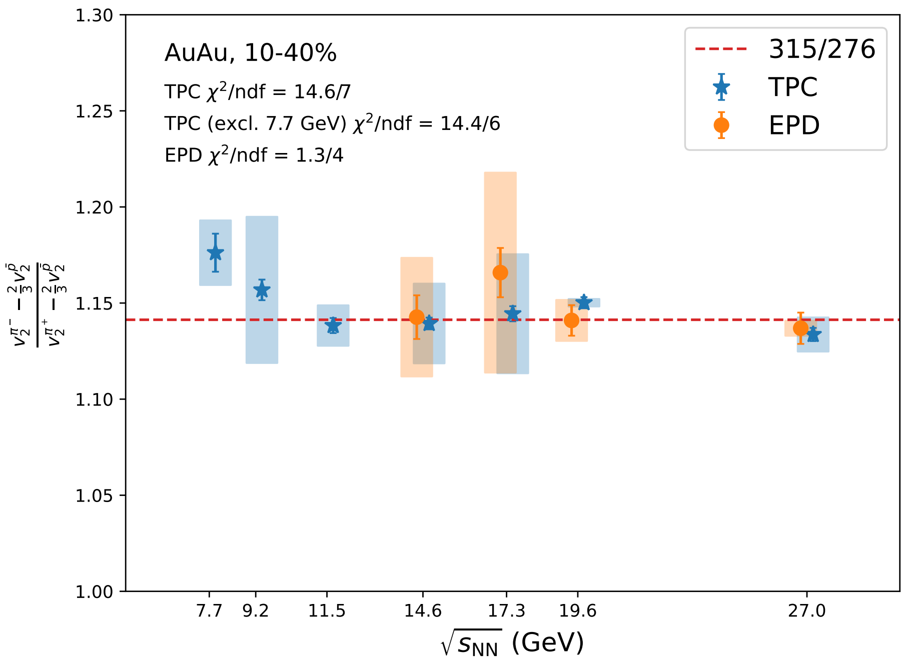

Figure 4.3: Transported quark ratio $N_d^{\text{tr}}/N_u^{\text{tr}}$ at 10–40% centrality as a function of $\sqrt{s_{NN}}$. The dashed horizontal line indicates the Au stoichiometric prediction $315/276 \approx 1.141$. Statistical uncertainties are shown as error bars; systematic uncertainties are shown as boxes.

4.1.4 Robustness Against Resonance Feeddown

A potential concern for the coalescence decomposition of Sect. 3.3 is that a large fraction of the measured pions at BES energies do not originate directly from coalescence but are decay products of short-lived resonances, including the $\rho \to \pi\pi$ channel. At first glance the isospin-symmetric decay $\rho^0 \to \pi^+\pi^-$ would appear to add equally to $v_2(\pi^-)$ and $v_2(\pi^+)$, diluting the splitting and driving the extracted ratio toward unity.

This argument, however, overlooks that the transported quark asymmetry is established *before* hadronization. If the elliptic flow develops during the deconfined partonic stage, the excess of transported d quarks over transported u quarks is already imprinted on the quark momentum distributions before any hadrons form. At hadronization, pions and ρ mesons are produced by coalescence alike, so the ρ mesons inherit the same quark-level asymmetry rather than introducing a contamination from outside the coalescence picture. This reasoning relies on the formation of a deconfined phase; where QGP production is not established, hadronic degrees of freedom persist throughout the evolution and the feeddown would require more careful treatment.

The neutral channel $\rho^0 \to \pi^+\pi^-$ produces one π^+ and one π^- in every decay, so it feeds the two charges equally and can only dilute an existing splitting, moving the measured ratio toward unity. The asymmetry can still be carried by the charged channels: because transported d quarks outnumber transported u quarks, ρ^- ($\bar{u}d$) mesons are more abundant than ρ^+ ($u\bar{d}$), and their decays feed π^- more than π^+.

4.2 Directed Flow Splitting of Λ Hyperons

The directed flow of Λ^0 and $\bar{\Lambda}^0$ is measured differentially in rapidity and centrality at each collision energy. The particle–antiparticle splitting is extracted from these measurements, characterized as a function of centrality and energy, and compared with the coalescence sum rule built from proton and kaon directed flow.

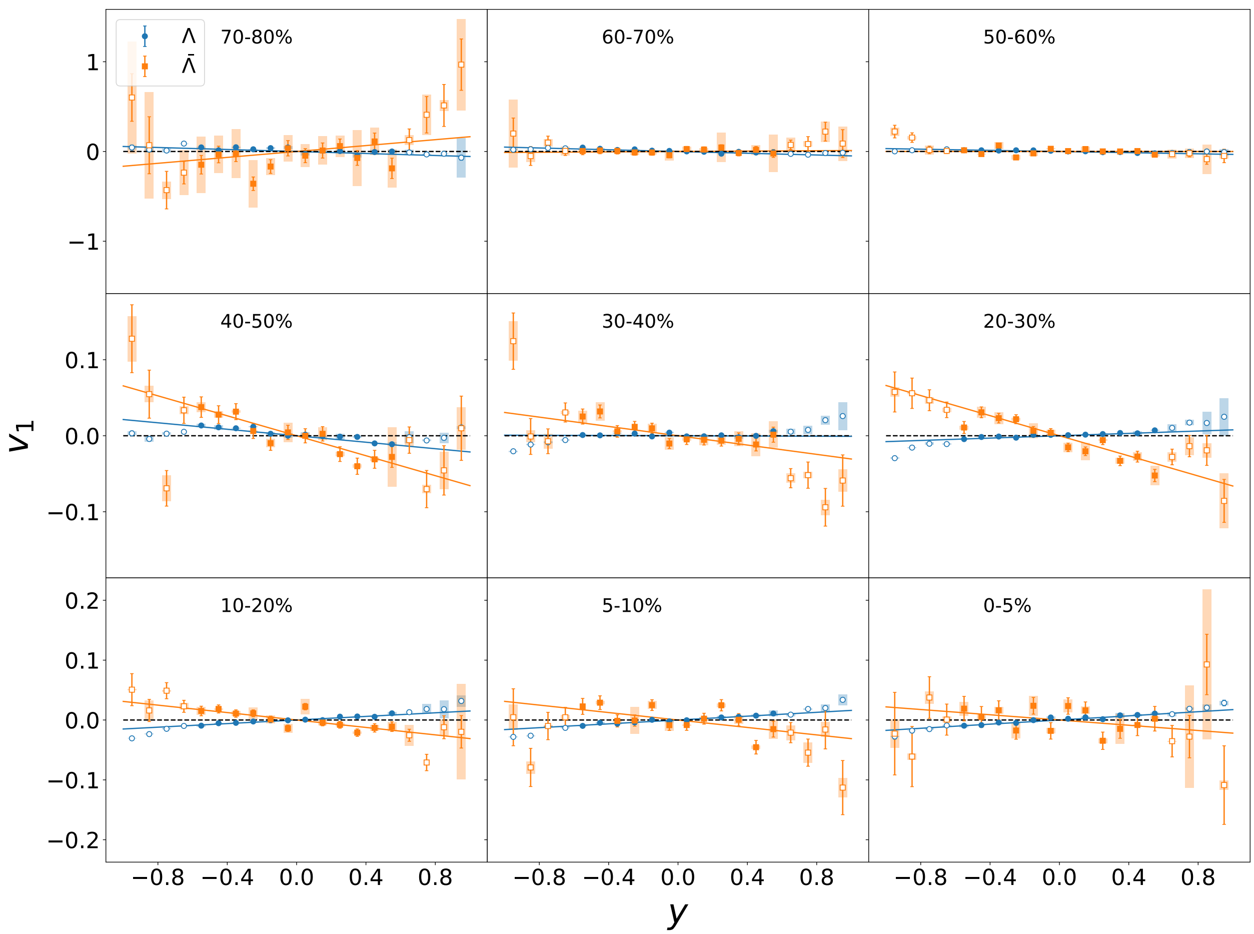

Figure 4.4: $v_1(y)$ for Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 7.7$ GeV. Λ^0 is drawn as circles and $\bar{\Lambda}^0$ as squares. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

4.2.1 $v_1(y)$ of Λ^0 and $\bar{\Lambda}^0$

Figs. 4.4–4.10 present measurements of $v_1(y)$ of Λ^0 and $\bar{\Lambda}^0$ as a function of centrality for each collision energy. Panels are omitted for centrality bins in which a reliable EPD first-order event-plane resolution could not be obtained. Within $|y| < 0.6$, the distributions are well described by a linear function; at larger rapidities, deviations consistent with a cubic structure become visible in some centrality bins, motivating the restriction of the fit range. The Λ^0 generally exhibits a stronger v_1 slope than the $\bar{\Lambda}^0$, reflecting the additional contribution from transported quarks carried over from the initial-state nuclei.

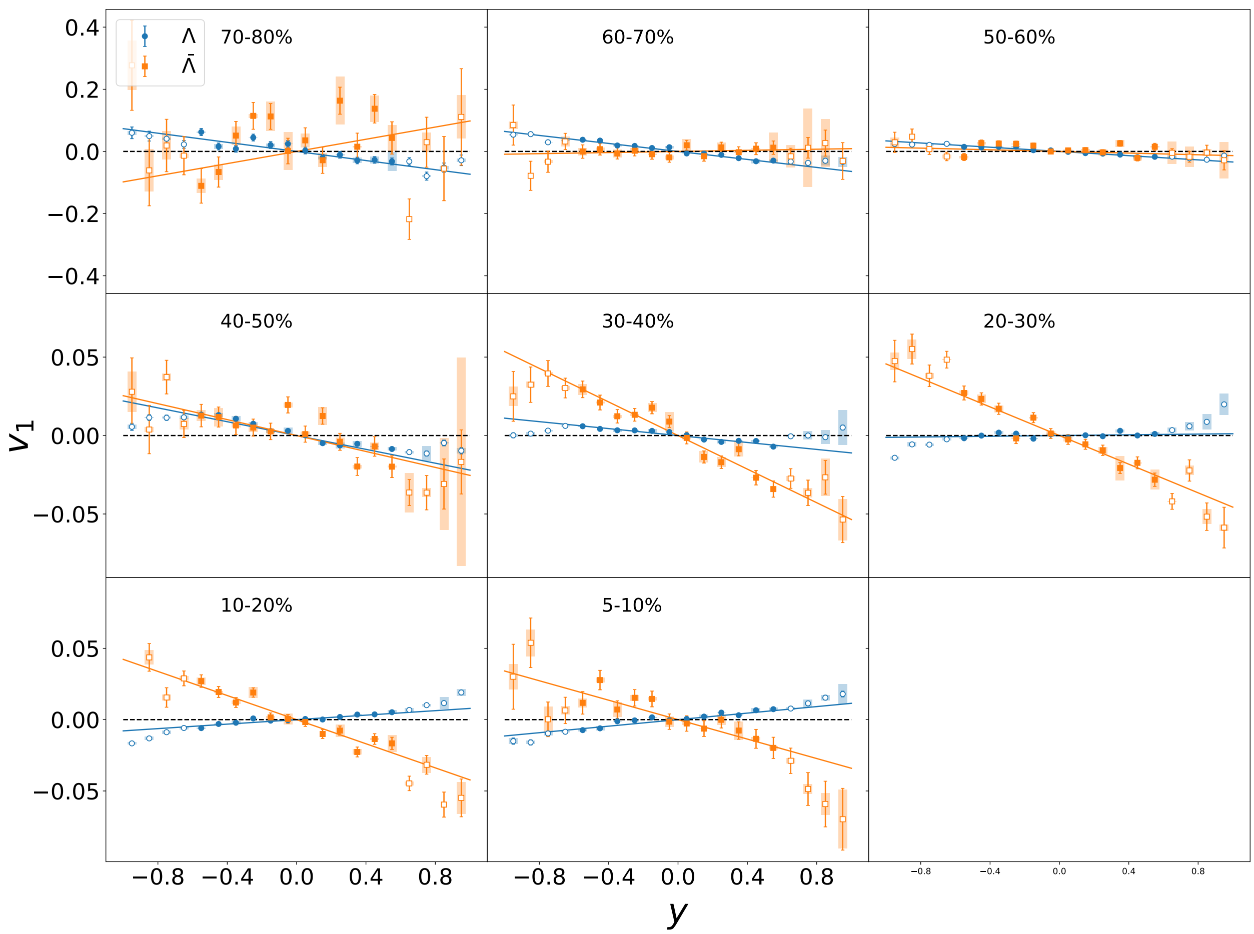

Figure 4.5: $v_1(y)$ for Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 9.2$ GeV. Λ^0 is drawn as circles and $\bar{\Lambda}^0$ as squares. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

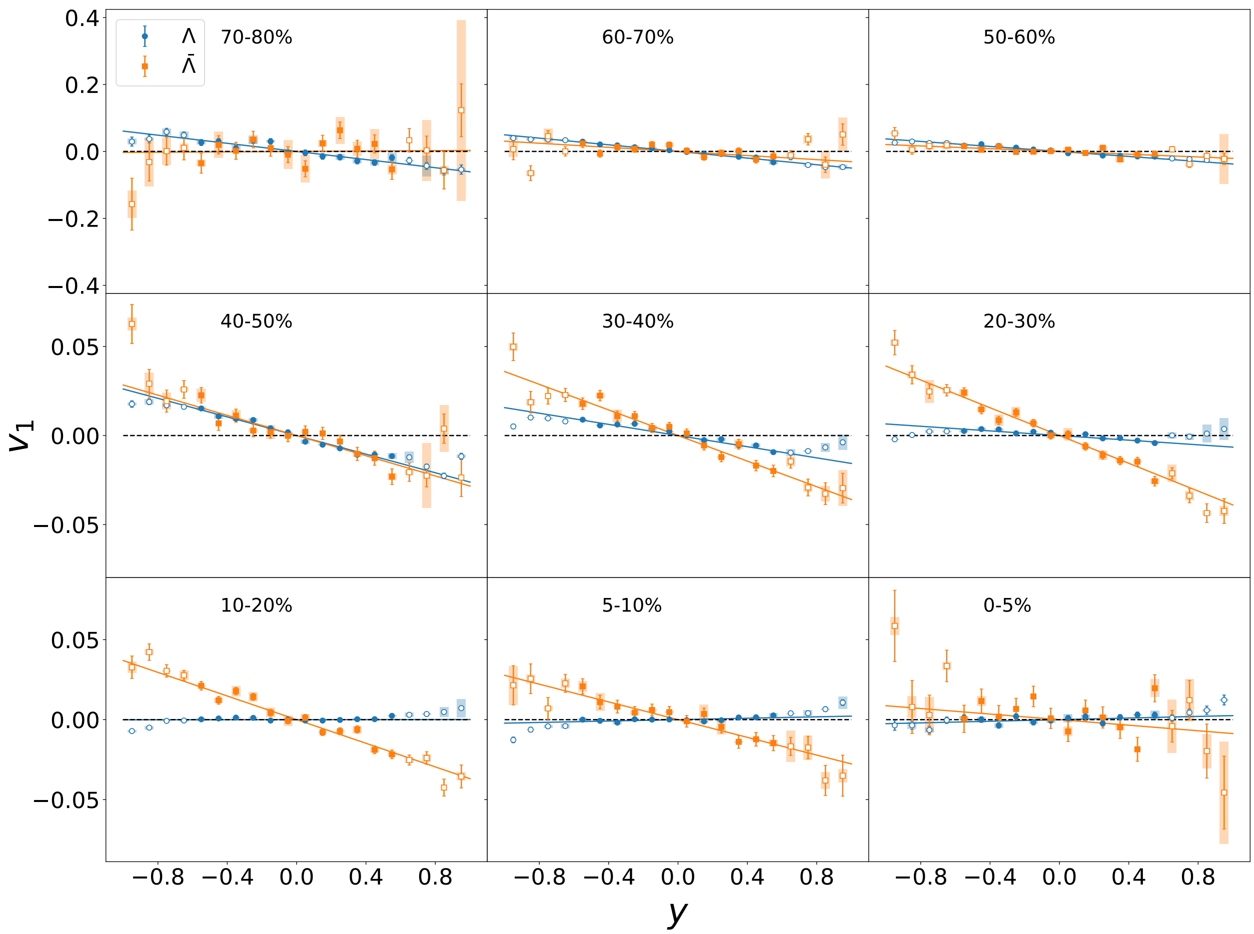

Figure 4.6: $v_1(y)$ for Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 11.5$ GeV. Λ^0 is drawn as circles and $\bar{\Lambda}^0$ as squares. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

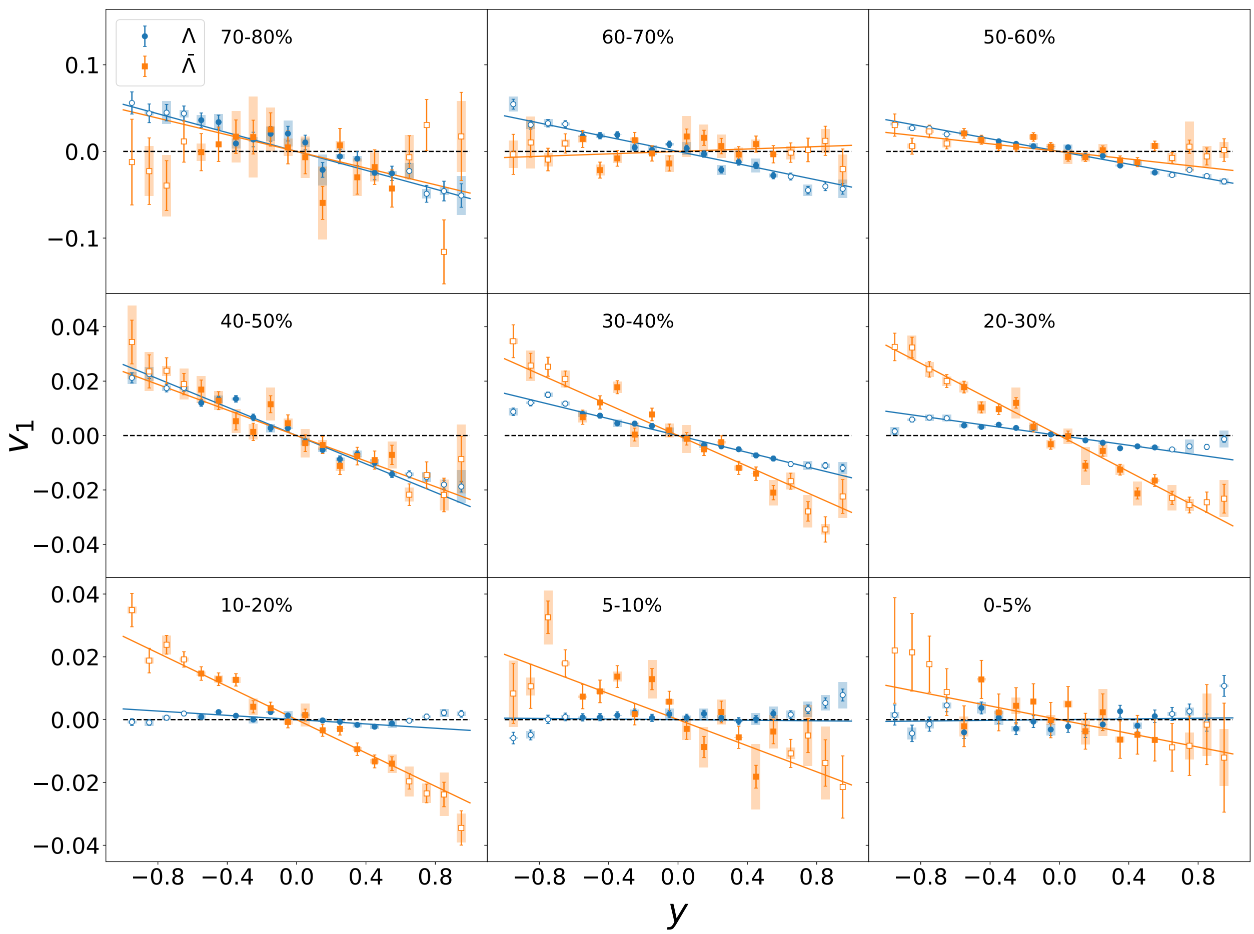

Figure 4.7: $v_1(y)$ for Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 14.6$ GeV. Λ^0 is drawn as circles and $\bar{\Lambda}^0$ as squares. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

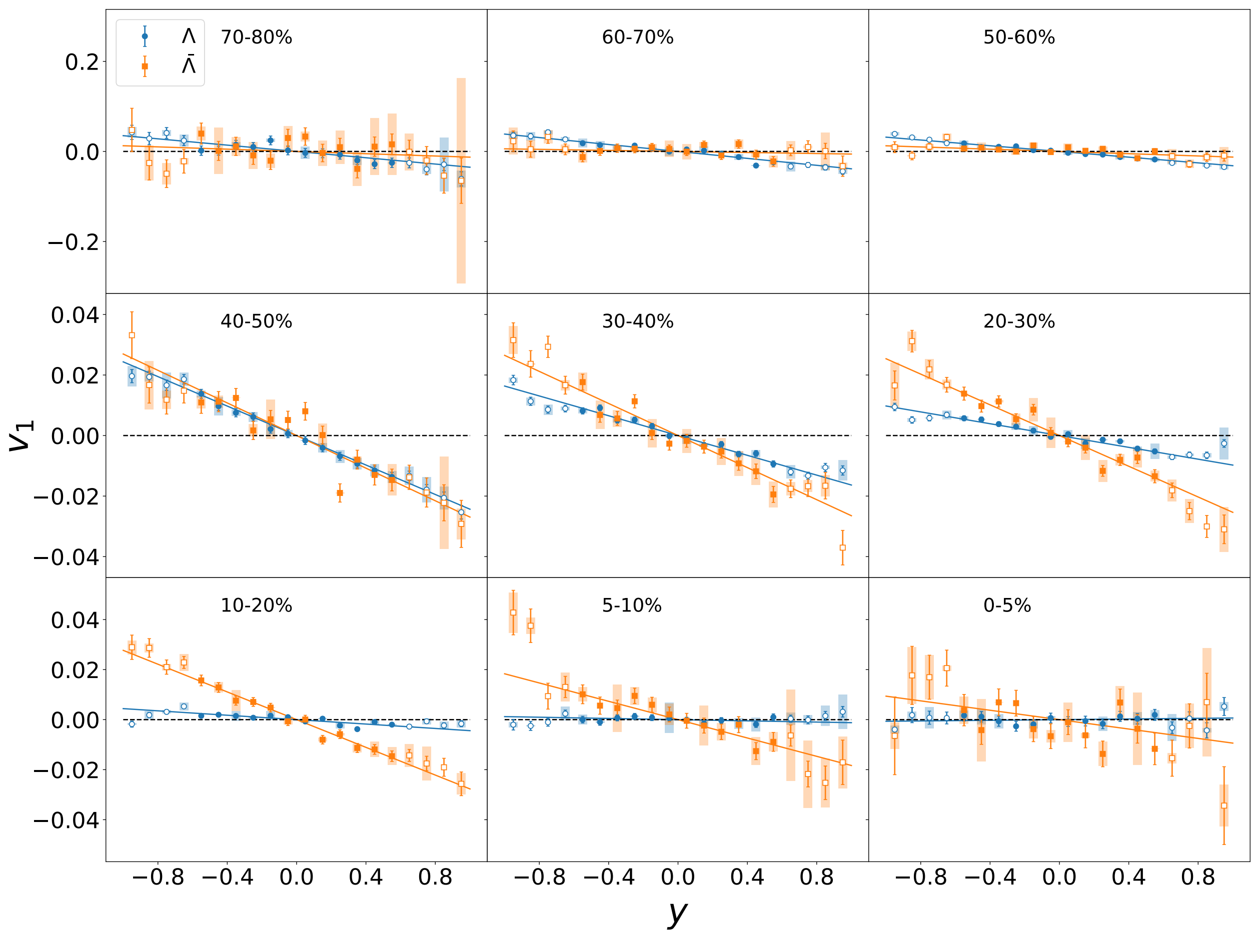

Figure 4.8: $v_1(y)$ for Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 17.3$ GeV. Λ^0 is drawn as circles and $\bar{\Lambda}^0$ as squares. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

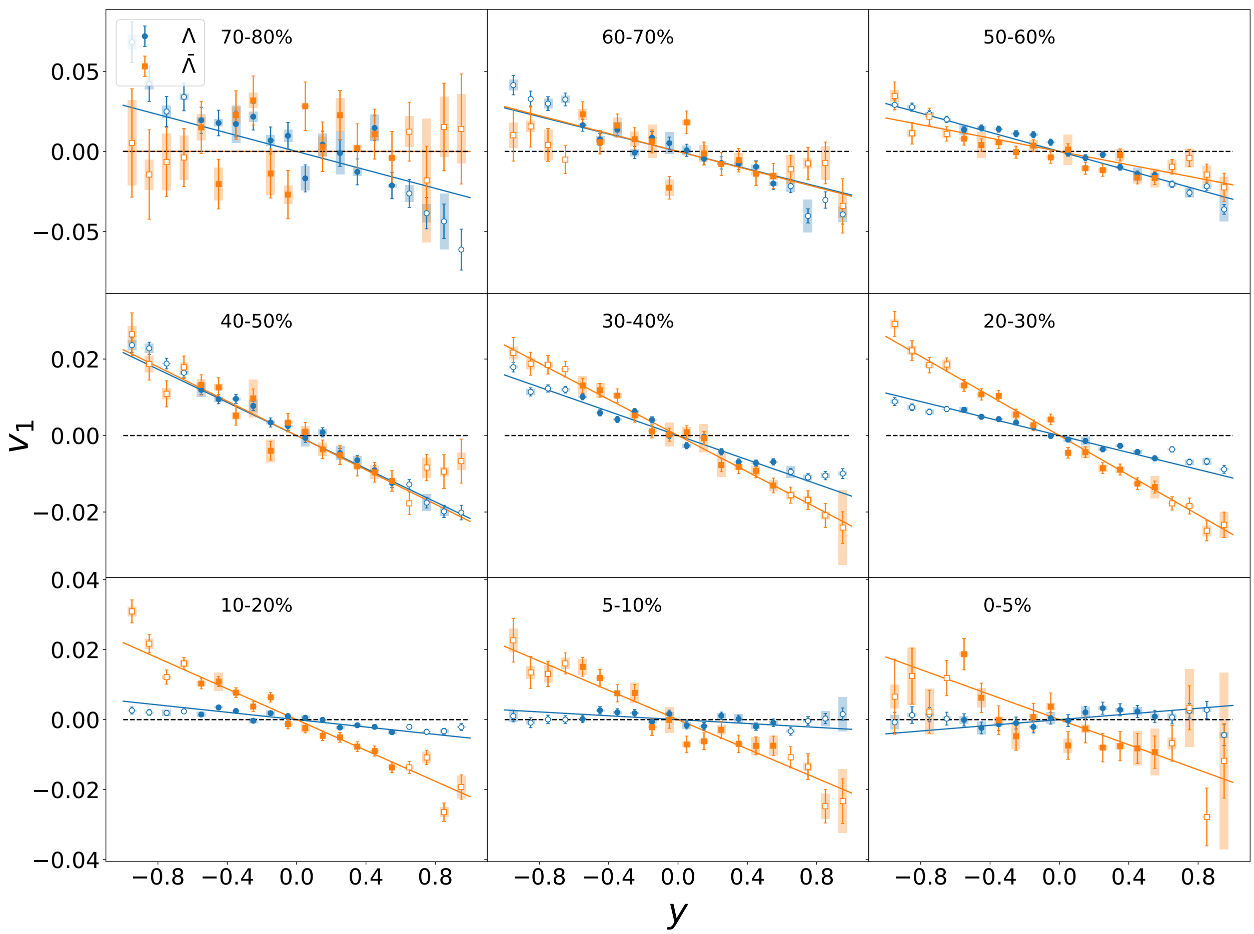

Figure 4.9: $v_1(y)$ for Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 19.6$ GeV. Λ^0 is drawn as circles and $\bar{\Lambda}^0$ as squares. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

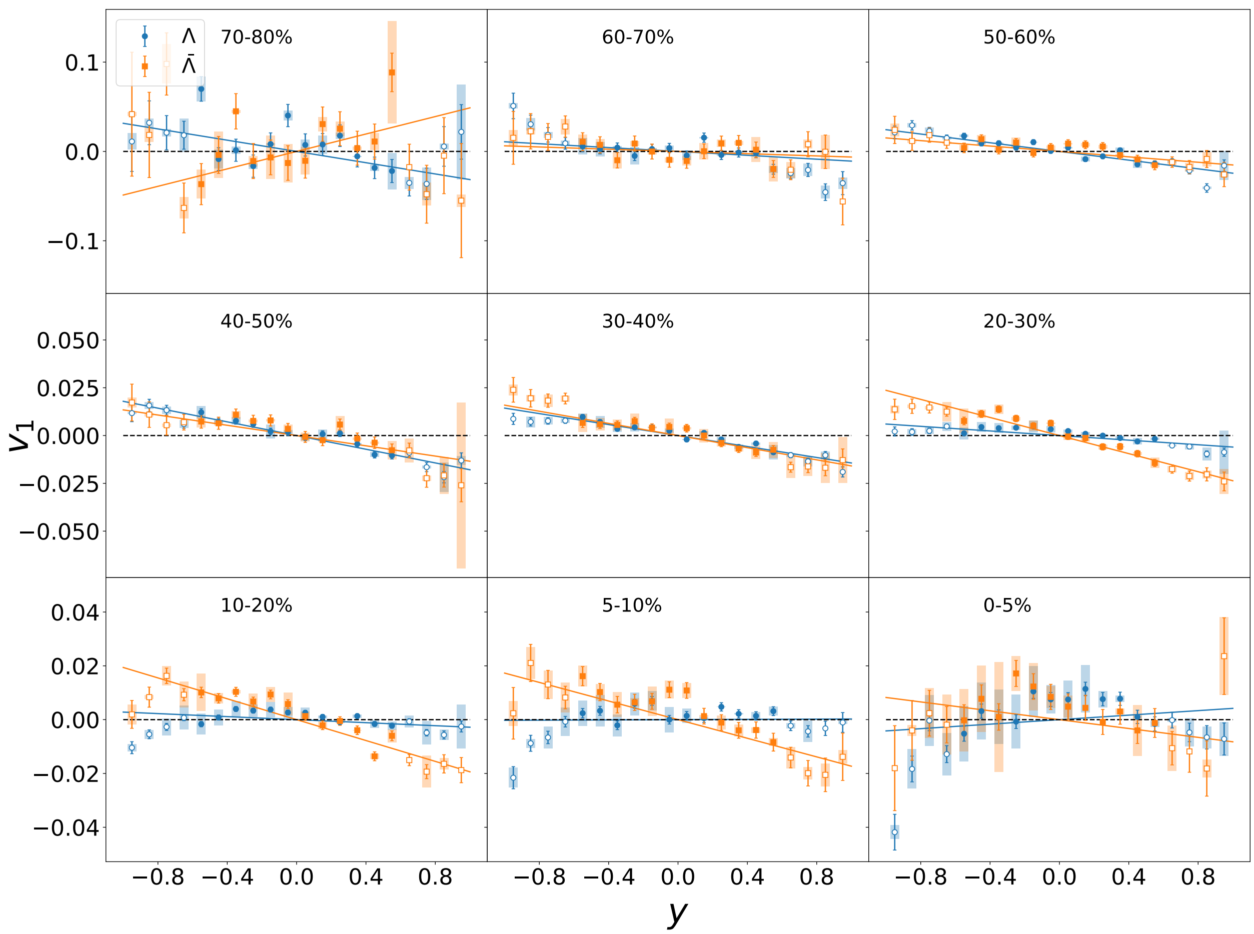

Figure 4.10: $v_1(y)$ for Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 27$ GeV. Λ^0 is drawn as circles and $\bar{\Lambda}^0$ as squares. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

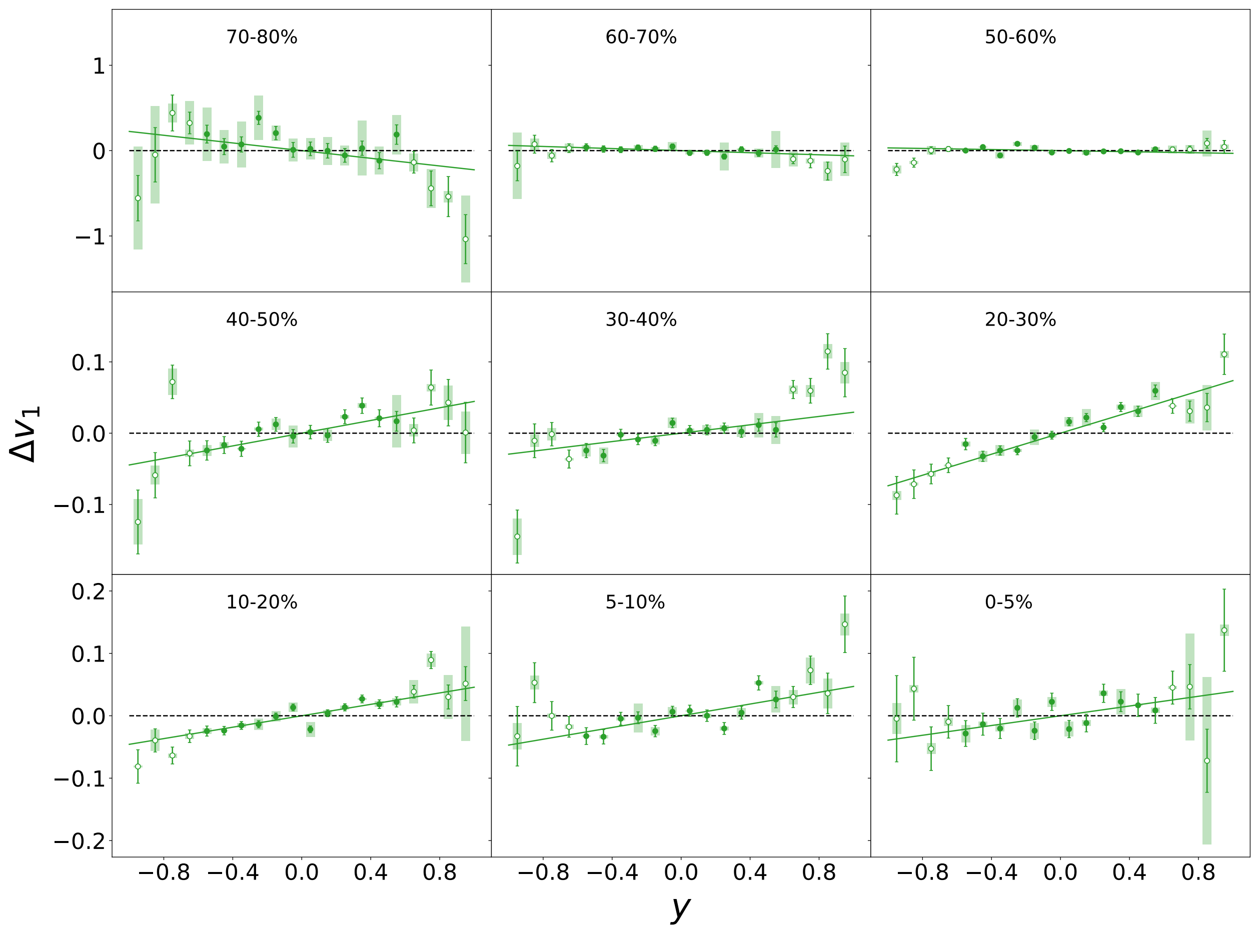

Figure 4.11: $\Delta v_1(y)$ between Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 7.7$ GeV. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

4.2.2 $\Delta v_1(y)$ between Λ^0 and $\bar{\Lambda}^0$

Figs. 4.11–4.17 present the directed flow splitting $\Delta v_1(y) = v_1^{\Lambda^0}(y) - v_1^{\bar{\Lambda}^0}(y)$ as a function of rapidity and centrality. The cubic structures observed in the individual $v_1(y)$ distributions largely cancel in the difference, and the splitting is consistent with a linear rapidity dependence across all centralities and energies. As for the individual $v_1(y)$, the slope of $\Delta v_1(y)$ is extracted from a linear fit over $|y| < 0.6$; the sensitivity of the slope to this range is included in the systematic uncertainty via separate one-sided fits over the positive-y and negative-y halves. The slopes extracted from these linear fits are presented in the following section.

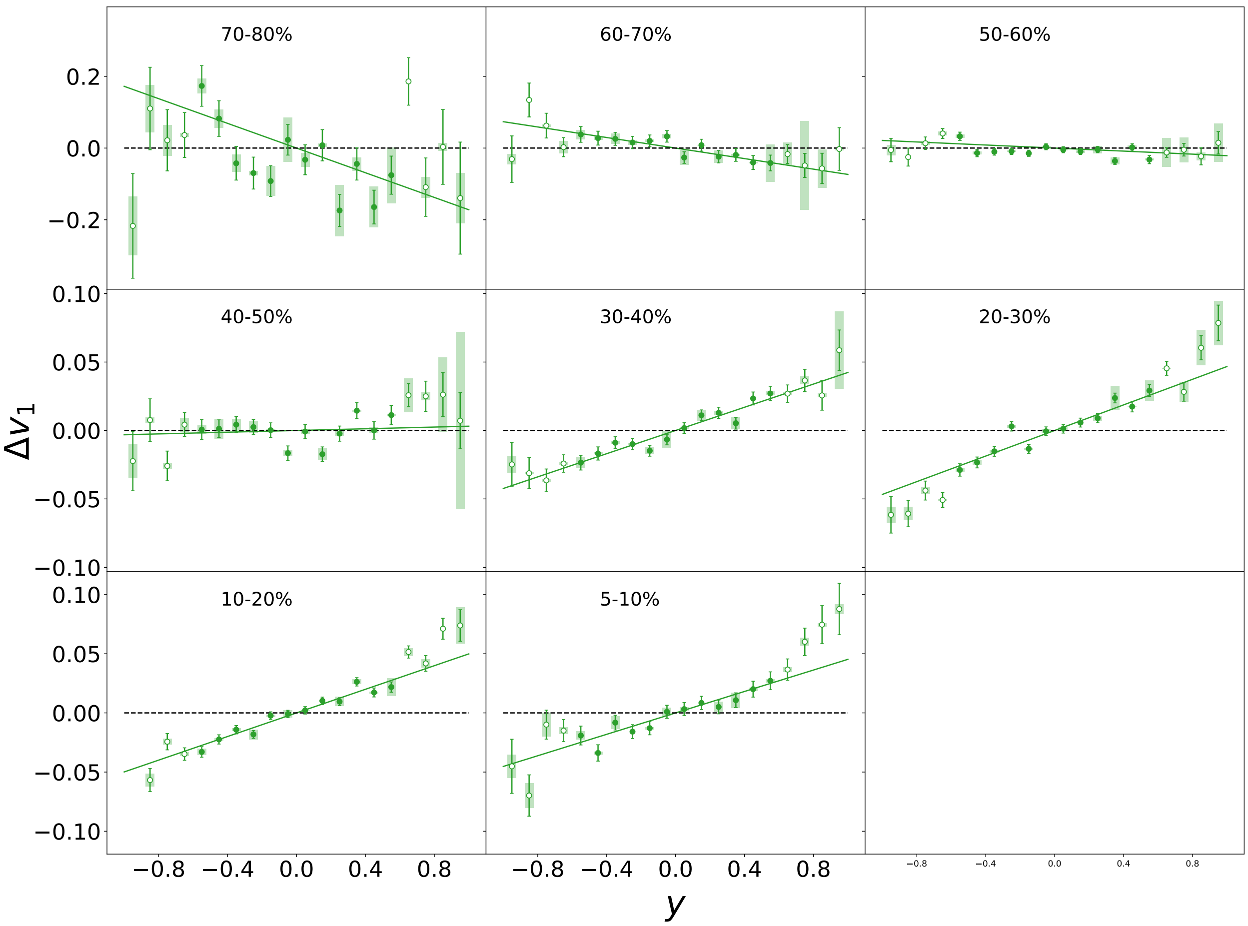

Figure 4.12: $\Delta v_1(y)$ between Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 9.2$ GeV. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

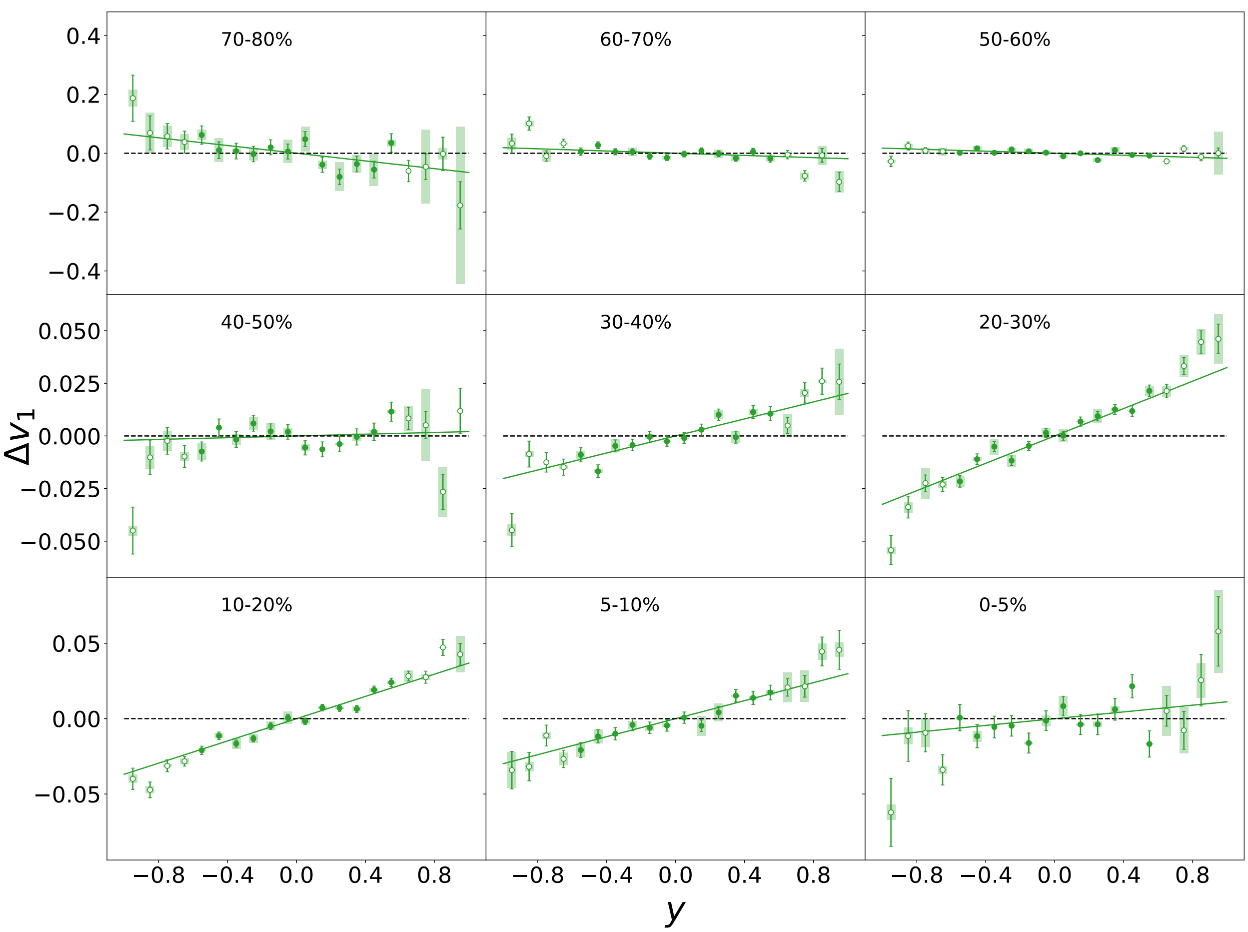

Figure 4.13: $\Delta v_1(y)$ between Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 11.5$ GeV. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

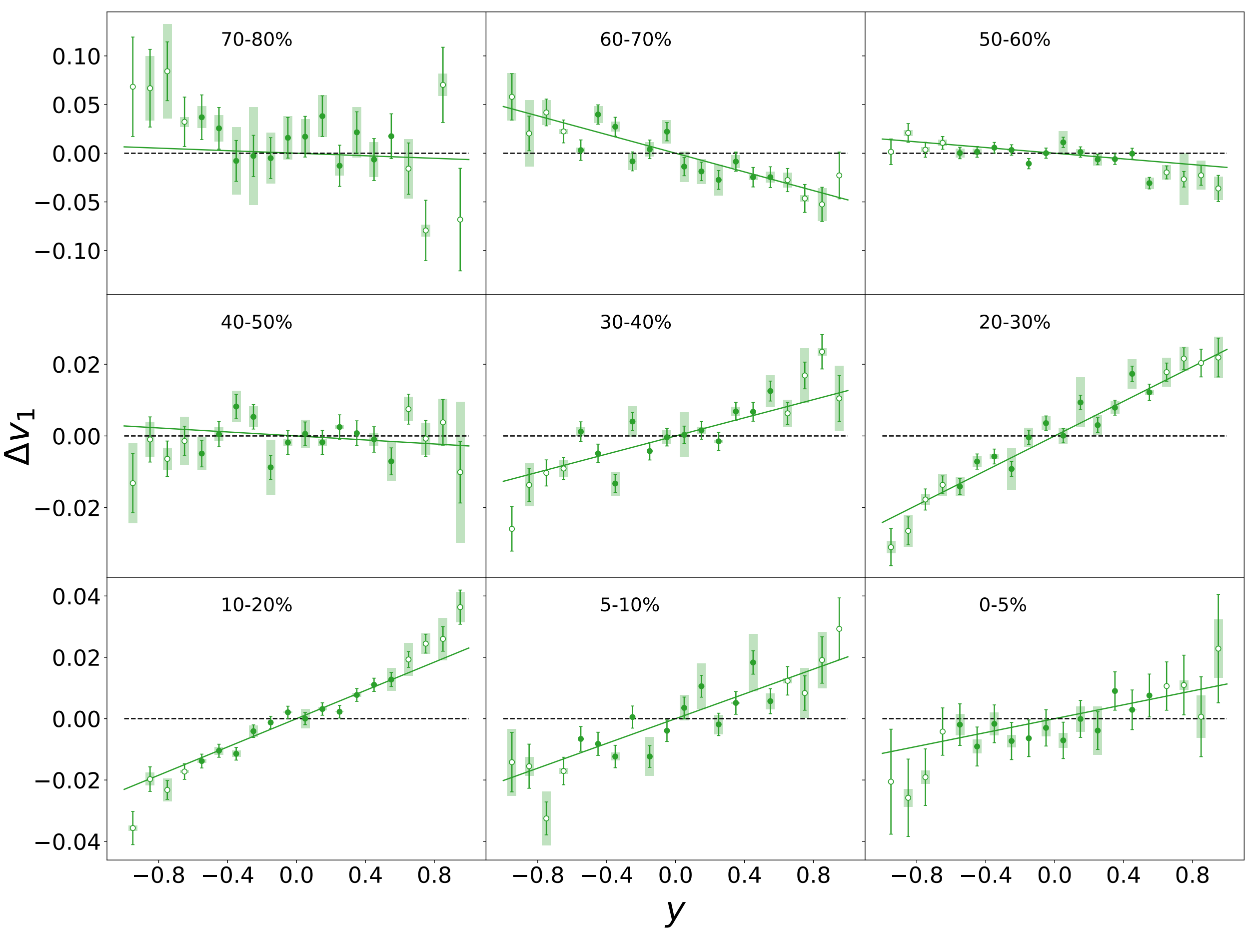

Figure 4.14: $\Delta v_1(y)$ between Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 14.6$ GeV. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

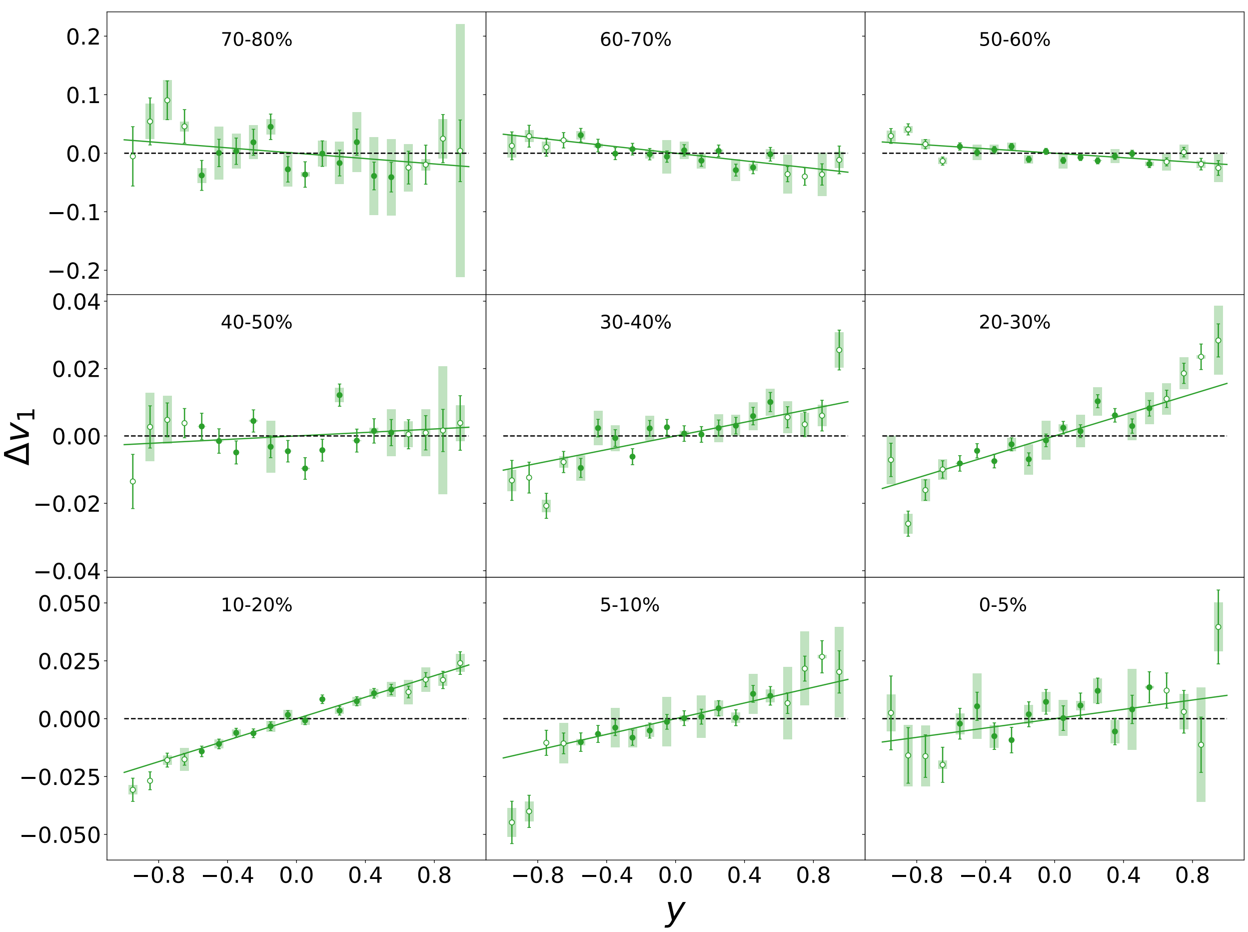

Figure 4.15: $\Delta v_1(y)$ between Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 17.3$ GeV. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

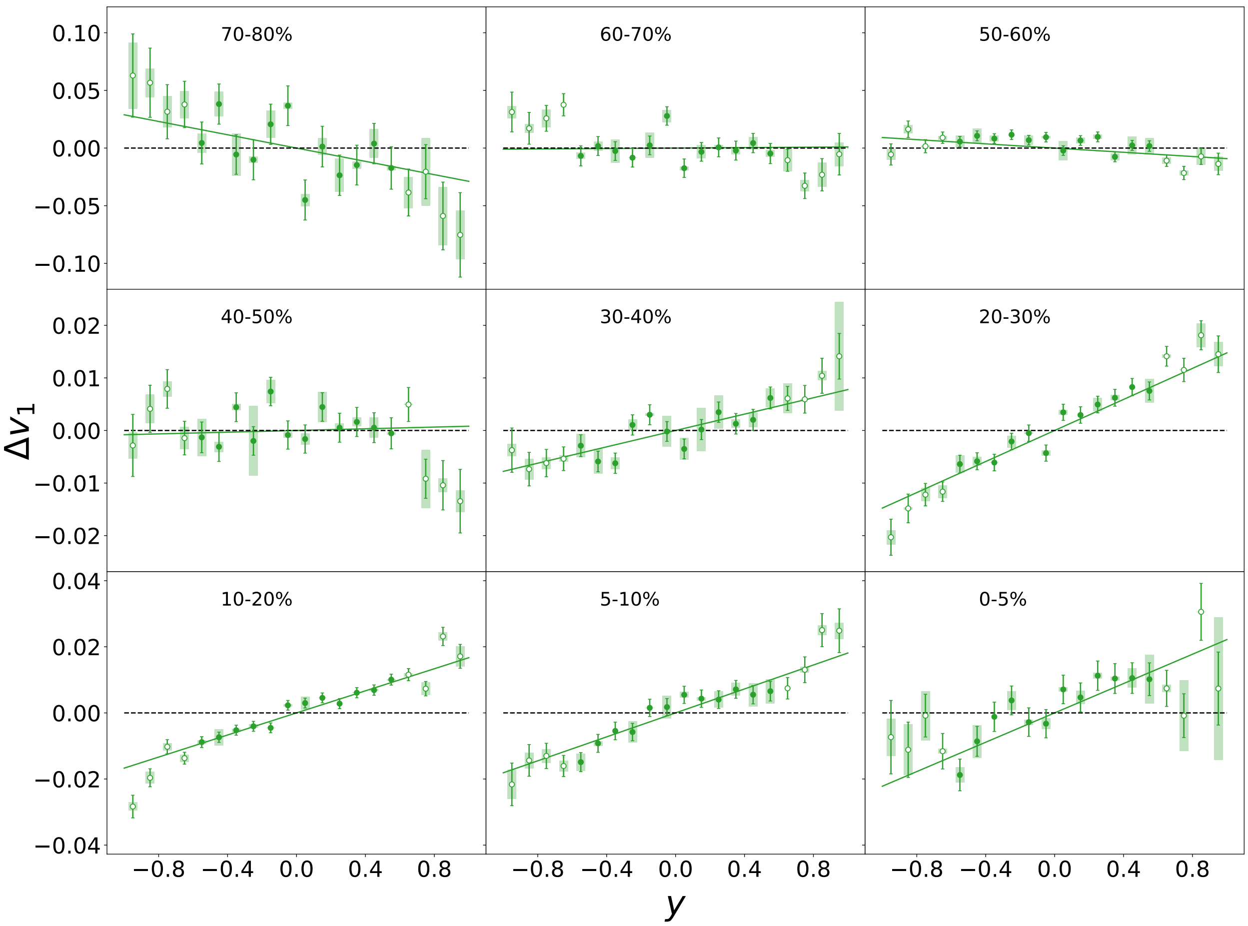

Figure 4.16: $\Delta v_1(y)$ between Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 19.6$ GeV. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

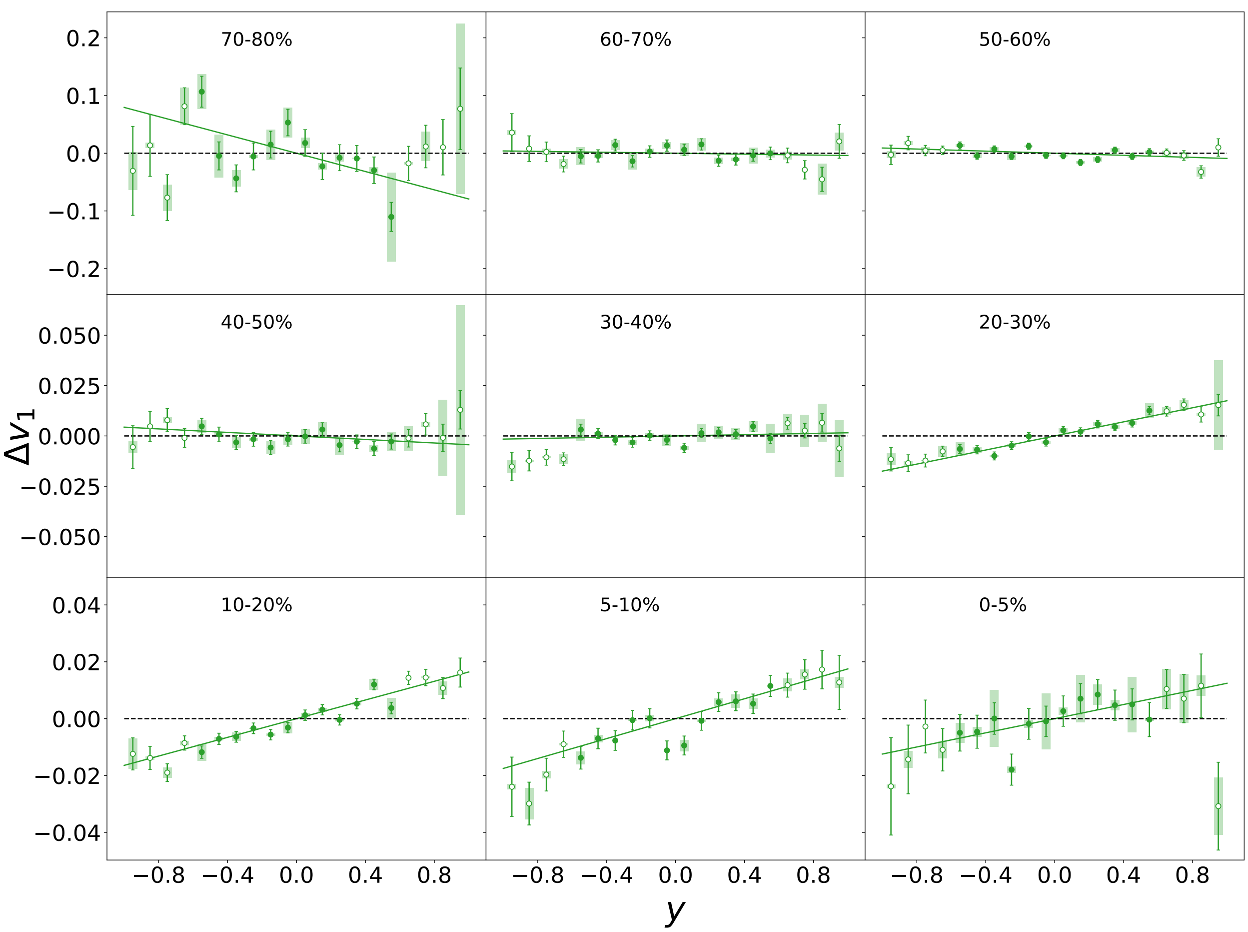

Figure 4.17: $\Delta v_1(y)$ between Λ^0 and $\bar{\Lambda}^0$ at $\sqrt{s_{NN}} = 27$ GeV. A linear fit is applied over the range $|y| < 0.6$; only the filled points, lying within this range, are included in the fit, while the open points at larger rapidity are excluded.

4.2.3 $\Delta(dv_1/dy)$ as a Function of Centrality

Figure 4.18 presents $\Delta(dv_1/dy)$ for $\Lambda^0(\bar{\Lambda}^0)$, $p(\bar{p})$, and $K^+(K^-)$ as a function of centrality at each collision energy. The proton and kaon splittings, measured on the same data set as part of this joint analysis, are shown alongside the Λ measurement to facilitate the coalescence comparison discussed in Sect. 4.2.4. The proton, kaon, and Λ slopes are evaluated over matching momentum-per-quark ranges, each extending to $p_T/n_q = 0.6$ GeV/c, so that the species enter the coalescence sum rule at comparable momentum per constituent quark. Three features are common across all collision energies:

1. **Positive splitting in central collisions, ordered p, Λ, K.** At all collision energies, the proton and Λ exhibit significant positive splittings in the most central events (0–10%), with the proton splitting the larger of the two. The kaon follows the same positive pattern at a smaller magnitude for $\sqrt{s_{NN}} \geq 11.5$ GeV; at 9.2 GeV its central splitting is consistent with zero and at 7.7 GeV it is slightly negative. The ordering tracks the number of transported quarks each species carries: the more transported quarks a hadron contains, the more directed flow it inherits from the incoming nuclei, while its antiparticle, composed of produced quarks, lags behind.

2. **Sign change toward peripheral collisions.** As centrality increases (more peripheral), $\Delta(dv_1/dy)$ for both protons and Λ hyperons decreases, crosses zero, and becomes negative. Baryon transport contributes positively at every centrality, as does the Hall (Lorentz) effect induced by the spectator fields, while the Faraday and Coulomb effects contribute negatively. The crossing marks the impact parameter at which the Coulomb and Faraday contributions, which grow with the increasing number of spectators in more peripheral collisions, overcome the other two.

3. **A neutral hadron whose splitting changes sign.** The Λ carries no net charge, and

its splitting nonetheless turns over from positive to negative with centrality. Baryon transport enters $\Delta(dv_1/dy)$ with a single sign, so it cannot produce that reversal on its own. The electromagnetic effects act on the charged constituent quarks that the Λ and $\bar{\Lambda}$ acquire at coalescence, and supply the negative contribution the reversal requires.

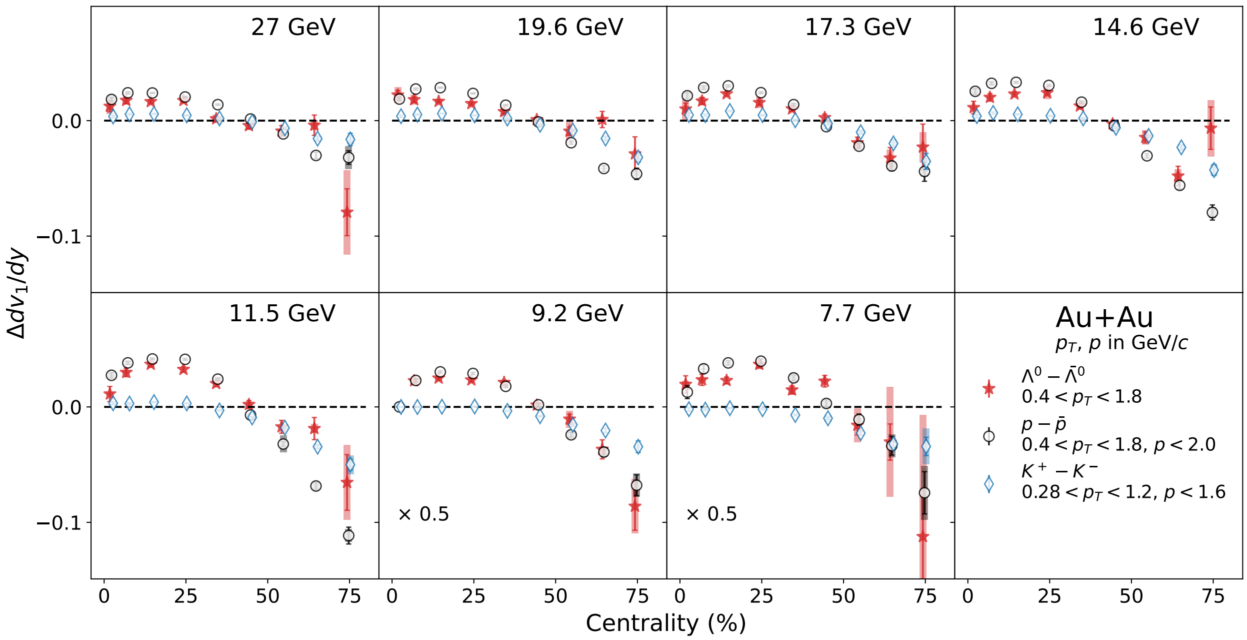

Figure 4.18: $\Delta(dv_1/dy)$ for $\Lambda^0(\bar{\Lambda}^0)$, $p(\bar{p})$, and $K^+(K^-)$ as a function of centrality for $\sqrt{s_{NN}}$ = 7.7, 9.2, 11.5, 14.6, 17.3, 19.6, and 27 GeV. Some panels are scaled for better visualization. Statistical uncertainties are shown as error bars; systematic uncertainties are represented by boxes.

4.2.4 Coalescence Prediction for the Λ Splitting

The central test of this analysis is whether the measured Λ directed flow splitting is consistent with the coalescence sum rule prediction:

$$\Delta\left(\frac{dv_1}{dy}\right)_\Lambda = \Delta\left(\frac{dv_1}{dy}\right)_p - \Delta\left(\frac{dv_1}{dy}\right)_K. \tag{4.1}$$

This relation can be read as a two-level comparison. At first order, the proton splitting $\Delta(dv_1/dy)_p$ alone serves as a natural reference: both the Λ and the proton are three-quark baryons, differing by only one quark ($u \to s$). At the lower BES-II energies, this kinship is reinforced by associative production near threshold ($N + N \to \Lambda + K + N$), which kinematically couples the Λ and proton populations. As a second-order test, the kaon splitting can be subtracted to complete the quark content, since the valence content of the Λ–$\bar{\Lambda}$ pair follows from that of the proton and kaon,

$$(uds) - (\bar{u}\bar{d}\bar{s}) = \big[(uud) - (\bar{u}\bar{u}\bar{d})\big] - \big[(u\bar{s}) - (\bar{u}s)\big], \tag{4.2}$$

where the right-hand side maps term by term onto $\Delta(dv_1/dy)_p$ and $\Delta(dv_1/dy)_K$: the K^+ $(u\bar{s})$ and K^- $(\bar{u}s)$ remove the surplus u of the proton and supply the s of the Λ, giving the prediction of Eq. 4.1.

Figure 4.19 presents both comparisons across three centrality groupings and all seven collision energies, and Table 4.2 lists the corresponding per-energy pulls and the resulting χ^2/ndf. The first-order comparison (proton alone) already reproduces the qualitative behavior of the Λ splitting: the sign, the centrality dependence, and the overall energy trend are all captured, confirming that baryon transport is the dominant driver. However, a systematic offset between the proton and Λ splittings remains visible at all centralities. The second-order coalescence prediction $\Delta(dv_1/dy)_p - \Delta(dv_1/dy)_K$ reduces this offset substantially: χ^2/ndf falls from 5.65 to 1.61 in central collisions (0–10%) and from 5.70 to 1.15 in peripheral collisions (40–80%), both indicating a good description. In mid-central collisions (10–40%) the kaon subtraction again improves the description significantly, more than halving χ^2/ndf from 17.39 to 7.76, despite a residual that remains in this centrality range.

The data prefer the quark-content-matched reference at every centrality, evidence for the coalescence picture. That the proton alone already tracks the Λ closely shows that the

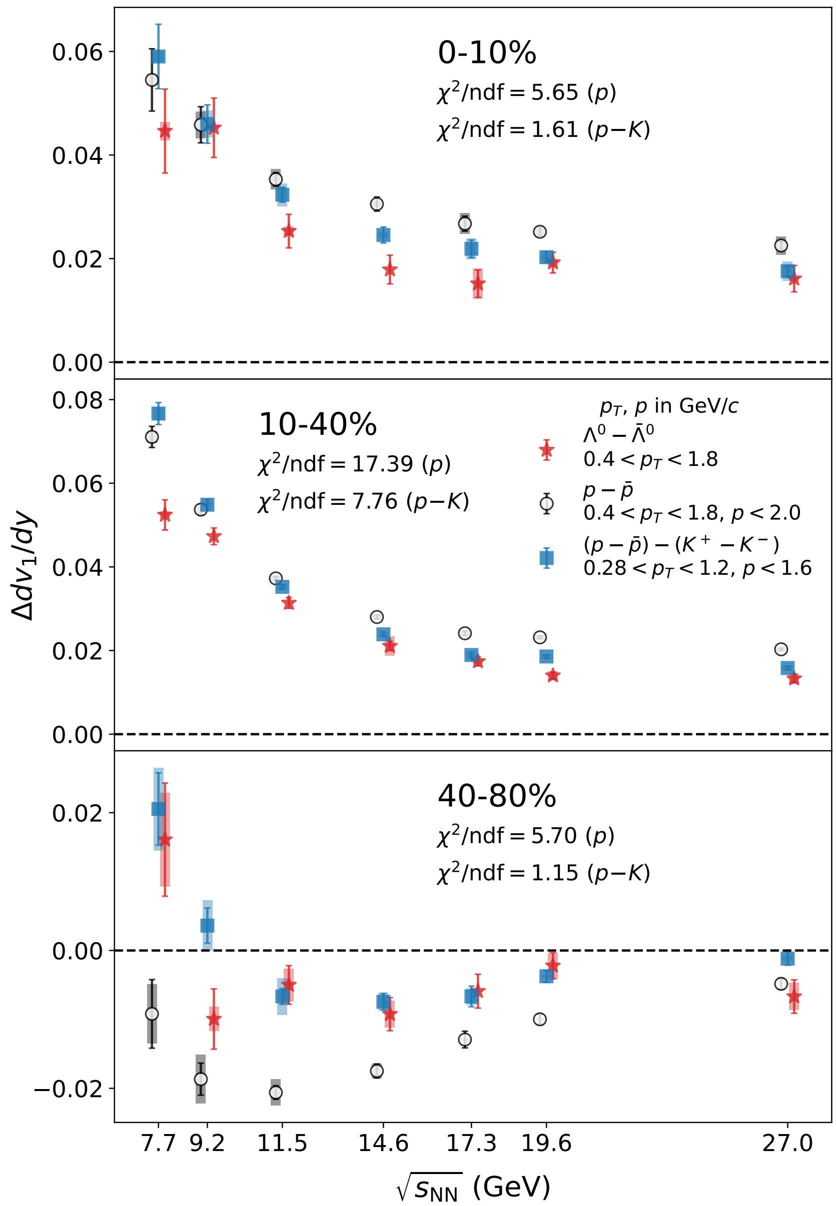

Figure 4.19: $\Delta(dv_1/dy)$ as a function of $\sqrt{s_{NN}}$ in three centrality groupings. Red stars: $\Lambda^0 - \bar{\Lambda}^0$; open circles: $p - \bar{p}$; blue squares: coalescence prediction $(p - \bar{p}) - (K^+ - K^-)$.

Table 4.2: Per-energy pulls of the $\Lambda^0 - \bar{\Lambda}^0$ directed flow splitting against the coalescence references $p \equiv p - \bar{p}$ and p-$K \equiv (p - \bar{p}) - (K^+ - K^-)$, for the three merged centralities (a)–(c). The pull is $(\Lambda - \text{ref})/\sigma$ with σ the total (stat. $\oplus$ sys.) error of the difference; the last column is $\chi^2/\text{ndf} = \sum \text{pull}^2/\text{ndf}$ with ndf = 7.

	$\sqrt{s_{NN}}$ (GeV)							
Reference	7.7	9.2	11.5	14.6	17.3	19.6	27	χ^2/ndf
(a) 0–10%								
Λ vs p	−0.96	−0.08	−2.44	−4.02	−2.50	−2.53	−1.98	5.65
Λ vs p-K	−1.38	−0.10	−1.64	−2.06	−1.43	−0.43	−0.43	1.61
(b) 10–40%								
Λ vs p	−3.99	−2.62	−4.02	−2.61	−4.35	−6.25	−4.24	17.39
Λ vs p-K	−5.08	−3.04	−2.42	−1.05	−0.96	−3.04	−1.46	7.76
(c) 40–80%								
Λ vs p	+2.02	+1.38	+3.67	+2.41	+2.56	+2.78	−0.57	5.70
Λ vs p-K	−0.33	−2.09	+0.35	−0.50	+0.24	+0.54	−1.67	1.15

splitting is governed by the transported quarks the two baryons share, which are present from the earliest stage of the collision; matching the valence content exactly, through the $K^+ - K^-$ term that removes the surplus u and supplies the s, is what a more precise description requires. Because $\Delta(dv_1/dy)_p - \Delta(dv_1/dy)_K$ is assembled from independently measured proton and kaon flow, its success indicates that the constituent quarks contribute additively to the directed flow, each carrying the transport and electromagnetic response of the early collision stage.

The Λ carries no net electric charge, which might suggest that its directed flow is insensitive to the electromagnetic field and that the Λ–$\bar{\Lambda}$ splitting should vanish; instead a robust splitting is observed that changes sign, becoming negative in peripheral collisions. Baryon transport contributes with a single sign and cannot produce this reversal, and an electromagnetic response governed by electric charge alone would cancel among the constituents of a neutral hadron, leaving no net splitting. The sign change therefore indicates that the electromagnetic response acts at the quark level and is not necessarily proportional to charge.

4.3 Summary

The two analyses presented in this chapter provide complementary, quantitative tests of constituent quark coalescence at BES-II energies. The pion elliptic flow splitting yields a transported quark ratio $N_d^{\mathrm{tr}}/N_u^{\mathrm{tr}} \approx 1.14$, consistent with the ^{197}Au nuclear stoichiometry across most energies and centralities: the isospin asymmetry of the colliding nuclei, a ground-state property of the gold nucleus, survives QGP formation, collective expansion, and hadronization, remaining encoded in the collective flow of the final-state pions. The Λ directed flow splitting is reproduced by the coalescence sum rule $\Delta(dv_1/dy)_\Lambda = \Delta(dv_1/dy)_p - \Delta(dv_1/dy)_K$ built from independently measured proton and kaon flow, demonstrating that the same coalescence picture governs directed flow and that the constituent quarks, including the strange quark, respond to the early-time electromagnetic field as independent degrees of freedom before hadronization.

In the directed-flow comparison the lowest energy stands apart: the Λ splitting departs most strongly from the coalescence prediction at $\sqrt{s_{NN}} = 7.7$ GeV in mid-central collisions (a pull of -5.1σ). The pion ratio at this energy remains consistent with the stoichiometric value once systematics are included.

Chapter 5

Ω–Hadron Correlations and Baryon Number Transport

5.1 Introduction

At BES-II energies ($\sqrt{s_{NN}} = 7.7$–27 GeV), baryon stopping transfers a significant fraction of the initial baryon number to midrapidity. The resulting baryon–antibaryon surplus reaches as far as the Ω, a hyperon built entirely of strange quarks and sharing no valence flavor with the incoming nuclei: at $\sqrt{s_{NN}} = 19.6$ GeV, $N(\Omega^-) \approx 2\,N(\bar{\Omega}^+)$. Because strangeness is produced in pairs, this surplus cannot arise from pair production; it represents net baryon number that has been transported from the incoming nuclei all the way into the multi-strange sector. A substantial fraction of the Ω^- observed at midrapidity is therefore associated with the transport of baryon number. Characterizing this transported component, where it sits in phase space and what it reveals about the production mechanism, is the central aim of this chapter.

The *excess distribution* is defined as that part of the Ω^- phase-space distribution which $\bar{\Omega}^+$ cannot account for. The Ω^- distribution is decomposed as

$$f_{\Omega}(x) = R \cdot f_{\bar{\Omega}}(x) + (1 - R) \cdot g(x), \tag{5.1}$$

where $f_{\bar{\Omega}}(x)$ is the pair-production-like component (identical to the $\bar{\Omega}^+$ distribution), $g(x)$ is the excess distribution reflecting where baryon transport manifests in phase space, and $R = N(\bar{\Omega}^+)/N(\Omega^-)$ is the measured $\bar{\Omega}^+/\Omega^-$ ratio. The decomposition applies to the density rather than to individual particles, and *excess* Ω^- is used throughout as shorthand for the component g. In the excess distribution, strangeness conservation requires that each unit of transported baryon number be accompanied by compensating anti-strangeness, produced as associated K^+ and K^0 mesons.

The Ω^- hyperon (sss) is a uniquely sensitive probe of baryon number transport. Because it contains no up or down valence quarks, the Ω^- shares no valence quark content with the incoming gold nuclei: the net baryon number associated with the excess Ω^- must have been transported from the initial nuclei. Two microscopic pictures have been proposed for how this transport occurs. In the *valence quark picture*, baryon number is carried by the valence quarks of the nucleon. During the collision, valence quarks lose energy through multiple soft interactions and can be stopped at midrapidity, where they recombine with other quarks to form baryons [13]. In this picture, the baryon number of the Ω^- ultimately traces back to the valence quarks of the incoming nuclei, even though the valence quark content is fully replaced by sea strange quarks during the hadronization process. In the *gluon junction picture* [35], baryon number resides not in the quarks but in a Y-shaped topological configuration of the gluon field where three color flux tubes meet. The junction is flavor-blind: it captures whatever sea quarks are available during hadronization, making Ω^- production natural in a strangeness-rich environment. The gluon junction is a soft, gluonic degree of freedom that can be transported over large rapidity intervals, a feature that distinguishes it kinematically from valence quark transport, which tends to keep the baryon closer to beam rapidity [38]. Distinguishing the relative contributions of these two mechanisms requires theoretical model input. The experimental goal of this analysis is to characterize the excess distribution $g(x)$ and its strangeness-conservation pattern: to determine where in phase space baryon transport manifests and how kaon associations differ in the excess-enriched regime, providing observables that can be compared to model predictions [52].

A standard femtoscopic correlation function $C(k^*)$, normalized to unity at large relative momentum, divides out the pair yield in which the baryon-transport signal resides. Because $K^+\Omega^-$ and $K^-\bar{\Omega}^+$ pairs share the same final-state interaction, the excess K^+ associated with the transported Ω^- leaves the normalized correlation unchanged, as confirmed by a previous study with A Multi-Phase Transport (AMPT) model [52]. Recovering it requires a yield-preserving observable; the simplest is a per-Ω correlation *difference*, in which the pair-produced $\bar{\Omega}^+$ combination is subtracted as a baseline from its Ω^- partner.

The Combinatorial Background Subtraction (CBS) method generalizes this first-order difference, exploiting the four charge combinations ($K^+\Omega^-$, $K^-\Omega^-$, $K^+\bar{\Omega}^+$, $K^-\bar{\Omega}^+$) together with the constraint that $\bar{\Omega}^+$ is a pure pair-produced sample. This algebraic decomposition preserves both the shape and the magnitude of the correlation, enabling the excess component to be isolated from the pair-production-like baseline at the ensemble level.

This chapter develops the CBS (Combinatorial Background Subtraction) correlation method (Sect. 5.4) to characterize the excess distribution and the pair-production-like baseline. The method is applied to K–Ω correlations as a function of the rapidity difference $\Delta y = y_K - y_\Omega$, the difference of absolute rapidities $\Delta|y| = |y_K| - |y_\Omega|$, and the Lorentz-invariant relative momentum k^*, decomposing the four charge combinations into the excess signal and its pair-production baseline at the ensemble level.

5.2 Data Sample and Event Selection

The analysis uses Au+Au collision data collected by the STAR detector during the BES-II program at $\sqrt{s_{NN}} = 19.6$ GeV. Only the most central 0–10% of collisions are analyzed, where baryon stopping and the Ω^- surplus are largest and where the Ω yield, and with it the number of K–Ω pairs available to the decomposition, is greatest.

Event-level selection criteria follow the same procedures described in Sect. 3.2 for the flow analyses. The primary vertex is required to satisfy the longitudinal and radial constraints summarized in Table 3.1, pile-up events are rejected via the TOF–TPC multiplicity correlation, and runs flagged on the official STAR bad-run list are excluded.

The bachelor kaon enters the Ω reconstruction as a KFParticle daughter track and is subject to the same track quality requirements as the Λ daughters (Sect. 3.4.2): a number of TPC fit hits nHitsFit ≥ 15 and a dE/dx error $\sigma_{dE/dx} \in [0.04, 0.12]$.

5.3 Particle Identification and Ω Reconstruction

The Ω^- is reconstructed via its dominant weak decay channel $\Omega^- \rightarrow \Lambda K^-$ ($BR \approx$ 67.8%) [44], followed by $\Lambda \rightarrow p\pi^-$. Reconstruction uses the KFParticle Kalman-filter toolkit [30, 34] integrated into the STAR analysis framework.

5.3.1 Particle Identification

The requirements collected in Table 5.1 apply to the daughter tracks of the Ω decay chain ($\Omega^- \to \Lambda K^- \to p\pi^- K^-$) at the KFParticle reconstruction stage; the selection of correlation-partner kaons is treated separately below. All daughter tracks are identified using TPC dE/dx in terms of the number of standard deviations from the expected Bethe–Bloch value ($n\sigma$). The bachelor kaon carries the tightest requirement of the three because it is the track that distinguishes an Ω from the far more abundant Ξ, whose bachelor is a pion; it additionally requires a TOF match for momenta $0.5 < p < 2.0$ GeV/c, where the kaon and pion dE/dx bands merge and the TPC alone no longer separates them.

Table 5.1: Particle-identification requirements applied to the daughters of the Ω decay chain.

Daughter	Requirement
Proton	$\lvert n\sigma_p \rvert < 3.0$
Pion	$\lvert n\sigma_\pi \rvert < 3.0$
Bachelor kaon	$\lvert n\sigma_K \rvert < 2.0$
Bachelor kaon, TOF	match required for $0.5 < p < 2.0$ GeV/c

5.3.2 Λ Reconstruction

Proton–pion pairs are fitted to a common secondary vertex by the same Kalman-filter procedure used in the directed-flow analysis (Sect. 3.4.2). A Λ entering a cascade originates at the Ω decay point rather than at the collision vertex, and only candidates displaced from the primary vertex are passed to the cascade fit. The requirements are listed in Table 5.2.

5.3.3 Ω Reconstruction

Each Λ candidate is combined with a bachelor K^- candidate, and the pair is fitted to a common cascade vertex. The Ω is produced at the collision vertex, so its trajectory

Table 5.2: Topological requirements applied to Λ candidates within the Ω decay chain.

Parameter	Value
Daughter distance of closest approach, DCA(p, π)	< 1.0 cm
χ^2 of each daughter to the primary vertex	> 3.0
χ^2/NDF of the Λ vertex fit	< 10.0
Decay-length significance, $l/\delta l$	> 5.0
Decay length, L_Λ	> 1.0 cm

is additionally constrained to originate there, and the χ^2/NDF of that constrained fit is required to be small. The selection criteria are listed in Table 5.3.

Table 5.3: Selection criteria applied at the Ω cascade-vertex fit.

Parameter	Value
χ^2 of the bachelor track to the primary vertex	> 3.0
χ^2/NDF of the cascade vertex fit	< 6.0
χ^2/NDF with the Ω constrained to the primary vertex	< 5.0
Decay-length significance, $l/\delta l$	> 10.0
Daughter sharing between Λ and bachelor	not allowed

The mass resolution is not uniform across the acceptance: it degrades away from midrapidity, and the signal-to-background ratio varies by more than an order of magnitude between the midrapidity and forward slices. A single global mass window would therefore hold a different signal fraction in different regions of phase space, so the signal and sideband windows are instead defined slice by slice. The candidate sample is divided into bins of (p_T, y) over the accepted range $|y| < 1.6$, with no explicit transverse-momentum requirement; the bin boundaries are $p_T = 0.8$, 1.4, 2.0 and 3.0 GeV/c and $y = -0.6$, 0 and 0.6, the outermost rapidity bins being wide because the yield there is sparse. Within each slice the peak position μ and the width are taken from the invariant-mass fit of that slice. The signal window $[\mu - w, \mu + w]$ is fixed by solving for the half-width w that contains the fraction 0.9973 of the fitted signal shape, the fraction a single Gaussian holds within $\pm 3\sigma$. Defining the window by containment rather than by a fixed multiple of the width keeps the accepted signal fraction the same in every slice even though the peak is not a single Gaussian; because the fitted shape carries broader tails than a single Gaussian, the resulting half-width exceeds three times the root-mean-square width of the peak. The sidebands are taken as the two intervals $\mu \pm (5\text{–}9)\,\sigma_\Omega$, with σ_Ω the root-mean-square width of the same slice. They sit far enough from the peak that the signal leaking into them is small, and close enough to follow the local shape of the background they are used to estimate.

5.3.4 Kaon Selection

Kaons used as correlation partners in the CBS analysis are selected independently of the Ω daughter tracks, with the requirements listed in Table 5.4. These are primary-track requirements: the partner kaons are prompt particles, not displaced decay products. All Ω daughters, including those of sideband candidates, are also vetoed from the partner sample.

Table 5.4: Selection requirements for the correlation-partner kaons.

Parameter	Value
Transverse momentum, p_T	0.15–2.0 GeV/c
Pseudorapidity, $\lvert\eta\rvert$	< 1.5
TPC PID, $\lvert n\sigma_K\rvert$	< 2.0, recentered in bins of p
TOF mass squared, m^2	0.15–0.34 GeV$^2/c^4$, required for $p_T > 0.4$ GeV/c
DCA to primary vertex	< 2.0 cm
Ω daughter veto	applied, signal and sideband candidates

5.4 CBS Correlation Method

To probe how the Ω is correlated with the kaons that balance its strangeness, a natural starting point would be the standard femtoscopic correlation function $C(X) = S(X)/M(X)$, in which division by the mixed-event distribution $M(X)$ removes single-particle acceptance and efficiency effects and isolates dynamical pair-level correlations: Coulomb, strong final-state interactions, and $s\bar{s}$ pair-production correlations. The femtoscopically relevant pairings are $K^+\Omega^-$ and $K^-\bar{\Omega}^+$, both of which carry the same $s\bar{s}$ pair-production correlation. Any baryon-transport-driven asymmetry between Ω^- and $\bar{\Omega}^+$ is therefore unresolvable at the femtoscopic level: a mixed-event-divided observable would be dominated by the common pair-production correlation and would obscure the signal of interest. Resolving it requires observables that preserve the pair yield rather than dividing it out. The remainder of this section constructs these from the same-event distributions, building from the simplest per-Ω yield difference up to the full CBS decomposition.

5.4.1 Observable Construction

The signal we wish to extract is the asymmetric component unique to Ω^-: correlations of the excess Ω^- population, associated with net baryon number transported from the colliding nuclei, with the kaons that balance its strangeness. This component has no counterpart in $\bar{\Omega}^+$, which is purely pair-produced. Both observables constructed below therefore work directly with the same-event pair distribution $S(X)$ for each of the four pair combinations ($K^+\Omega^-$, $K^-\Omega^-$, $K^+\bar{\Omega}^+$, $K^-\bar{\Omega}^+$), and differ only in how they subsequently remove the pair-production baseline. Here X denotes the differential variable in which the correlation is analyzed: the rapidity difference $\Delta y = y_K - y_\Omega$ in the laboratory frame, the longitudinal separation of the pair; the difference of absolute rapidities $\Delta|y| = |y_K| - |y_\Omega|$, which registers whether the kaon sits farther from midrapidity than the Ω ($\Delta|y| > 0$) or nearer to it ($\Delta|y| < 0$); the Lorentz-invariant relative momentum k^* in the pair rest frame; or single-particle variables (y, p_T, ϕ) of the Ω. Denoting by $\vec{p}^{\,*}_K$ and $\vec{p}^{\,*}_\Omega$ the momenta of the two particles in the pair rest frame, where $\vec{p}^{\,*}_K = -\vec{p}^{\,*}_\Omega$, the relative momentum is

$$k^* = \tfrac{1}{2}\,|\vec{p}^{\,*}_K - \vec{p}^{\,*}_\Omega| = |\vec{p}^{\,*}_K|, \tag{5.2}$$

with the boost to that frame constructed from the measured track momenta and the PDG masses $m_K = 493.7$ MeV$/c^2$, $m_\Omega = 1672.5$ MeV$/c^2$ [44].

The raw correlation functions contain a contamination from combinatorial background under the Ω invariant-mass peak, which is removed by sideband subtraction. The correction is carried out in the same (p_T, y) slices in which the mass windows are defined (Sect. 5.3.3), so that each Ω is classified by the windows of its own slice and subtracted with the signal-to-background ratio measured there. Writing $S_{\rm sig}(X)$ and $S_{\rm SB}(X)$ for the same-event pair distributions built from the signal-window and sideband candidates of a slice, the purity-corrected distribution is

$$S(X) = S_{\rm sig}(X) - \frac{N_{\rm bkg}}{N_{\rm SB}}\, S_{\rm SB}(X), \qquad N_{\rm bkg} = \frac{N_{\rm sig}}{(S/B)+1}, \tag{5.3}$$

where $N_{\rm sig}$ and $N_{\rm SB}$ are the candidate counts in the signal and sideband windows of that slice and S/B is its fitted signal-to-background ratio. The scale factor $N_{\rm bkg}/N_{\rm SB}$ transports the sideband pair distribution to the amount of background actually present beneath the peak. The corrected distributions are then summed over slices, which allows a single set of correlations to be built from a sample whose purity ranges from below unity in the forward slices to well above it at midrapidity, without a global window imposing either extreme on the whole acceptance.

5.4.2 First-Order Correlation Difference

The simplest observable built on these same-event distributions is a per-Ω correlation *difference*. Each purity-corrected combination is normalized to its parent-hyperon yield (N_Ω for the Ω^- triggers and $N_{\bar\Omega}$ for the $\bar\Omega^+$ triggers, so that the normalized distribution is the mean number of kaons per hyperon in the bin X), and the same-strangeness-sign combination is subtracted from the opposite-strangeness-sign one,

$$\begin{aligned} D_+(X) &= \frac{S(K^+\Omega^-;\, X)}{N_\Omega} - \frac{S(K^+\bar{\Omega}^+;\, X)}{N_{\bar{\Omega}}}, \\ D_-(X) &= \frac{S(K^-\bar{\Omega}^+;\, X)}{N_{\bar{\Omega}}} - \frac{S(K^-\Omega^-;\, X)}{N_\Omega}, \end{aligned} \tag{5.4}$$

where $S(\cdot\,; X)$ is the purity-corrected same-event pair distribution of the given charge combination in the bin X. Each normalized combination is dominated by combinatorial pairings rather than genuinely associated kaons. Since $\bar{\Omega}^+$ is purely pair-produced, the mean number of K^+ per $\bar{\Omega}^+$ reproduces the pair-production baseline, and subtracting it in D_+ leaves a proxy for the excess K^+ associated with the transported Ω^-. Its charge conjugate D_- is built on the pair-produced $\bar{\Omega}^+$ and serves as a null reference. This difference, however, subtracts one raw same-event combination directly from another without an independent estimate of the combinatorial baseline underlying each. The CBS decomposition developed in the remainder of this section mitigates this by replacing the single pairwise subtraction with an independent mixed-event estimate of the baseline for all four charge combinations.

5.4.3 CBS Decomposition

The CBS method decomposes the four purity-corrected same-event correlations into physically motivated components. As established in the chapter introduction (Eq. 5.1), the Ω^- phase-space density is the sum of a pair-production-like component, identical in shape to the $\bar{\Omega}^+$ density, and an excess component g associated with the excess Ω^- population. Because a pair distribution is linear in the single-particle densities, every measured K–Ω^- correlation inherits this split. As an example, the same-event $K^-\Omega^-$ distribution is an integral over the single-particle densities weighted by a pair kernel $\mathcal{K}$ that carries the genuine final-state and strangeness correlations,

$$S(K^-\Omega^-;X) = \int \rho_{K^-}(p_K)\,\rho_{\Omega^-}(p_\Omega)\,\mathcal{K}(p_K,p_\Omega)\,\delta(X - X(p_K,p_\Omega))\;\mathrm{d}p_K\,\mathrm{d}p_\Omega. \tag{5.5}$$

Since K^- has no excess component, the split here comes entirely from the Ω^- density. Inserting $\rho_{\Omega^-} = Rf_{\bar{\Omega}} + (1-R)g$ from Eq. 5.1 factors the correlation into the K^- correlation with the pair-production-like Ω component and with the excess component,

$$S(K^-\Omega^-) = \underbrace{R\int \rho_{K^-} f_{\bar{\Omega}}\,\mathcal{K}}_{\text{pair-production-like}} + \underbrace{(1-R)\int \rho_{K^-} g\,\mathcal{K}}_{\text{excess}}. \tag{5.6}$$

The kaons are decomposed in the same way: an excess-associated component (subscript t, "transported"), which balances the strangeness of the excess Ω^-, and a pair-production component (subscript p).

Writing each kaon and Ω density as the sum of its t and p parts, the four measured same-event correlations (denoted $S(X)$ for charge combination X) expand into products of sub-densities,

$$\begin{aligned}
S(K^+\Omega^-) &= K_t^+\Omega_t^- + K_p^+\Omega_t^- + K_t^+\Omega_p^- + K_p^+\Omega_p^-,\\
S(K^+\bar{\Omega}^+) &= K_t^+\bar{\Omega}_p^+ + K_p^+\bar{\Omega}_p^+,\\
S(K^-\bar{\Omega}^+) &= K_p^-\bar{\Omega}_p^+,\\
S(K^-\Omega^-) &= K_p^-\Omega_t^- + K_p^-\Omega_p^-,
\end{aligned} \tag{5.7}$$

where K^- and $\bar{\Omega}^+$, being pair-produced, carry no excess component. If pair production is taken to be symmetric under charge conjugation, so that pair-produced particles and antiparticles share the same single-particle densities and their correlations with other produced particles are invariant under conjugating both members of a pair, the decomposition yields the identities $K_p^+\Omega_p^- = K_p^-\bar{\Omega}_p^+$ and $K_p^+\bar{\Omega}_p^+ = K_p^-\Omega_p^-$. Conjugating only one member reverses the relative charge of the pair, and with it the Coulomb and strong final-state interaction. This clearly separates the fully pair-produced $K_p^+\Omega_p^-$ and $K_p^-\Omega_p^-$ (the opposite- and like-sign combinations).

With these identities the sub-densities collapse to seven independent terms, for which we introduce the shorthands

$$\begin{aligned}
&a = K_t^+\Omega_t^-, \\
&b = K_p^+\Omega_t^-, \qquad b' = K_p^-\Omega_t^-, \\
&c = K_t^+\Omega_p^-, \qquad c' = K_t^+\bar{\Omega}_p^+, \\
&d = K_p^+\Omega_p^- = K_p^-\bar{\Omega}_p^+, \\
&e = K_p^-\Omega_p^- = K_p^+\bar{\Omega}_p^+,
\end{aligned} \tag{5.8}$$

where a is the excess signal of interest, (b, b') a pair-produced kaon paired with a transported Ω, and (c, c') an excess kaon paired with a pair-produced Ω. Substituting into Eq. 5.7, the four measured correlations take the compact form

$$\begin{aligned}
&S(K^+\Omega^-) = a + b + c + d, \\
&S(K^+\bar{\Omega}^+) = c' + e, \\
&S(K^-\bar{\Omega}^+) = d, \\
&S(K^-\Omega^-) = b' + e.
\end{aligned} \tag{5.9}$$

5.4.4 Event Mixing

Isolating the excess signal a in the system of Eq. 5.9 requires knowledge of the pair-production baseline terms. Those not measured directly in a single charge combination are estimated from mixed events, in which a kaon and an Ω taken from two different collisions are combined into a pair. Because the two particles come from unrelated events, the mixed-event distribution $M(X)$ reproduces the single-particle acceptance and phase space of each species while carrying none of the pair-level correlations (Coulomb, strong final-state, or strangeness), and thus provides the uncorrelated reference from which the baseline terms are built. Each mixed-event distribution is normalized to its same-event counterpart in the high-k^* region $k^* \in [0.8, 1.5]$ GeV/c, under the assumption that above this k^* threshold the strangeness and final-state correlations have died out; the factor fixed in this band is then applied over the full k^* range.

For observables differential in a rapidity variable (the pair rapidity difference Δy, the difference of absolute rapidities $\Delta|y|$, or the single-Ω rapidity $|y|$), the same-event and mixed-event distributions are accumulated as two-dimensional histograms in (k^*, var). The normalization and the subsequent baseline subtraction both depend on k^*, so they are applied in each k^* bin first, and only then is the distribution integrated over k^* to give its dependence on the rapidity variable.

5.4.5 Signal and Baseline Extraction

The quantity of physical interest is $a = K_t^+\Omega_t^-$, the correlation of the excess kaons with the excess Ω^-. It is compared against the baseline $c = K_t^+\Omega_p^-$, the correlation of the excess kaons with a pair-production-like Ω. The baseline replaces the excess Ω^- with an ordinary produced Ω, so comparing a/N_t with c/N_p (per excess and per pair-production-like Ω, with $N_t = N_\Omega - N_{\bar{\Omega}}$ and $N_p = N_{\bar{\Omega}}$) asks directly whether the transported Ω^- shares the strangeness-conservation pattern of an ordinary produced Ω: a difference in shape between them would show that baryon transport gives the excess Ω^- a distinct kaon-association pattern. Extracting a and c for this comparison is the goal of the decomposition.

The reconstruction rests on three assumptions:

1. Pair production is symmetric under charge conjugation: pair-produced particles and antiparticles share the same single-particle densities, and their correlations with other produced particles are invariant under conjugating both members of a pair. This is the symmetry that collapses the sub-densities to the seven terms of Eq. 5.8.

2. The uncorrelated reference of the e pairs is that of the d pairs, $M(K_p^-\Omega_p^-) \equiv M_e = M_d \equiv M(K_p^-\bar{\Omega}_p^+) = M(K^-\bar{\Omega}^+)$, so that the single measured mixed-event distribution supplies the phase space of both.

3. The short-range final-state interaction within a baseline pair, one joining an independently produced particle with a transported one, as in b, b', c and c', is negligible.

The four correlations of Eq. 5.9 form a system of four equations in seven unknowns (a, b, b', c, c', d, e), and neither a nor c is measured directly: both are locked inside the four-term correlation $S(K^+\Omega^-) = a + b + c + d$, which no charge combination separates. They must instead be reconstructed from the other three correlations. The pair-production term $d = S(K^-\bar{\Omega}^+)$ is measured directly, while b and c are accessed through the partner terms b' and c', each formed by conjugating a single member of the pair: these appear in the two-term correlations $S(K^-\Omega^-) = b' + e$ and $S(K^+\bar{\Omega}^+) = c' + e$, and so are determined once the lone pair-production term e is known. The extraction therefore proceeds by reconstructing e, from which b' and c' follow and serve as estimates of the unmeasured b and c.

The term $e = K_p^-\Omega_p^-$ pairs a pair-produced kaon with a pair-produced Ω; sharing a valence s quark, the two carry a genuine correlation (a same-strangeness anti-correlation and a like-sign Coulomb repulsion). By the second assumption above, the measured mixed event $M(K^-\bar{\Omega}^+)$ supplies the uncorrelated reference M_e of the e pairs, carrying their single-particle densities but not the correlation, so e is estimated as this mixed event modified by the measured correlation,

$$e(X) = C_e(k^*)\, M(K^-\bar{\Omega}^+; X), \qquad C_e \equiv \frac{S_e}{M_e}, \tag{5.10}$$

where $C_X \equiv S_X/M_X$ is the same- to mixed-event correlation of any pair X, normalized to unity at large k^*. As the pair-produced pairs cannot be isolated in the same event, C_e is taken from the ratio

$$C(K^-\Omega^-) \equiv \frac{S(K^-\Omega^-)}{M(K^-\Omega^-)}. \tag{5.11}$$

This ratio mixes the correlated e with the transported term b' (uncorrelated, $C_{b'} \approx 1$). Being the R-weighted average of the two, it inverts to the un-diluted correlation

$$C(K^-\Omega^-) = R\, C_e + (1 - R) \quad \Longrightarrow \quad C_e = 1 + \frac{C(K^-\Omega^-) - 1}{R}, \tag{5.12}$$

with $R = N_{\bar{\Omega}}/N_\Omega$ the pair-production fraction.

With e known, the two partners follow by subtraction,

$$b'(X) = S(K^-\Omega^-; X) - e(X), \qquad c'(X) = S(K^+\bar{\Omega}^+; X) - e(X). \tag{5.13}$$

The excess signal is obtained by removing the baselines from $S(K^+\Omega^-)$, with b' and c' in place of b and c,

$$\begin{aligned} \hat{a} &= S(K^+\Omega^-) - d - b' - c' \\ &= S(K^+\Omega^-) - S(K^-\bar{\Omega}^+) - \big[S(K^-\Omega^-) - e\big] - \big[S(K^+\bar{\Omega}^+) - e\big] \\ &= S(K^+\Omega^-) - S(K^+\bar{\Omega}^+) - S(K^-\Omega^-) - S(K^-\bar{\Omega}^+) + 2\, C_e\, M(K^-\bar{\Omega}^+), \end{aligned} \tag{5.14}$$

so that the four same-event correlations enter directly and the only mixed-event input is the modified e term; the plotted baseline is correspondingly

$$c' = S(K^+\bar{\Omega}^+) - C_e\, M(K^-\bar{\Omega}^+). \tag{5.15}$$

The extracted excess signal $\hat{a}$ and the baseline c' are the two components of the *CBS correlation*, written C_{CBS} and reported per relevant Ω as $\hat{a}/N_t$ and c'/N_p. Unlike the femtoscopic $C(X) = S/M$ of Sect. 5.4, C_{CBS} is not a same- to mixed-event ratio but a yield-preserving combination of the four same-event correlations, so it retains the pair-yield asymmetry that a normalized correlation discards.

The substitution of b' for b and c' for c relies on the third of these. Each partner is related to the term it estimates by conjugating a single member of the pair (the kaon for b and b', the Ω for c and c'), so the two share the same single-particle content and, joining an independently produced particle with a transported one, carry no genuine correlation; they differ only in the short-range final-state interaction, and neglecting it also sets $C_{b'} = 1$ in Eq. 5.12. Under this assumption the partners coincide, and the estimator returns

$$\hat{a} = a + (b - b') + (c - c'), \tag{5.16}$$

with residuals $b - b'$ and $c - c'$. These arise from the charge-dependent final-state interaction and are confined to low k^*; in the rapidity-differential observables Δy and $|y|$, where the correlations are integrated over k^*, their effect is strongly suppressed.

Because the Ω yield per event is far below one, a pair-production-like Ω event contains essentially no transported Ω: the baseline c'/N_p, with $N_p = N_{\bar{\Omega}}$ the number of pair-production-like Ω, then counts only the ambient excess kaons that balance the far more abundant transported Λ and Ξ. The signal $\hat{a}/N_t$, with $N_t = N_\Omega - N_{\bar{\Omega}}$ the number of transported Ω, contains those same ambient kaons and, in addition, the excess kaons that balance the transported

Ω^- itself. Their integrated difference

$$\frac{\int \hat{a}\, dX}{N_t} - \frac{\int c'\, dX}{N_p} \tag{5.17}$$

is therefore the Ω-specific transported kaon balancing, nonzero whenever baryon transport enriches the kaon association of the excess Ω^-. The ambient term common to both cancels in the difference up to a small event-level contribution, since selecting a transported Ω^- favors events carrying marginally more transported Λ and Ξ.

5.4.6 Systematic Uncertainties

The systematic uncertainty on the integrated correlation difference $\hat{a}/N_t - c'/N_p$ is assessed by repeating the full analysis chain with varied selection and normalization criteria; the sources and their variations are listed in Table 5.5. For each variation the decomposition is re-extracted and the integrated difference is compared with the default, and the significance of the deviation is judged with the Barlow test as described in Sect. 3.3.5. For the vertex-position variation the mixed-event vertex binning is matched to the tightened acceptance, so the baseline is not distorted by the change in coverage. The total is taken as the quadrature sum of the significant deviations, divided by $\sqrt{3}$ under a uniform-variation assumption.

Table 5.5: Systematic-uncertainty sources and their variations for the CBS K–Ω correlation analysis.

#	Source	Default	Variation
1	Vertex z position	$\|V_z\| < 70$ cm	< 35 cm
2	nHitsFit	≥ 15	≥ 20
3	Ω mass window	$3\sigma_\Omega$	$2.5\sigma_\Omega$
4	Mixed-event normalization band	$k^* \in [0.8, 1.5]$ GeV/c	$k^* \in [1.0, 1.5]$ GeV/c

One source intrinsic to the CBS method is included in the table (row 4): the mixed-event normalization anchor, whose uncertainty is the spread of the extraction over the correlation-free bands above the strangeness tail.

The final-state-interaction assumption of the method, that the short-range interaction between the uncorrelated baseline pairs is negligible, cannot be assigned a systematic by variation. A point-Coulomb estimate would bound its electromagnetic part, but the accompanying strong K–Ω interaction, for which no independent determination exists, would remain unconstrained. Its effect is instead bounded by the confinement of the residuals $(b - b')$ and $(c - c')$ to low k^*, where they are strongly suppressed in the rapidity-differential observables.

Chapter 6

Results: Ω–Hadron Correlations

6.1 Overview

This chapter presents the results of the CBS combinatorial-background-subtraction correlation analysis (Sect. 5.4) developed in Chapter 5, reported at $\sqrt{s_{NN}} = 19.6$ GeV. The signal extraction (Sect. 6.2), the femtoscopic correlation functions (Sect. 6.3), the first-order per-Ω correlation difference (Sect. 6.4), and the extracted CBS correlations (Sect. 6.5) are reported.

6.2 Signal Extraction

The Ω candidates are reconstructed as described in Sect. 5.3. The yield is extracted differentially: the sample is divided into bins of (p_T, y), the invariant-mass spectrum in each bin is fitted with a double-Gaussian signal on a quadratic background using the IMINUIT minimizer, and the fitted signal yields are summed over all (p_T, y) bins. Within each slice the peak is described by a double Gaussian, a sharp core with a wider resolution tail, and the signal window is set to contain 99.7% of the fitted signal. Figure 6.2 shows a representative slice, at midrapidity and intermediate p_T, where the peak sits sharply on a flat background; the full set of per-slice fits across the (p_T, y) acceptance is collected in Appendix A. Figure 6.1 shows the summed fit overlaid on the acceptance-integrated spectra for Ω^- and $\bar{\Omega}^+$ at $\sqrt{s_{NN}} = 19.6$ GeV, with summed signal yields

$$
\begin{aligned}
N(\Omega^-) &= (7.00 \pm 0.05) \times 10^4, \\
N(\bar{\Omega}^+) &= (3.84 \pm 0.03) \times 10^4,
\end{aligned}
\tag{6.1}
$$

quoted in the signal window defined above, with mean effective resolution $\langle\sigma\rangle \approx 2.6$ MeV$/c^2$. The pair-production fraction is therefore $R = N(\bar{\Omega}^+)/N(\Omega^-) = 0.55$, so that the excess (transported) component accounts for $1 - R \approx 45\%$ of the Ω^- sample, corresponding to $N_t = N(\Omega^-) - N(\bar{\Omega}^+) \approx 3.2 \times 10^4$ excess Ω^-. The clear surplus of Ω^- over $\bar{\Omega}^+$ is the direct manifestation of net baryon-number transport into the multi-strange sector at this energy.

6.3 Femtoscopic Correlation Functions

The four charge combinations are first examined at the femtoscopic level, before the CBS decomposition. Figure 6.3 shows the standard correlation functions $C_X(k^*) = S_X/M_X$ (Sect. 5.4), each formed as the ratio of the purity-corrected same-event distribution to the mixed-event reference and normalized to unity in the band $k^* \in [0.8, 1.5]$ GeV$/c$, above the range of the final-state and strangeness correlations. The same-sign pairs $K^-\Omega^-$ and $K^+\bar{\Omega}^+$ fall below unity in the lowest k^* bins, consistent with like-sign Coulomb repulsion together with the same-strangeness anti-correlation of a shared valence s quark, while the opposite-sign pairs $K^+\Omega^-$ and $K^-\bar{\Omega}^+$ track unity across the measured range at the present statistics. Each pair agrees with its charge conjugate, as pair-production symmetry requires.

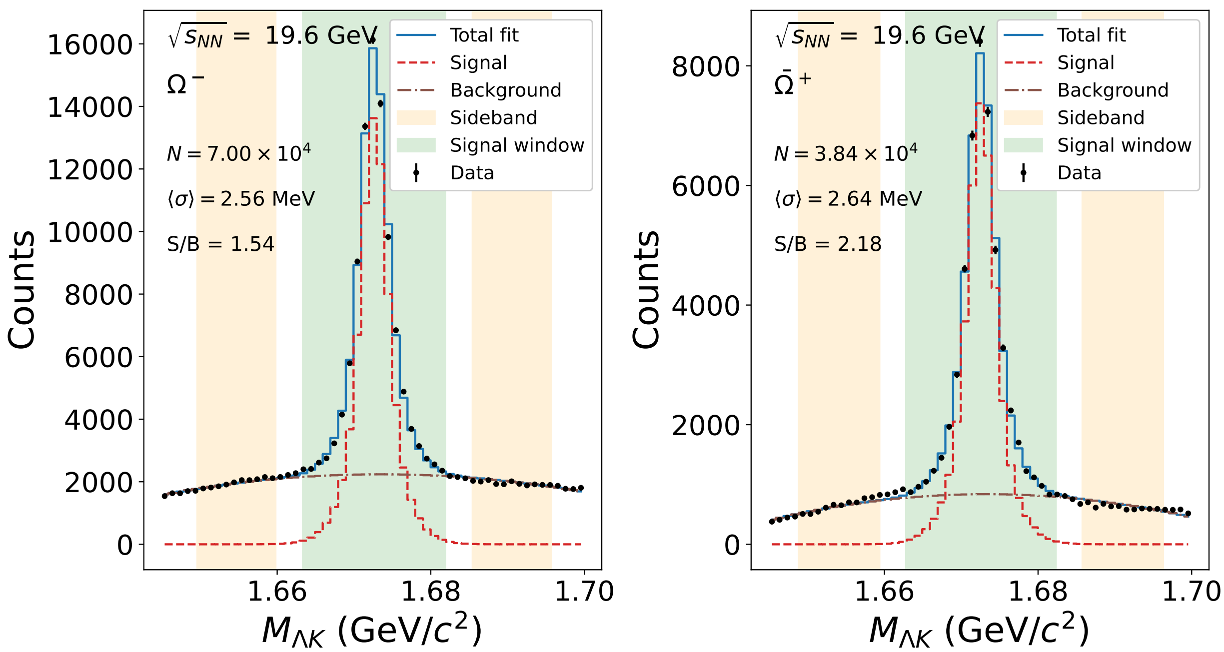

Figure 6.1: Acceptance-integrated $K\Lambda$ invariant-mass spectra for Ω^- (left) and $\bar{\Omega}^+$ (right) at $\sqrt{s_{NN}} = 19.6$ GeV. The total fit (blue solid) is the sum of the per-(p_T, y)-slice fits, each a double-Gaussian signal (red dashed) on a quadratic background (brown dash-dotted). Shaded regions mark the signal window (containing 99.7% of the fitted signal) and the sidebands used for the purity correction.

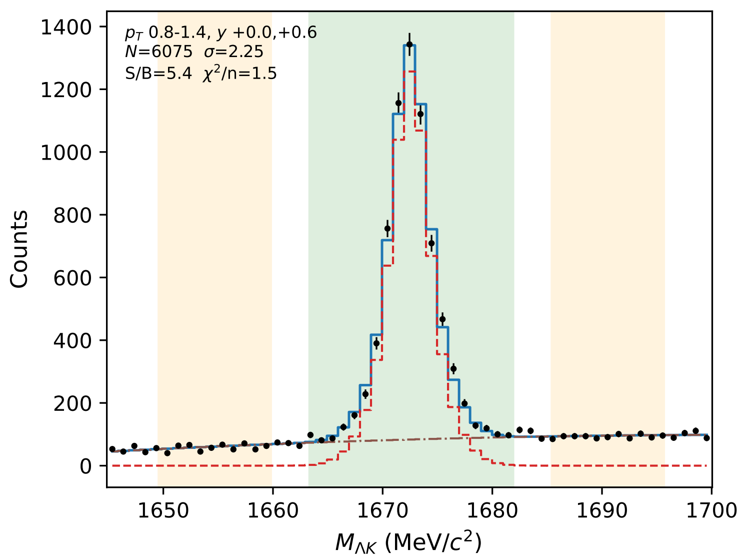

Figure 6.2: Representative single-slice $K\Lambda$ mass fit for Ω^- at $\sqrt{s_{NN}} = 19.6$ GeV, for $p_T \in [0.8, 1.4]$ GeV/c and $y \in [0.0, 0.6]$. Data (black), total fit (blue solid), double-Gaussian signal (red dashed), and quadratic background (brown dash-dotted) are shown; the green and orange bands mark the signal and sideband windows. The fitted yield, effective resolution σ, signal-to-background ratio, and χ^2/ndf are quoted in the panel. The full grid of per-slice fits is given in Appendix A.

The two femtoscopically relevant pairs, $K^+\Omega^-$ and $K^-\bar{\Omega}^+$, are charge conjugates carrying the same $s\bar{s}$ correlation, so a mixed-event-divided observable is dominated by the common pair-production correlation and returns a similar shape for Ω^- and $\bar{\Omega}^+$. The observation that the two shapes indeed agree affirms the need for an observable that preserves the pair yield rather than dividing it out. The baryon-transport asymmetry between the two is carried by the additional yield of excess-associated kaons rather than by the correlation shape, and so remains hidden in a normalized C; recovering it is the purpose of the yield-preserving C_{CBS} of Sect. 5.4.5.

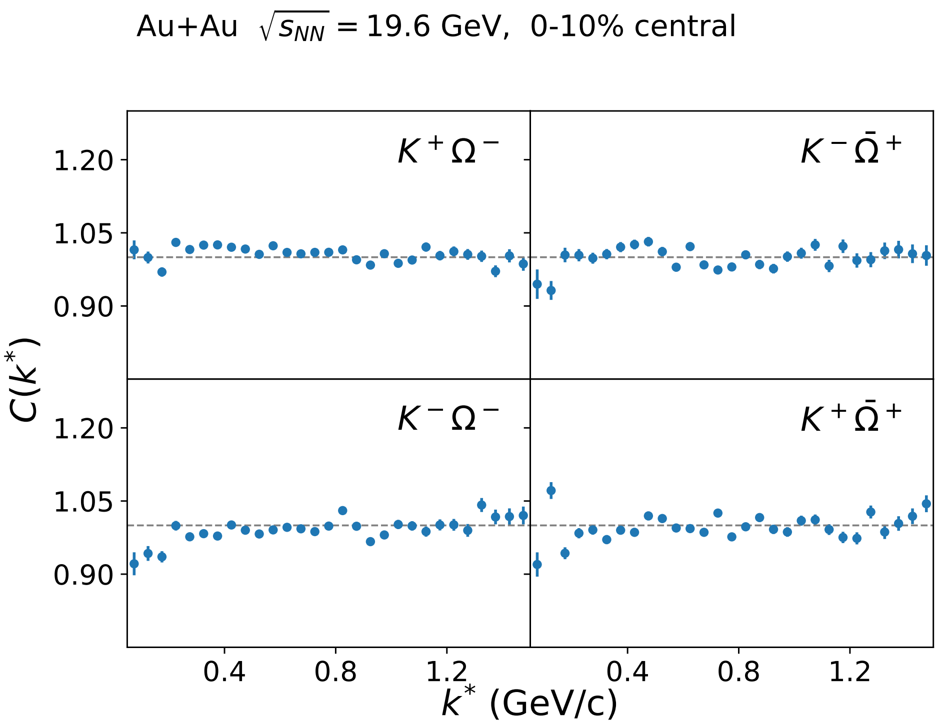

Figure 6.3: Femtoscopic correlation functions $C_X(k^*) = S_X/M_X$ for the four kaon–Ω charge combinations in 0–10% central Au+Au collisions at $\sqrt{s_{NN}} = 19.6$ GeV, normalized to unity in $k^* \in [0.8, 1.5]$ GeV/c. Top row: opposite-sign pairs; bottom row: same-sign pairs. The lowest k^* bin ($k^* < 0.05$ GeV/c), where the pair statistics are sparse and the Coulomb factor diverges, is omitted. The dashed line marks $C = 1$.

6.4 First-Order Correlation Difference

Before the full CBS decomposition, the transported excess can be exposed with the simpler per-Ω correlation difference $D_{\pm}$ defined in Sect. 5.4.2 (Eq. 5.4), in which each purity-corrected same-event combination is normalized to its parent-hyperon yield and the same-strangeness-sign combination is subtracted as a baseline from the opposite-strangeness-sign one. Figure 6.4 shows D_+ and D_- at $\sqrt{s_{NN}} = 19.6$ GeV as functions of the rapidity difference Δy, the difference of absolute rapidities $\Delta|y|$, and the relative momentum k^*, with adjacent bins merged to suppress combinatorial fluctuations.

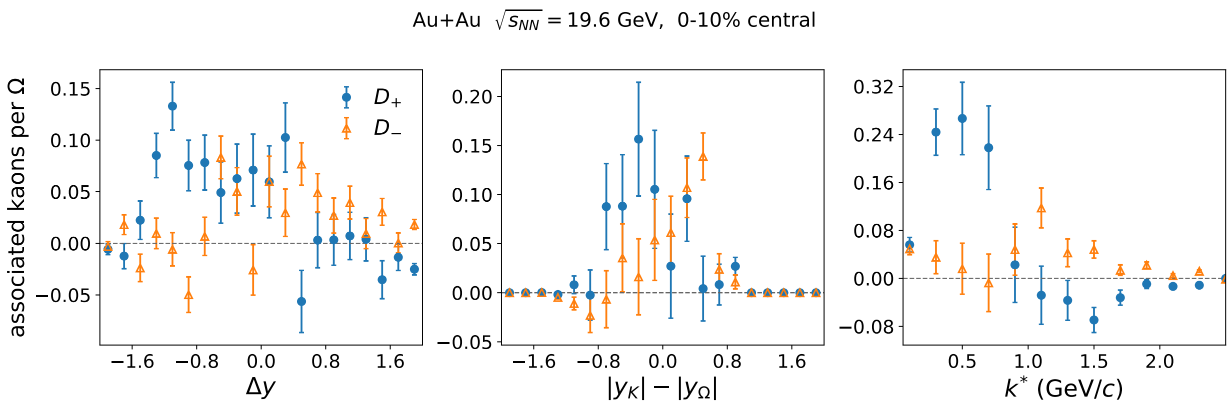

Figure 6.4: First-order per-Ω correlation difference at $\sqrt{s_{NN}} = 19.6$ GeV: D_+ (blue filled circles), the mean associated K^+ per Ω^- minus per $\bar{\Omega}^+$, and D_- (orange open triangles), the charge-conjugate K^- difference (Eq. 5.4), as functions of Δy, $\Delta|y|$, and k^*. Adjacent bins are merged by a factor of four; the lowest k^* bin ($k^* < 0.05$ GeV$/c$) is omitted. The dashed line marks zero.

The two differences are read through the excess decomposition of Eq. 5.1. The Ω^- sample is a pair-produced component (fraction R, identical in every respect to $\bar{\Omega}^+$) together with a transported excess (fraction $1 - R$), and strangeness conservation requires the excess to be accompanied by additional associated K^+ and K^0. Because $\bar{\Omega}^+$ is purely pair-produced, subtracting its K^+ association removes the pair-production baseline and leaves the excess: D_+ measures the extra strangeness-balancing kaons associated with the transported Ω^-, whereas D_- is the charge-conjugate construction on the pair-produced $\bar{\Omega}^+$ side and serves as a null reference. Integrated over each variable, D_+ lies systematically above D_-, about 0.60 against 0.40 associated kaons per Ω, and the surplus is concentrated at small k^*, where the balancing partners are produced close in phase space. This excess of K^+ around the Ω^- relative to the $\bar{\Omega}^+$ reference is consistent with the transported-strangeness signature that the CBS excess signal $\hat{a}$ isolates.

The first-order difference is nonetheless a blunt instrument. It carries no mixed-event baseline: the $\bar{\Omega}^+$ subtraction cancels the bulk of the combinatorial K^+ association but not its dependence on pair kinematics, so residual acceptance mismatch survives as bin-to-bin scatter and even drives D_+ below zero at intermediate k^*, where a genuine associated yield cannot be negative. The signal is robust only in the integral and, differentially, in k^*; the rapidity projections remain noisy even after merging bins. The CBS decomposition of Sect. 5.4.5 is built to remedy exactly these shortcomings: it estimates the pair-production term e from mixed events and uses all four charge states rather than a single pairwise difference, turning the first-order difference into the baseline-corrected excess signal presented next.

6.5 Extracted CBS Correlations

Applying the CBS decomposition (Eqs. 5.14–5.15) to the four purity-corrected same-event correlations yields the excess signal $\hat{a}/(N_\Omega - N_{\bar{\Omega}})$ and the pair-production-like baseline $c'/N_{\bar{\Omega}}$, each normalized per relevant Ω, the two components of the CBS correlation C_{CBS} (Sect. 5.4.5). Figure 6.5 shows them as functions of two rapidity variables, both extracted with the full method of Chapter 5. The rapidity difference $\Delta y = y_K - y_\Omega$ measures the longitudinal separation of the pair; both $\hat{a}$ and c' peak at $\Delta y = 0$ and fall to zero by $|\Delta y| \approx 2$. The difference of absolute rapidities $\Delta|y| \equiv |y_K| - |y_\Omega|$ asks a different question, of whether the kaon sits farther from midrapidity than the Ω or nearer to it. A transported Ω carries baryon number from one of the colliding nuclei and keeps some of that beam's forward motion, so it lies off midrapidity, on the side it came from, at a finite $|y_\Omega|$. Its strangeness-balancing kaon may sit even farther out, between the Ω and that beam, in which case $|y_K| > |y_\Omega|$ and $\Delta|y| > 0$; or it may sit back toward midrapidity, with $|y_K| < |y_\Omega|$ and $\Delta|y| < 0$. Absolute rapidities are compared, rather than the kaon's position relative to a particular beam, because the two colliding nuclei are identical: for a single Ω one cannot tell which of the two beams supplied it, only how far it has ended up from the center. Distance from midrapidity is therefore what separates "toward the beam" from "toward the center" here. Both $\hat{a}$ and c' peak at $\Delta|y| = 0$ and vanish by $\Delta|y| \approx \pm 1$: within the current acceptance, the observed balancing kaons sit alongside the Ω in rapidity.

Across both variables the excess signal $\hat{a}$ lies above the baseline c' near the peak, where the balancing kaons cluster, and the two merge in the tails; this ordering is consistent with the expected transported-kaon signal. The baseline c'/N_p measures the ambient excess kaons an Ω sits among, those balancing the far more abundant transported Λ and Ξ, while $\hat{a}/N_t$ adds the excess kaons that balance the transported Ω^- itself. Because an Ω is rare, a pair-production-like Ω event carries essentially none of the latter, so the difference isolates the Ω-specific balancing (Eq. 5.17). Integrated over the correlation variable, $\hat{a}/N_t = 13.7$ and $c'/N_p = 12.6$ associated kaons per Ω, differing, for the Δy projection, by

$$\frac{\int \hat{a}\, dX}{N_t} - \frac{\int c'\, dX}{N_p} = 1.10 \pm 0.74\,(\text{stat}) \pm 0.01\,(\text{syst}), \tag{6.2}$$

a positive difference of about 9% of the common baseline. Binning the same pairs in $\Delta|y|$ gives 1.09 in place of 1.10, the small shift following from the different integration range. The measurement is limited by statistics rather than by systematics: the full set of selection and normalization variations (Sect. 5.4.6) together is negligible against the statistical uncertainty. The total is dominated by the mixed-event normalization band, while the vertex-position, track-quality, and mass-window variations are each consistent with the default within the statistical precision of the comparison. That statistical reach is in turn bounded by the kaon acceptance. The kaons that balance the strangeness of a transported Ω^- are not required to fall within the acceptance used here, so a part of the balancing yield may lie outside the measured window and go uncounted. Yet the positive difference indicates that the transported Ω^- tends to balance its strangeness more with kaons than an ordinary produced Ω does, as expected.

6.6 Summary

At $\sqrt{s_{NN}} = 19.6$ GeV the Ω^- sample carries a $\sim 45\%$ transported excess over $\bar{\Omega}^+$. The CBS decomposition, extracted with the full method of Chapter 5, separates the excess signal $\hat{a}$ from the ambient baseline c': integrated, $\hat{a}/N_t$ exceeds c'/N_p by 1.10 ± 0.74 (statistical), about 9%, a positive difference in the direction expected if baryon transport shifts the compensating anti-strangeness of the excess Ω^- toward kaons and away from co-produced anti-baryons. Selection and normalization systematics are negligible in comparison, so the measurement is limited by statistics, and in particular by the kaon acceptance: the kaons balancing the transported Ω^- need not fall within the measured window, so part of the associated yield may be missed. Nonetheless, the CBS approach provides a path to extract the strangeness balancing specific to the transported Ω^-, which can be used for future theoretical and model comparisons.

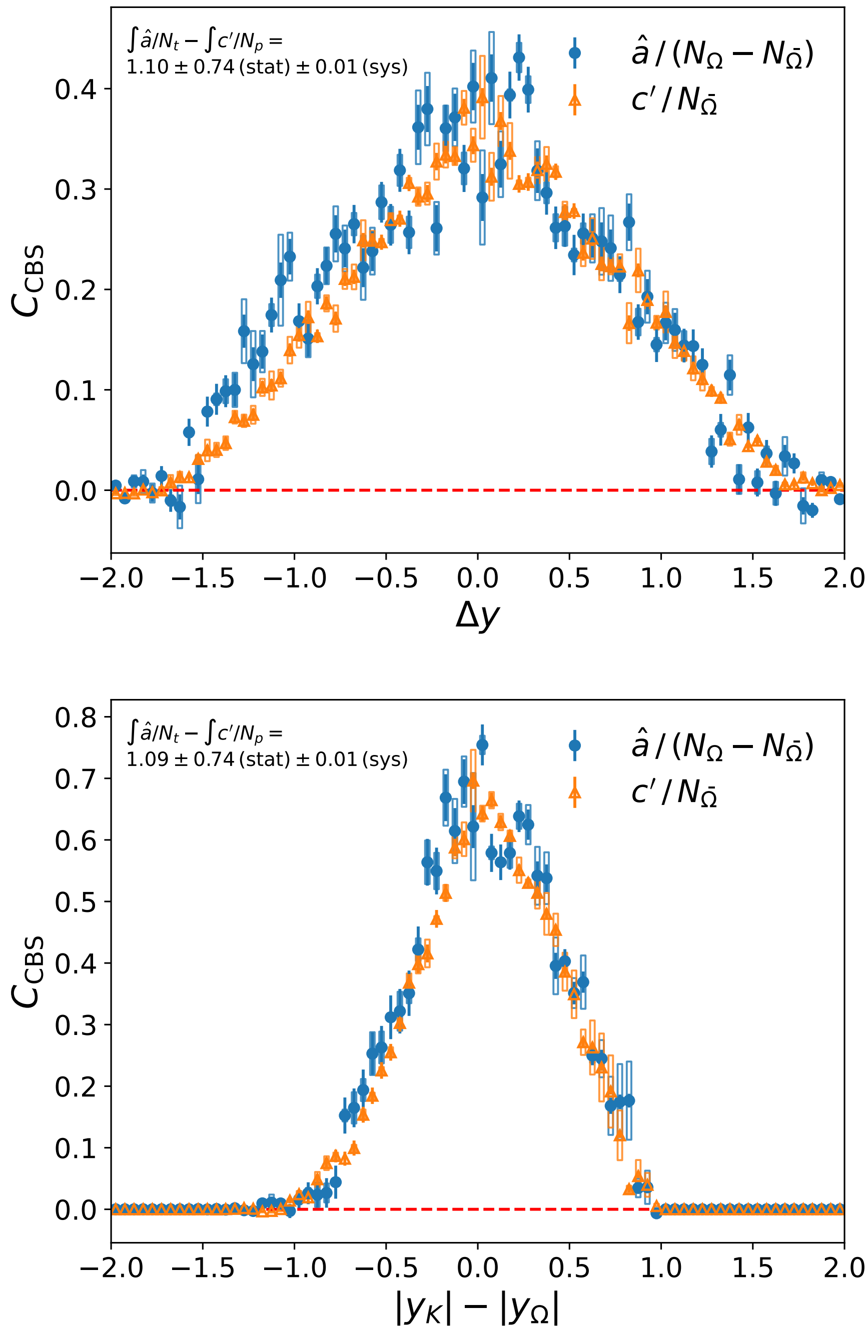

Figure 6.5: Full-method CBS correlation C_{CBS} at $\sqrt{s_{NN}} = 19.6$ GeV: excess signal $\hat{a}/(N_\Omega - N_{\bar{\Omega}})$ (blue filled circles) and pair-production-like baseline $c'/N_{\bar{\Omega}}$ (orange open triangles), as functions of the rapidity difference Δy (top) and the difference of absolute rapidities $\Delta|y| = |y_K| - |y_\Omega|$ (bottom). The extraction uses high-k^* mixed-event normalization and the C_e-modified e term. Error bars are statistical; open boxes are the systematic uncertainty from the variations of Table 5.5. Integrated, $\hat{a}/N_t$ exceeds c'/N_p by 1.10 ± 0.74 (statistical) $\pm\, 0.01$ (systematic), about 9%, comparable to the statistical uncertainty.

Chapter 7

Conclusion

This thesis has presented three measurements from the STAR BES-II program in Au+Au collisions at $\sqrt{s_{NN}}$ = 7.7–27 GeV. At these energies baryon stopping deposits a substantial fraction of the incoming baryon number at midrapidity, and the quarks transported in from the colliding nuclei, having experienced the full collective evolution of the medium, leave measurable signatures on the final-state hadrons. The first two analyses are closely connected: the elliptic flow of charged pions and the directed-flow splitting of Λ hyperons are both collective-flow observables, interpreted within the same quark-coalescence framework, that access these transported quarks through momentum-space anisotropies. The third, the kaon-associated correlations of the multi-strange Ω^-, stands on its own, using a correlation technique rather than flow to isolate the transported baryon population directly and to develop a methodology of broader applicability. This chapter draws together what they establish.

Transported-quark flavor from pion elliptic flow. The elliptic flow of charged pions provides a direct, flavor-sensitive measure of the transported light-quark content. Across the BES-II range, π^- was found to carry systematically larger v_2 than π^+, the natural consequence of the greater availability of transported d quarks than u quarks in the neutron-rich Au nucleus. Casting this difference into a coalescence sum rule, with the antiproton flow as a proxy for the produced-quark component, yields the transported-quark ratio $N_d^{\mathrm{tr}}/N_u^{\mathrm{tr}}$. The extracted ratio is consistent with the Au stoichiometric expectation $(2N+Z)/(N+2Z) = 315/276 \approx 1.14$ across the energies studied, and shows no significant dependence on collision energy, as expected, since the initial nuclear geometry that fixes the d/u ratio is independent of $\sqrt{s_{NN}}$. This constitutes a direct quantitative link between the isospin content of the colliding nuclei and the collective expansion of the matter they produce, and a clean confirmation that the constituent quarks of the final-state pions inherit the flavor composition of the transported population.

Quark-level electromagnetic response from Λ directed flow. The directed-flow splitting between Λ^0 and $\bar{\Lambda}^0$ opens a complementary window onto the early-time dynamics. The slope difference $\Delta(dv_1/dy)$ is positive in central collisions, following the baryon-transport hierarchy in which species richer in transported quarks acquire stronger directed flow, and changes sign to become negative toward peripheral collisions, where the intense electromagnetic fields of the spectator protons, through the Faraday and Coulomb effects, overcome the positive transport and Hall contributions. Interpreting the splitting through a coalescence relation, the proton splitting alone already reproduces the sign, centrality, and energy dependence of the Λ result, confirming baryon transport as the dominant driver; adding the kaon term to complete the quark-content match improves the description at every centrality. The comparison indicates that the electromagnetic response acts at the level of the constituent quarks and is not set by electric charge alone, giving access to the field evolution at hadronization. Separating, at the quark level, the response of the transported quarks from that of the pair-produced ones, which enter the splitting together, calls for precision beyond what the present measurement reaches.

Baryon transport into the multi-strange sector from Ω correlations. Whereas the flow observables probe the transported population through momentum-space anisotropies, the Ω^- hyperon offers a more direct handle on the transported *baryon* itself. Carrying no valence light quarks, the Ω shares no valence flavor with the incoming nuclei, so any net baryon number in this sector must have been transported, and at $\sqrt{s_{NN}} = 19.6$ GeV the Ω^- yield shows a marked excess over $\bar{\Omega}^+$, of order half the sample. To characterize this excess, the Combinatorial Background Subtraction (CBS) method was established, using the four kaon–Ω charge combinations and the purely pair-produced $\bar{\Omega}^+$ reference to decompose the Ω^- correlations into a transported-excess component and a pair-production-like baseline at the ensemble level. The measured difference is positive, about 9% and comparable to the present statistical uncertainty: the transported Ω^- balances its strangeness somewhat more with kaons than an ordinary produced Ω does, in the direction expected if baryon transport carries the compensating anti-strangeness into kaons rather than into co-produced anti-baryons. Though the forward acceptance limits how much of the balancing yield is counted, the method delivers an ensemble-level measure of the kaon association of the transported Ω^-, a reference against which future theoretical work and model simulations can be compared directly. Carrying the decomposition to the other BES-II energies, where the transported fraction varies with μ_B, and to asymmetric collision systems, where the colliding nuclei deliver baryon number unevenly, is underway.

Appendix A

Ω Mass-Fit Quality Assurance

The Ω yields used in Chapter 6 are extracted differentially, by fitting the $K\Lambda$ invariant-mass spectrum in each (p_T, y) bin and summing the fitted signal yields (Sect. 6.2). Figures A.1 and A.2 document the quality of these per-slice fits for Ω^- and $\bar{\Omega}^+$ at $\sqrt{s_{NN}} = 19.6$ GeV. Each panel shows the invariant-mass distribution of one slice with the double-Gaussian signal (red) on a quadratic background (green) and their sum (blue), together with the fitted yield N, effective resolution σ, signal-to-background ratio, and χ^2/ndf. The shaded bands mark the signal window (containing 99.7% of the fitted signal, green) and the sidebands (orange) used in the purity correction. At midrapidity the peaks are sharp and sit on a flat background (effective width $\sigma \approx 2.5$ MeV$/c^2$, high S/B); toward forward rapidity the peaks broaden slightly and the background grows, but the fits remain stable across the acceptance.

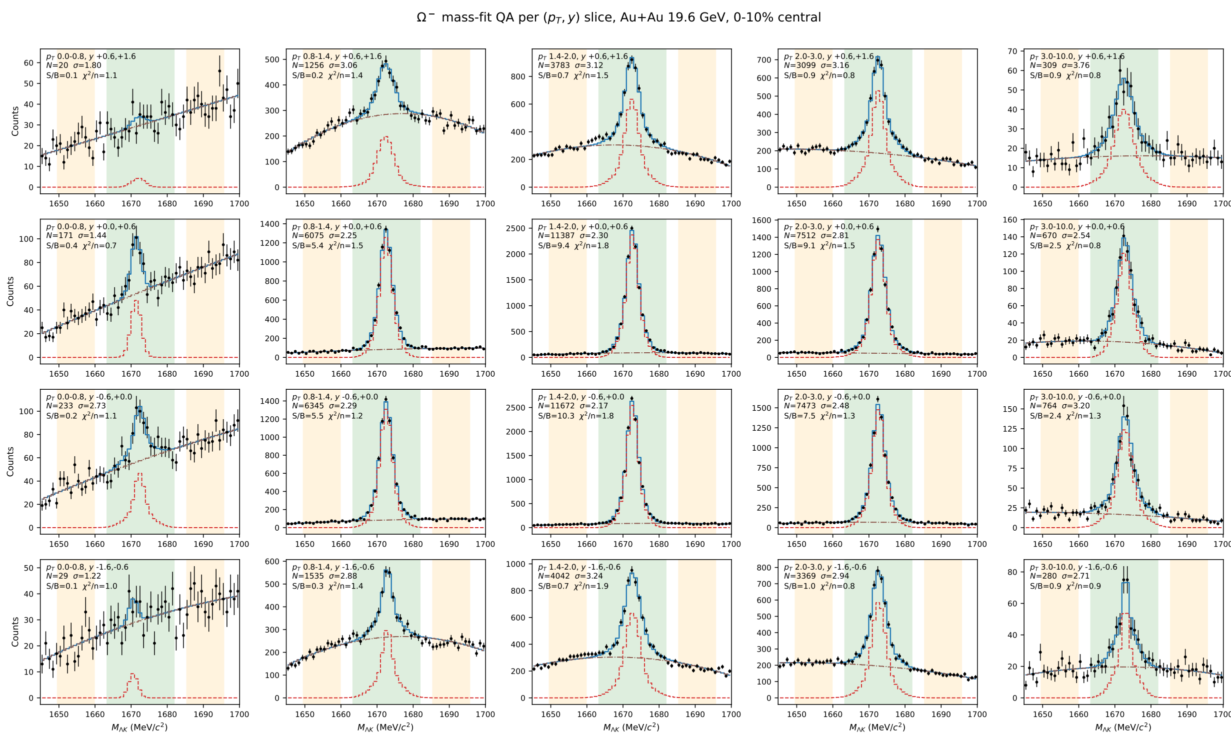

Figure A.1: Per-(p_T, y)-slice $K\Lambda$ mass fits for Ω^- at $\sqrt{s_{NN}} = 19.6$ GeV. Rows are rapidity intervals (forward at top and bottom, midrapidity in the middle), columns are p_T intervals. Data (black), total fit (blue solid), double-Gaussian signal (red dashed), quadratic background (brown dash-dotted); green and orange bands mark the signal and sideband windows.

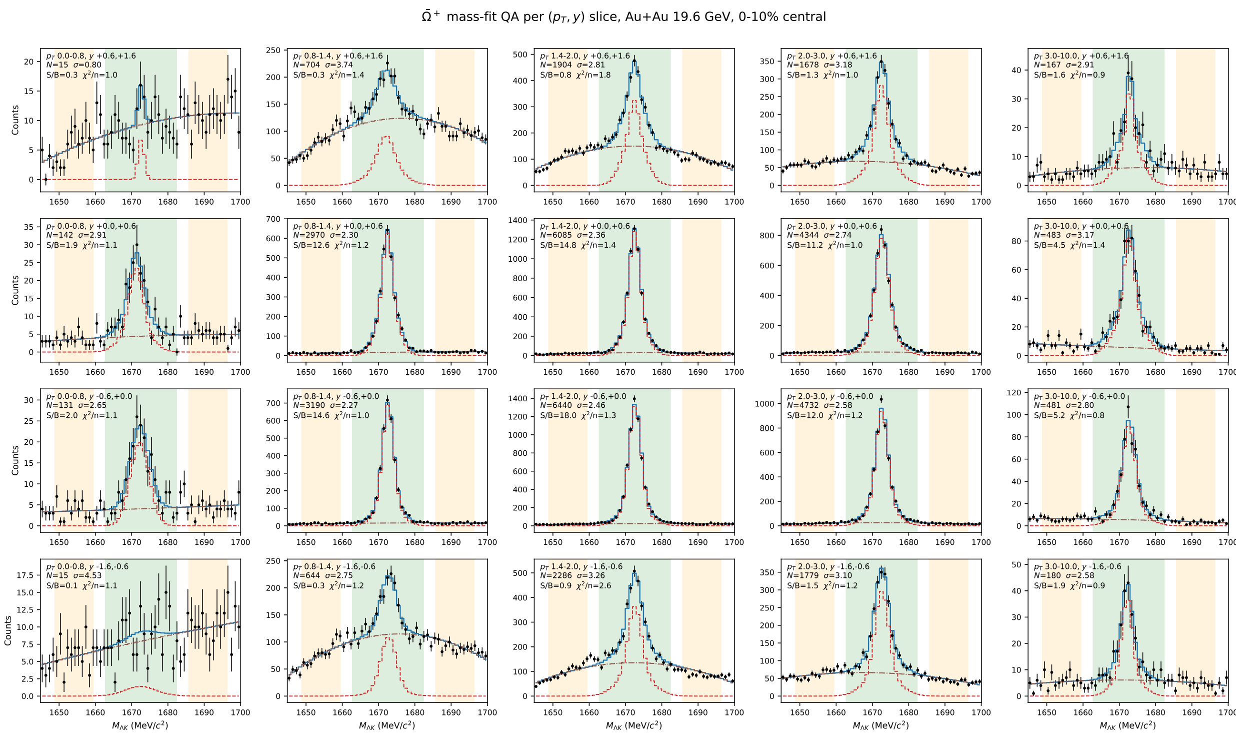

Figure A.2: Per-(p_T, y)-slice $K\Lambda$ mass fits for $\bar{\Omega}^+$ at $\sqrt{s_{NN}} = 19.6$ GeV, in the same layout as Figure A.1.

Bibliography

[1] M. I. Abdulhamid et al. "Observation of the Electromagnetic Field Effect via Charge-Dependent Directed Flow in Heavy-Ion Collisions at the Relativistic Heavy Ion Collider". In: *Phys. Rev. X* 14 (1 2024), p. 011028. DOI: 10.1103/PhysRevX.14.011028. arXiv: 2304.03430 [nucl-ex]. URL: https://link.aps.org/doi/10.1103/PhysRevX.14.011028.

[2] B. I. Abelev et al. "Centrality dependence of charged hadron and strange hadron elliptic flow from $\sqrt{s_{NN}} = 200$ GeV Au+Au collisions". In: *Phys. Rev. C* 77 (5 May 2008), p. 054901. DOI: 10.1103/PhysRevC.77.054901. URL: https://link.aps.org/doi/10.1103/PhysRevC.77.054901.

[3] J. Adam et al. "Strange hadron production in Au+Au collisions at $\sqrt{s_{NN}} = 7.7$, 11.5, 19.6, 27, and 39 GeV". In: *Phys. Rev. C* 102 (3 Sept. 2020), p. 034909. DOI: 10.1103/PhysRevC.102.034909. URL: https://link.aps.org/doi/10.1103/PhysRevC.102.034909.

[4] L. Adamczyk et al. "Observation of an Energy-Dependent Difference in Elliptic Flow between Particles and Antiparticles in Relativistic Heavy Ion Collisions". In: *Phys. Rev. Lett.* 110 (14 2013), p. 142301. DOI: 10.1103/PhysRevLett.110.142301. arXiv: 1301.2347 [nucl-ex]. URL: https://link.aps.org/doi/10.1103/PhysRevLett.110.142301.

[5] L. Adamczyk et al. "Observation of Charge Asymmetry Dependence of Pion Elliptic Flow and the Possible Chiral Magnetic Wave in Heavy-Ion Collisions". In: *Phys. Rev. Lett.* 114 (25 2015), p. 252302. DOI: 10.1103/PhysRevLett.114.252302. arXiv: 1504.02175 [nucl-ex]. URL: https://link.aps.org/doi/10.1103/PhysRevLett.114.252302.

[6] L. Adamczyk et al. "Bulk Properties of the Medium Produced in Relativistic Heavy-Ion Collisions from the Beam Energy Scan Program". In: *Phys. Rev. C* 96 (4 Oct. 2017), p. 044904. DOI: 10.1103/PhysRevC.96.044904. URL: https://link.aps.org/doi/10.1103/PhysRevC.96.044904.

[7] J. Adams et al. "The STAR event plane detector". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 968 (July 2020), p. 163970. ISSN: 0168-9002. DOI: 10.1016/j.nima.2020.163970. URL: http://dx.doi.org/10.1016/j.nima.2020.163970.

[8] AGS Booster Project. *Booster Design Manual*. Technical Report. Upton, NY: Brookhaven National Laboratory, Oct. 1988. URL: https://www.bnl.gov/cad/accelerator/docs/pdf/boosterdesignmanual.pdf.

[9] F.S. Bieser et al. "The STAR trigger". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 499.2–3 (2003), pp. 766–777. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(02)01974-5.

[10] K.H. Ackermann et al. "STAR detector overview". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 499.2 (2003), pp. 624–632. ISSN: 0168-9002. DOI: https://doi.org/10.1016/S0168-9002(02)01960-5. URL: https://www.sciencedirect.com/science/article/pii/S0168900202019605.

[11] J.G. Alessi et al. "Design of an EBIS for RHIC". In: *Proceedings of the 2003 Particle Accelerator Conference*. Vol. 1. 2003, 89–91 Vol.1. DOI: 10.1109/PAC.2003.1288848.

[12] M. Anderson et al. "The STAR time projection chamber: a unique tool for studying high multiplicity events at RHIC". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 499.2–3 (Mar. 2003), pp. 659–678. ISSN: 0168-9002. DOI: 10.1016/s0168-9002(02)01964-2. URL: http://dx.doi.org/10.1016/S0168-9002(02)01964-2.

[13] B. Andersson, G. Gustafson, and T. Sjöstrand. "A model for baryon production in quark and gluon jets". In: *Nucl. Phys. B* 197.1 (1982), pp. 45–54. DOI: 10.1016/0550-3213(82)90153-5.

[14] B. Andersson et al. "Parton fragmentation and string dynamics". In: *Phys. Rep.* 97.2 (1983), pp. 31–145. DOI: 10.1016/0370-1573(83)90080-7.

[15] Mauro Anselmino et al. "Diquarks". In: *Rev. Mod. Phys.* 65 (4 Oct. 1993), pp. 1199–1233. DOI: 10.1103/RevModPhys.65.1199. URL: https://link.aps.org/doi/10.1103/RevModPhys.65.1199.

[16] W.K. van Asselt et al. "The transition jump system for the AGS". In: *Proceedings Particle Accelerator Conference*. Vol. 5. 1995, 3022–3024 vol.5. DOI: 10.1109/PAC.1995.505771.

[17] Dennis Barak, Beau F. Harrison, and Adam Watts. *Concepts Rookie Book*. 2025. arXiv: 2510.20880 [physics.acc-ph]. URL: https://arxiv.org/abs/2510.20880.

[18] A. Bazavov et al. "Chiral crossover in QCD at zero and non-zero chemical potentials". In: *Physics Letters B* 795 (2019), pp. 15–21. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2019.05.013. URL: https://www.sciencedirect.com/science/article/pii/S0370269319303223.

[19] F. Bergsma et al. "The STAR detector magnet subsystem". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 499.2 (2003), pp. 633–639. ISSN: 0168-9002. DOI: https:

//doi.org/10.1016/S0168-9002(02)01961-7. URL: https://www.sciencedirect.com/science/article/pii/S0168900202019617.

[20] Hans Bichsel. “A method to improve tracking and particle identification in TPCs and silicon detectors”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 562.1 (2006), pp. 154–197. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2006.03.009. URL: https://www.sciencedirect.com/science/article/pii/S0168900206005353.

[21] Brookhaven National Laboratory. *Aerial view of the Relativistic Heavy Ion Collider (RHIC) complex*. https://www.bnl.gov/newsroom/news.php?a=221190. Accessed: Wednesday 19th August, 2026. 2023.

[22] Y. Burnier et al. “Chiral Magnetic Wave at Finite Baryon Density and the Electric Quadrupole Moment of the Quark-Gluon Plasma in Heavy Ion Collisions”. In: *Phys. Rev. Lett.* 107 (5 2011), p. 052303. DOI: 10.1103/PhysRevLett.107.052303. arXiv: 1103.1307 [hep-ph]. URL: https://link.aps.org/doi/10.1103/PhysRevLett.107.052303.

[23] Wei-Tian Deng and Xu-Guang Huang. “Event-by-event generation of electromagnetic fields in heavy-ion collisions”. In: *Phys. Rev. C* 85 (4 2012), p. 044907. DOI: 10.1103/PhysRevC.85.044907. arXiv: 1201.5108 [nucl-th]. URL: https://link.aps.org/doi/10.1103/PhysRevC.85.044907.

[24] Lipei Du, Agnieszka Sorensen, and Mikhail Stephanov. “The QCD phase diagram and Beam Energy Scan physics: A theory overview”. In: *International Journal of Modern Physics E* 33.07 (2024), p. 2430008. DOI: 10.1142/S021830132430008X. eprint: https://doi.org/10.1142/S021830132430008X. URL: https://doi.org/10.1142/S021830132430008X.

[25] Alexei Fedotov et al. "Operational Electron Cooling in the Relativistic Heavy Ion Collider". In: *12th International Particle Accelerator Conference* . Aug. 2021. DOI: 10.18429/JACoW-IPAC2021-WEXA02.

[26] R. J. Fries, V. Greco, and P. Sorensen. "Coalescence models for hadron formation from quark-gluon plasma". In: *Annu. Rev. Nucl. Part. Sci.* 58 (2008), pp. 177–205. DOI: 10.1146/annurev.nucl.58.110707.171134.

[27] R. J. Fries et al. "Hadronization in Heavy-Ion Collisions: Recombination and Fragmentation of Partons". In: *Phys. Rev. Lett.* 90 (20 May 2003), p. 202303. DOI: 10.1103/PhysRevLett.90.202303. URL: https://link.aps.org/doi/10.1103/PhysRevLett.90.202303.

[28] Chris Gardner et al. "Operation of the RHIC Injector Chain with Ions from EBIS". In: *6th International Particle Accelerator Conference*. 2015, THPF046. DOI: 10.18429/JACoW-IPAC2015-THPF046.

[29] Christopher J. Gardner. *FY2020-21 parameters for Gold ions in Booster, AGS, and RHIC*. Tech. rep. Brookhaven National Laboratory (BNL), Upton, NY (United States), Feb. 2021. DOI: 10.2172/1765655. URL: https://www.osti.gov/biblio/1765655.

[30] S. Gorbunov and I. Kisel. "Reconstruction of decayed particles based on the Kalman filter". In: CBM-SOFT-note-2007-003 (2007). URL: https://www.gsi.de/documents/DOC-2007-May-14-1.pdf.

[31] V. Greco, C. M. Ko, and P. Lévai. "Parton Coalescence and the Antiproton/Pion Anomaly at RHIC". In: *Phys. Rev. Lett.* 90 (20 May 2003), p. 202302. DOI: 10.1103/PhysRevLett.90.202302. URL: https://link.aps.org/doi/10.1103/PhysRevLett.90.202302.

[32] Umut Gürsoy, Dmitri Kharzeev, and Krishna Rajagopal. "Magnetohydrodynamics, charged currents, and directed flow in heavy ion collisions". In: *Phys. Rev. C* 89 (5

2014), p. 054905. DOI: 10.1103/PhysRevC.89.054905. arXiv: 1401.3805 [hep-ph]. URL: https://link.aps.org/doi/10.1103/PhysRevC.89.054905.

[33] Umut Gürsoy et al. “Charge-dependent flow induced by magnetic and electric fields in heavy ion collisions”. In: *Phys. Rev. C* 98 (5 2018), p. 055201. DOI: 10.1103/PhysRevC.98.055201. arXiv: 1806.05288 [hep-ph]. URL: https://link.aps.org/doi/10.1103/PhysRevC.98.055201.

[34] R. E. Kalman. “A New Approach to Linear Filtering and Prediction Problems”. In: *J. Basic Eng.* 82.1 (1960), pp. 35–45. DOI: 10.1115/1.3662552.

[35] D. Kharzeev. “Can gluons trace baryon number?” In: *Physics Letters B* 378.1 (1996), pp. 238–246. ISSN: 0370-2693. DOI: https://doi.org/10.1016/0370-2693(96)00435-2. URL: https://www.sciencedirect.com/science/article/pii/0370269396004352.

[36] P. F. Kolb et al. “Momentum anisotropies in the quark coalescence model”. In: *Phys. Rev. C* 69 (5 May 2004), p. 051901. DOI: 10.1103/PhysRevC.69.051901.

[37] P. K. Kovtun, D. T. Son, and A. O. Starinets. “Viscosity in Strongly Interacting Quantum Field Theories from Black Hole Physics”. In: *Phys. Rev. Lett.* 94 (11 Mar. 2005), p. 111601. DOI: 10.1103/PhysRevLett.94.111601. URL: https://link.aps.org/doi/10.1103/PhysRevLett.94.111601.

[38] N. Lewis et al. “Search for baryon junctions in photonuclear processes and isobar collisions at RHIC”. In: *Eur. Phys. J. C* 84 (2024), p. 590. DOI: 10.1140/epjc/s10052-024-12834-2. arXiv: 2205.05685 [nucl-ex].

[39] W.J Llope et al. “The TOFp/pVPD time-of-flight system for STAR”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 522.3 (2004), pp. 252–273. ISSN: 0168-9002. DOI: https://doi.org/10.1016/j.nima.2003.11.414. URL: https://www.sciencedirect.com/science/article/pii/S016890020303393X.

[40] W.J. Llope. "The large-area time-of-flight upgrade for STAR". In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms* 241.1 (2005), pp. 306–310. ISSN: 0168-583X. DOI: https://doi.org/10.1016/j.nimb.2005.07.089. URL: https://www.sciencedirect.com/science/article/pii/S0168583X05012267.

[41] Brennan J M, M Blaskiewicz, and K Mernick. "Stochastic cooling in RHIC". In: Brookhaven National Laboratory (BNL) Relativistic Heavy Ion Collider. May 2012. URL: https://www.osti.gov/biblio/1047664.

[42] Michael L. Miller et al. "Glauber Modeling in High-Energy Nuclear Collisions". In: *Annual Review of Nuclear and Particle Science* 57.Volume 57, 2007 (2007), pp. 205–243. ISSN: 1545-4134. DOI: https://doi.org/10.1146/annurev.nucl.57.090506.123020. URL: https://www.annualreviews.org/content/journals/10.1146/annurev.nucl.57.090506.123020.

[43] P. Möller et al. "Nuclear ground-state masses and deformations: FRDM(2012)". In: *Atomic Data and Nuclear Data Tables* 109-110 (2016), pp. 1–204. ISSN: 0092-640X. DOI: https://doi.org/10.1016/j.adt.2015.10.002. URL: https://www.sciencedirect.com/science/article/pii/S0092640X1600005X.

[44] S. Navas et al. "Review of Particle Physics". In: *Phys. Rev. D* 110 (3 2024), p. 030001. DOI: 10.1103/PhysRevD.110.030001.

[45] Jean-Yves Ollitrault. "Anisotropy as a signature of transverse collective flow". In: *Phys. Rev. D* 46 (1 July 1992), pp. 229–245. DOI: 10.1103/PhysRevD.46.229. URL: https://link.aps.org/doi/10.1103/PhysRevD.46.229.

[46] Satoshi Ozaki and T. Roser. "Relativistic Heavy Ion Collider, its construction and upgrade". In: *Progress of Theoretical and Experimental Physics* 2015.3 (2015), 03A102. DOI: 10.1093/ptep/ptu093.

[47] A. M. Poskanzer and S. A. Voloshin. "Methods for analyzing anisotropic flow in relativistic nuclear collisions". In: *Phys. Rev. C* 58 (3 Sept. 1998), pp. 1671–1678. DOI: 10.1103/PhysRevC.58.1671. URL: https://link.aps.org/doi/10.1103/PhysRevC.58.1671.

[48] Johann Rafelski and Berndt Müller. "Strangeness Production in the Quark-Gluon Plasma". In: *Phys. Rev. Lett.* 48 (16 Apr. 1982), pp. 1066–1069. DOI: 10.1103/PhysRevLett.48.1066. URL: https://link.aps.org/doi/10.1103/PhysRevLett.48.1066.

[49] STAR Collaboration. *Technical Design Report for the iTPC Upgrade*. Tech. rep. STAR Note 0644. Brookhaven National Laboratory, 2015. URL: https://drupal.star.bnl.gov/STAR/system/files/TDRiTPCUpgrade_final.pdf.

[50] M Thiel et al. "Neutron skins of atomic nuclei: per aspera ad astra". In: *Journal of Physics G: Nuclear and Particle Physics* 46.9 (Aug. 2019), p. 093003. DOI: 10.1088/1361-6471/ab2c6d. URL: https://doi.org/10.1088/1361-6471/ab2c6d.

[51] Sergei A. Voloshin, Arthur M. Poskanzer, and Raimond Snellings. "Collective phenomena in non-central nuclear collisions". In: *arXiv:0809.2949* (2008). Published in: Relativistic Heavy Ion Physics, Landolt-Börnstein, Vol. I/23 (Springer, 2010). DOI: 10.1007/978-3-642-01539-7_10.

[52] Xiatong Wu et al. "Baryon Number Transport, Strangeness Conservation and Ω-hadron Correlations". In: *EPJ Web of Conferences*. Vol. 276. 2023, p. 03002. DOI: 10.1051/epjconf/202327603002. arXiv: 2212.10568 [nucl-th].

[53] Zhiwan Xu. "Search for the Chiral Magnetic Effect from RHIC Beam Energy Scan II data with STAR". PhD thesis. University of California, Los Angeles, 2024.

[54] C. Yang and Q. Yang. "The STAR Detector Upgrades for the BES-II and at Forward Rapidity". In: *Journal of Instrumentation* 15.07 (2020), p. C07040. DOI: 10.1088/1748-0221/15/07/C07040.

[55] Chi Yang. “The STAR beam energy scan phase II physics and upgrades”. In: *Nuclear Physics A* 967 (2017), pp. 800–803. ISSN: 0375-9474. DOI: https://doi.org/10.1016/j.nuclphysa.2017.05.042. URL: https://www.sciencedirect.com/science/article/pii/S0375947417301598.