

© 2016 by Randall Evan McClellan. All rights reserved.

ANGULAR DISTRIBUTIONS OF HIGH-MASS DILEPTON PRODUCTION IN
HADRON COLLISIONS

BY

RANDALL EVAN MCCLELLAN

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Physics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2016

Urbana, Illinois

Doctoral Committee:

Professor Matthias Grosse-Perdekamp, Chair
Professor Jen-Chieh Peng, Director of Research
Professor Scott Willenbrock
Professor James Eckstein

Abstract

The SeaQuest experiment is a fixed-target dimuon experiment currently running at the Fermi National Accelerator Laboratory (FNAL). By utilizing the high-intensity, 120 GeV proton beam delivered by the FNAL Main Injector (MI), SeaQuest is able to measure proton-induced Drell-Yan dimuon production off of various nuclear targets in kinematic regions inaccessible to previous similar experiments. A suitably large fraction of the final dataset has been recorded, reconstructed, and analyzed. Very preliminary results from light-sea flavor asymmetry, nuclear dependence, and partonic energy loss analyses have been presented at numerous international conferences.

A novel, FPGA-based trigger system has been designed, implemented, and optimized for the SeaQuest experiment. By implementing the trigger decision logic in FPGA firmware, it is more adaptable to changing experimental conditions. Additionally, the peripheral tasks of timing alignment, “trigger matrix” generation, and firmware uploading have been mostly automated, reducing the likelihood of user error in the maintenance and operation of the trigger system. Significant upgrades to hardware and firmware have greatly improved the performance of the trigger system since the 2012 commissioning run of SeaQuest. Four additional v1495 modules were added to facilitate thorough pulser testing of the firmware designs and *in-situ* pulser tests of all compiled firmware. These pulser tests proved crucial for diagnosing many errors that may have otherwise gone unnoticed. A significant change to the internal clocking of the trigger system eliminated a subtle source of rate-dependent trigger efficiency. With this upgrade, the trigger finally meets the “dead-time free” design specification.

Drell-Yan dimuon data have been collected and analyzed for central θ_{CS} , with nearly flat acceptance in ϕ_{CS} , in the mass range $5.0 \text{ GeV} < M_{\gamma^*} < 10.0 \text{ GeV}$ at forward x_F with the SeaQuest spectrometer at FNAL. A very preliminary extraction of λ has been performed, and the remaining difficulties in extracting ν have been evaluated. Although the results are not yet publishable, significant progress has been made in developing this very challenging angular distributions analysis. A simple scheme for correcting for the angular acceptances of the spectrometer, trigger, and reconstruction has been developed and demonstrated. A generally applicable correction for the kinematically-dependent, rate-dependent reconstruction efficiency

has been developed and applied to all current analyses on SeaQuest data. This rate-dependence correction was the first major hurdle in the path to publication of many preliminary SeaQuest results. The last remaining major correction for all analyses, but especially important for the angular parameter extraction, is the full characterization, rate-dependence correction, and subtraction of the combinatoric background contribution to the reconstructed dimuon sample.

Independently, an intuitive, kinematic derivation of the single-event definitions of the Drell-Yan angular parameters has been developed under the assumption of unpolarized annihilating quarks within unpolarized nuclei. At $\mathcal{O}(\alpha_s)$, where the quarks remain co-planar with the hadrons in the photon rest frame, this kinematic method reproduces the Lam-Tung relation and derives an additional equality for μ^2 , which is only interpretable for single-event parameters. This method has been extended to the case of quark non-coplanarity, and the coplanar equalities become inequalities. A new equality was discovered, which should be obeyed by single-event parameters even in the case of a non-coplanar quark axis. The non-coplanar parameter relations have been used to derive constraints on the experimentally accessible values of λ and ν . These constraints are compared with existing data and have been found consistent, except in the cases where significant contributions from non-zero Boer-Mulders functions are expected. Finally, the kinematically-derived parameter definitions have been applied to high-precision CMS data. The relative contributions of the $q\bar{q}$ and qg processes to the Z-boson ‘‘Drell-Yan’’ cross-section have been extracted. Further, an average measure of non-coplanarity, likely caused by $\mathcal{O}(\alpha_s^2)$ and higher processes, has been extracted.

Table of Contents

List of Tables	vii
List of Figures	viii
Chapter 1 Introduction	1
1.1 Background	2
1.1.1 Quantum Chromodynamics	4
1.1.2 Parton Distribution Functions	5
1.2 The Drell-Yan Process	7
1.2.1 Kinematics	7
1.2.2 Dynamics	8
1.2.3 Universality	10
1.2.4 QCD-Improved Drell-Yan	10
1.3 The Nucleon Light-Flavor Sea Asymmetry	12
1.4 Drell-Yan Angular Distributions	13
1.4.1 Photon Rest Frames	14
1.4.2 General Form of Drell-Yan Angular Distributions	15
1.4.3 The Lam-Tung Relation	15
1.4.4 The Boer-Mulders Function	16
1.4.5 The Sivers Function	18
Chapter 2 The SeaQuest Experiment	20
2.1 Overview	20
2.2 The Proton Beam	21
2.2.1 The Main Injector	21
2.2.2 Slow-Spill Extraction	23
2.2.3 Beam-Intensity Monitor	23
2.3 Targets	26
2.4 Magnets	26
2.5 Tracking Detectors	28
2.5.1 Hodoscopes	29
2.5.2 Wire Chambers	30
2.5.3 Proportional Tubes	32
2.6 Trigger	33
2.7 Data Acquisition	33
2.7.1 Main DAQ	33
2.7.2 Scaler DAQ	34
2.7.3 Beam DAQ	35
2.7.4 Slow Control	35
2.8 Data Curation	36

Chapter 3	The SeaQuest Trigger System	39
3.1	Overview	39
3.2	Initial Design and Implementation	40
3.2.1	Overall Structure	40
3.2.2	TDC Block	43
3.2.3	Delay Adjustment Pipeline	43
3.2.4	Trigger Matrix	44
3.3	Trigger System Upgrades	46
3.3.1	Level 0 v1495 Modules	46
3.3.2	Trigger System Testing	48
3.3.3	Synchronization with the MI Beam RF-Clock	50
3.4	Peripheral Software	51
3.4.1	Timing Adjustment	52
3.4.2	Trigger Matrix Generation	53
3.5	Performance of the Trigger System	54
3.5.1	Signal Acceptance and Background Rejection	54
3.5.2	Input/Output Consistency	57
3.6	Issues in the Trigger System	58
3.7	Future Upgrades	59
Chapter 4	Data Analysis	61
4.1	Track Reconstruction	61
4.1.1	Pre-Tracking Cuts	61
4.1.2	Track Finding	64
4.1.3	Track Fitting	66
4.1.4	Vertex Fitting	68
4.2	Rate Dependence	69
4.2.1	Rate-Dependent Reconstruction Efficiency	69
4.2.2	Chamber and Hodoscope Rate-Dependent Efficiencies	72
4.3	Drell-Yan Dimuon Selection	83
4.3.1	Quality Cuts	83
4.3.2	Mass Spectrum Fitting	85
4.3.3	Mass Cut	87
4.4	Acceptance Correction	87
4.4.1	Geant-based Monte Carlo (GMC)	88
4.4.2	Bin-level, Dimuon-based Acceptance Correction	91
4.5	Collins-Soper θ Distribution Results	91
4.5.1	Corrected θ Distribution	91
4.5.2	Extracted λ	94
4.6	ϕ Distribution	96
4.7	Future Improvements	99
4.7.1	Rate-Dependent Reconstruction Efficiency for Mix Background	99
Chapter 5	Phenomenology of Drell-Yan Angular Distributions	104
5.1	Interpretation of CMS Z -boson ‘Drell-Yan’ Results	104
5.1.1	Kinematic Derivation of Angular Distribution	104
5.1.2	Deriving Angular Parameter p_T -dependence	106
5.1.3	$q\bar{q}$ and qg Contributions and Quark Axis Non-Coplanarity from the CMS Data	108
5.2	Derivation of Parameter Relations	112
5.2.1	Deriving the Lam-Tung Relation	112
5.2.2	Non-co-planar Quark Axis	113
5.2.3	Non-Transverse Photon	114
5.2.4	Fully General	114
5.3	Comparison with Data	114

5.3.1	Constraints on Parameters	115
5.3.2	Experimental Measurements	118
Chapter 6	Conclusions	121
References	124

List of Tables

2.1	The seven SeaQuest targets. “Spill/Cycle” is a typical configuration. The exact spill cycle varied somewhat throughout SeaQuest’s data-taking periods.	26
2.2	Hodoscope plane specifications.	29
2.3	Wire Chamber specifications. The “primed” and unprimed planes in each view are slightly separated in z-position. U- and V-view z-positions are given relative to the associated X-view.	31
2.4	Run history and installed chamber combinations.	31
2.5	Proportional tube detector plane specifications.	32
3.1	The five triggers output by the v1495 trigger system. Matrix 2 was turned off (with a high prescale factor) because it was overwhelmingly dominated by background. Matrix 3 is used to estimate the rate of combinatoric backgrounds in the Drell-Yan sample.	46
3.2	Trigger roadset signal acceptance and “Random RF” background acceptance estimation. Signal acceptance is calculated from DY GMC and represents the Matrix 1 acceptance vs a hypothetical “all roads” trigger. The DY analysis mass cut, $Mass > 4.2 \text{ GeV}$, is applied to the GMC. Background acceptance is also Matrix 1 relative to ‘all roads’, but using Random RF data from Run 3 (Run 4). The “Signal Amplification” is the Signal Acceptance divided by the Background Acceptance.	57
4.1	Analyzed Datasets	62
4.2	Dimuon Counts	62
4.3	Multiplicity Cuts	64
4.4	Rate-Dependent Reconstruction Efficiency Parameters, “a”, θ -bins	72
4.5	Rate-Dependent Reconstruction Efficiency Parameters, “a”, ϕ -bins	82
4.6	Spill-Level Cuts	84
4.7	Bad Spill Ranges	84
4.8	Track Cuts	84
4.9	Dimuon Cuts	85
4.10	Target/Dump Selection	85
4.11	Extracted λ_{CS} results from the four analyzed datasets.	95
4.12	Extracted α results from the four analyzed datasets. α is qualitatively related to ν_{CS}	98

List of Figures

1.1	Leading Order DIS Process Diagram.	2
1.2	Left: Meson Octet. Right: Baryon Octet. Vertical position indicates strange quark content.	3
1.3	Left: Charge screening in QED. Center: Color screening in QCD. Right: Color “camouflage” in QCD. The camouflage effect dominates the color screening effect, and quarks within hadrons become asymptotically free in the high- Q^2 limit.[14]	4
1.4	Leading Order Drell-Yan Process Diagram.	8
1.5	$\mathcal{O}(\alpha_s)$ diagrams for QCD improved Drell-Yan. (a) One-loop corrections to the LO diagram. Only the interference terms contribute at $\mathcal{O}(\alpha_s)$. (b) Gluon emission diagrams, $q\bar{q}$ -diagrams (c) QCD “bremsstrahlung” diagrams, qg -process [3].	11
1.6	Experimental Results for $\langle p_T \rangle$ vs \sqrt{s} for proton- and pion-induced Drell-Yan. The linear dependence is due to QCD corrections, with the $\sqrt{s} = 0$ intercept verifying the intrinsic k_T prediction.	12
1.7	NA51 and E866 data, with various PDF fits. [20]	13
1.8	Photon center of mass frame for naive Drell-Yan	14
1.9	Collins-Soper Frame	15
1.10	Boer model fit to NA10 data.[27]	17
1.11	E866 data on ν vs p_T compared with NA10 and E615 results. The curves are fits using Boer’s functional form[32].	18
1.12	E866 data on ν in $p + p$ and $p + d$. [33].	19
2.1	SeaQuest spectrometer schematic. The Main Injector proton beam enters on the left.	21
2.2	FNAL accelerator complex schematic. High-lighted from start to finish: Linac, Booster, Main Injector, Tevatron segment, Switchyard.	22
2.3	Čerenkov instrumentation package.	25
2.4	Target Table schematic, top view.	27
2.5	Left: FMAG. Right: KMAG.	28
2.6	Left: xz-plane view of prop tubes. Right: yz-plane view of prop tubes.	32
3.1	SeaQuest v1495 Trigger System initial design schematic [47]	40
3.2	Station 3 x-measuring hodoscopes, in front of the iron wall [47]	41
3.3	CAEN v1495 VME module [47]	42
3.4	Four-phase sampling schematic in v1495 firmware TDC implementation [47]	43
3.5	Digitized hit TDC pipeline [47]	44
3.6	Trigger System Schematic with addition of Level 0. Level 0 output can choose between pass-through hodoscope signals or pulser signals. Level 0 copies the on-board TDC functionality of Level 1 and Level 2.	47
3.7	Signal+Background Mass Spectrum Fit to the small subset of data first released at international meetings.	55
3.8	Comparison of Mass Acceptance for Run 2/Run 3 Trigger	56
3.9	Comparison of x_2 Acceptance for Run 2/Run 3 Trigger	57
4.1	Flow chart of Hit Cluster removal algorithm	62

4.2	Triplet reconstruction in a drift chamber	65
4.3	Sagitta-Ratio Projection to Station 1	66
4.4	Monte Carlo Sagitta-Ratio Distribution	66
4.5	kTracker’s track fitting algorithm flow	67
4.6	Kalman Filter reverse propagation in kTracker track fitting	68
4.7	Single slice of reverse FMAG ‘swim’ in kTracker	68
4.8	Mass residuals for Messy, Clean, and Overlap components of the GMC+NIM3 embedded data	71
4.9	kTracker Rate-dependent Efficiency for Roadset 67, Deuterium target	71
4.10	The reconstruction rate-dependent efficiency fits for Roadset 67, Deuterium target, in 10 bins of θ . Fit parameters are listed in Table 4.4	73
4.11	The reconstruction rate-dependent efficiency fits for Roadset 67, Hydrogen target, in 10 bins of θ . Fit parameters are listed in Table 4.4	74
4.12	The reconstruction rate-dependent efficiency fits for Roadset 62, Deuterium target, in 10 bins of θ . Fit parameters are listed in Table 4.4	75
4.13	The reconstruction rate-dependent efficiency fits for Roadset 62, Hydrogen target, in 10 bins of θ . Fit parameters are listed in Table 4.4	76
4.14	The reconstruction rate-dependent efficiency fits for Roadset 67, Deuterium target, in 10 bins of ϕ . Fit parameters are listed in Table 4.5	77
4.15	The reconstruction rate-dependent efficiency fits for Roadset 67, Hydrogen target, in 10 bins of ϕ . Fit parameters are listed in Table 4.5	78
4.16	The reconstruction rate-dependent efficiency fits for Roadset 62, Deuterium target, in 10 bins of ϕ . Fit parameters are listed in Table 4.5	79
4.17	The reconstruction rate-dependent efficiency fits for Roadset 62, Hydrogen target, in 10 bins of ϕ . Fit parameters are listed in Table 4.5	80
4.18	The reconstruction rate-dependent efficiency fits for Roadset 67, Deuterium. Three equal-statistics bins are plotted for each kinematic variable. The kinematic dependence of the efficiency of the reconstruction is likely linked dominantly to x_F	81
4.19	The reconstruction rate-dependent efficiency is independent of the chamber efficiency	83
4.20	Component-fitted mass spectrum for reconstructed dimuon events from Roadset 62, Deuterium target	86
4.21	Single-track momentum acceptance for high-mass DY Monte Carlo dimuons	88
4.22	Geometric and trigger acceptances in θ . “Matrix 1” is the main physics trigger: opposite sign muons, on opposite sides (Top/Bottom) of the spectrometer. “Matrix 2” is the opposite sign, same side trigger. Therefore, the difference in acceptance between these triggers demonstrates the effect of the Top/Bottom separation in the trigger system.	89
4.23	Geometric and Trigger Acceptances in ϕ . “Matrix 1” is the main physics trigger: opposite sign muons, on opposite sides (Top/Bottom) of the spectrometer. “Matrix 2” is the opposite sign, same side trigger. Therefore, the difference in acceptance between these triggers demonstrates the effect of the Top/Bottom separation in the trigger system.	90
4.24	Uncorrected high-mass dimuon data distribution in θ	92
4.25	Total Acceptance in θ , from GMC	93
4.26	Mix Background distribution in θ	94
4.27	High-Mass DY θ distribution components	95
4.28	Application of Intensity and Acceptance Corrections to the θ distribution for Roadset 67, Deuterium target	96
4.29	Application of Intensity and Acceptance Corrections to the θ distribution for Roadset 67, Hydrogen target	96
4.30	Application of Intensity and Acceptance Corrections to the θ distribution for Roadset 62, Deuterium target	97
4.31	Application of Intensity and Acceptance Corrections to the θ distribution for Roadset 62, Hydrogen target	97
4.32	Total Acceptance in ϕ , from GMC	98
4.33	Uncorrected DY-region Data in ϕ	99

4.34	DY-region Mix Background in ϕ	100
4.35	High-Mass DY ϕ distribution components	101
4.36	Application of intensity and acceptance corrections to the ϕ distribution for Roadset 67, Deuterium target	102
4.37	Application of intensity and acceptance corrections to the ϕ distribution for Roadset 67, Hydrogen target	102
4.38	Application of intensity and acceptance corrections to the ϕ distribution for Roadset 62, Deuterium target	103
4.39	Application of intensity and acceptance corrections to the ϕ distribution for Roadset 62, Hydrogen target	103
5.1	Relationship between the ‘natural axis’ (z') and the CS-frame	105
5.2	CS λ and ν vs p_T for qq and qg processes. Using $M = M_Z = 91.2$ GeV.	109
5.3	Comparison of the CMS data to the calculations described in Section 5.1.3. The figures demonstrate (a) extraction of the $q\bar{q}$ and qG contributions from the λ p_T -dependence, (b) extraction of the average degree of non-coplanarity from the ν p_T -dependence, and (c) comparison of the Lam-Tung violation between the data and calculations. From Ref. [53]	110
5.4	Diagrams contributing to the annihilation process at $\mathcal{O}(\alpha_s)$	111
5.5	Diagrams contributing to the quark bremsstrahlung process at $\mathcal{O}(\alpha_s)$	111
5.6	Flat bounds on A_0 and A_2	116
5.7	Bounds on A_0 and A_2 from the combined single-event parameter definitions.	117
5.8	Kinematically derived bounds on A_0 and A_2 , defining the “allowed triangle”.	118
5.9	CMS Data from Ref. [50]	119
5.10	FNAL E866 Data from Ref. [32, 33]	119
5.11	FNAL E615 Data from Ref. [30, 31]	119
5.12	NA10 Data from Ref. [28, 29]	119

Chapter 1

Introduction

In the quest for a complete understanding of the structure of hadrons, the theory of Quantum Chromodynamics (QCD) [1, 2] revolutionized the field, providing a rigorous Quantum Field Theory (QFT) framework for the existing model predictions and experimental observations. Since its initial development in the 1970s, significant theoretical progress has been made in generating predictions based on QCD for the high-energy interactions of fundamental particles. In parallel, experiments have set to work confirming predictions of the theory, and measuring non-perturbative effects intractable to analytical theory calculations.

While much of the existing experimental results come from lepton scattering experiments, di-lepton production (Drell-Yan) experiments have played a complementary and increasingly important role. The advent of QCD verified the Drell-Yan (DY) model for di-lepton production in spectacular fashion, explaining all the theoretically challenging experimental observations [3]. Drell-Yan experiments are especially important for aspects of hadron structure which are difficult or impossible with lepton scattering, such as sea quark distributions in nucleons and valence quark distributions in pions, kaons, and anti-protons.

The Drell-Yan angular distributions have been an important component of the verification of QCD-improved Drell-Yan. The angular distributions are uniquely equipped to disentangle the contributions from the $q\bar{q}$ and qg subprocesses to the Drell-Yan cross-section at $\mathcal{O}(\alpha_s)$. Drell-Yan angular distributions will become even more important as the study of Transverse Momentum Dependent PDFs (TMDs) matures. Thus far, non-perturbative effects due to the Boer-Mulders function have only been seen in the angular distributions of unpolarized Drell-Yan experiments. In the near future, Single-Spin Asymmetries in polarized Drell-Yan should provide the first experimental verification (or refutation) of the expected DIS/DY sign-change in the Sivers function. If confirmed, the sign-change would be the first confirmation of violation of PDF universality and a powerful vindication of the TMD formalism.




Figure 1.1: Leading Order DIS Process Diagram.

1.1 Background

Detailed study of the structure of the nucleon began in the 1960s at the Stanford Linear Accelerator Center (SLAC)¹. At SLAC, a high-energy electron beam is used to probe the structure of nucleons in fixed targets, by detecting the energy and direction of the scattered beam electron. This experimental technique, in which a lepton beam interacts with a hadronic target via a single, high-energy exchanged photon, is called Deep Inelastic Scattering (DIS). Figure 1.1 depicts the DIS interaction. DIS experiments at SLAC extracted the *inelastic form factors* of the proton and neutron, $F_1(x)$ and $F_2(x)$, by measuring the differential cross-section for ep- and en-scattering:

$$\frac{d\sigma}{dE'd\Omega} = \frac{\alpha^2}{4E^2 \sin^4 \frac{\theta}{2}} \left[\frac{1}{\nu} F_2(x) \cos^2 \frac{\theta}{2} + \frac{1}{M} F_1(x) \sin^2 \frac{\theta}{2} \right] \quad (1.1)$$

where E and E' are the incident and scattered electron energies, respectively, α is the fine structure constant, $\nu = E - E'$ (in the lab frame), M is the proton rest mass, θ is the lab-frame scattering angle, and x is the dimensionless Lorentz scalar defined by $x = \frac{Q^2}{2M\nu}$. $Q^2 = -q^2$, where q^2 is the square of the invariant mass of the exchanged photon.

In general, $F_1(x)$ and $F_2(x)$ should be replaced by $MW_1(\nu, Q^2)$ and $\nu W_2(\nu, Q^2)$, such that the proton structure functions depend independently on both ν , and Q^2 . However, for “high enough” Q^2 , the exchanged photon interacts with only a single parton within the nucleon. In this regime, the structure functions become independent of Q^2 , and can be rewritten as a function of only the “scaling variable”, x . This feature of nucleon structure, called “Bjorken Scaling” [4], is a consequence of the *point-like* constituents of the nucleon.

The interpretation of Q^2 -independent nucleon structure functions as evidence of point-like constituents is

¹Now known as the “SLAC National Accelerator Laboratory”




Figure 1.2: Left: Meson Octet. Right: Baryon Octet. Vertical position indicates strange quark content.

predicated on the “Parton Model” [5], proposed by Feynman in 1969. The parton model considers a hadron in the “infinite-momentum” frame, a limit in which the hadron’s momentum is infinitely large. In this limit, the hadron is Lorentz contracted and time-dilated, such that its constituent partons are at rest. Therefore, the interaction of incident high-energy photons, as in DIS, can be described by an incoherent sum over photon interactions with individual partons. This freedom of partons to interact with high-energy probes, independently of their parent nucleon, is called “asymptotic freedom”. The infinite momentum frame and the parton model are approximately valid for high- Q^2 interactions.

Separately, Gell-Mann [6] and Zweig [7] proposed the “Quark” [8] Model to explain the observed spectrum of hadrons generated by the disrupted proton in DIS interactions. The quark model proposed spin- $\frac{1}{2}$ constituents with fractional charge, permanently confined within hadrons by a “strong” force. ‘Quark confinement’ is required to simultaneously explain the rich spectrum of hadrons, and the non-observation of free quarks. At the time of the original proposal, the quark model only required the up, down, and strange quarks to describe all known hadrons. Later experiments confirmed the existence of six quarks: up, down, charm [9, 10], strange, top [11, 12], and bottom [13].

The quark-model can quite successfully explain the observed “Particle Zoo” of hadrons by organizing them into two groups. The first group, the mesons, are composed of one quark and one anti-quark. The second group, the baryons, are composed of three quarks. (Similarly, anti-baryons are composed of three anti-quarks.) Figure 1.2 depicts the “Eight-Fold Way”, the predicted groups of ground-state baryons and mesons that can be formed with the up, down, and strange quarks.




Figure 1.3: Left: Charge screening in QED. Center: Color screening in QCD. Right: Color “camouflage” in QCD. The camouflage effect dominates the color screening effect, and quarks within hadrons become asymptotically free in the high- Q^2 limit.[14]

1.1.1 Quantum Chromodynamics

Quantum Chromo-Dynamics [2, 1] (QCD), the Quantum Field Theory (QFT) responsible for the interaction of color-charged fundamental particles, is a pillar of the Standard Model of Particle Physics [14]. QCD formally marries the Quark Model with the Parton Model, identifying the quarks responsible for the hadron spectrum as identical to the partons responsible for Bjorken scaling in DIS. Famously, QCD unifies the seemingly-inconsistent observations of quark confinement and asymptotic freedom.

QCD is a copy of the fantastically successful electromagnetic quantum field theory, QED [15], adding two very important alterations:

First, “color charge” replaces electric charge. There are six color charges, $(r, \bar{r}, g, \bar{g}, b, \bar{b})$, in contrast to two electric charges $(+, -)$. The “strong” coupling strength between color charges is the origin of quark confinement and asserts that all observable particles must be “color-less”. Thus, mesons are formed from the color-neutral combinations $(r\bar{r}, b\bar{b}, g\bar{g})$, baryons are (rgb) , and anti-baryons are $(\bar{r}\bar{g}\bar{b})$.

Second, the photon, the gauge boson responsible for mediating QED interactions, is replaced by eight gluons. Whereas the photon is not electrically charged, the eight gluons have independent color-charges, formed from all possible combinations of color-anticolor, excluding the color-neutral pairings. This feature of the gluons has profound consequences for the QCD theory. In each interaction with a gluon, a quark must *change its color*. Because gluons are both the interaction mediators and color-charged themselves, gluons can self-interact.

These features of QCD cause a fascinating reversal of the “charge-screening” effect, observed as the running of the coupling constant in QED. In QED, a bare charge will polarize the vacuum around itself, aligning the short-lived pairs of virtual particles to cloak itself in positive charge, as seen in the left panel of

Figure 1.3. Long-distance (low-energy) probes will see a reduced effective charge. In the high-energy (small distance) limit, the full charge is resolved. This energy-scale-dependent modification of effective charge is parameterized by a Q^2 -dependence of the coupling constant, α . QCD also experiences an equivalent color-screening. The polarization of the color-charged vacuum causes an effective reduction in the observed strength of the coupling to a hypothetical quark. This is depicted in the central panel of Figure 1.3. However, the quark will also continuously emit and absorb virtual gluons. Since gluons are themselves color-charged, they effectively “carry away” the true color of quark, spreading it out over a larger volume. This “camouflage” effect overpowers the screening effect, and the strong coupling constant α_s , *decreases* with increasing Q^2 . Therefore, quarks are simultaneously confined within hadrons by a strong-binding force *and* behave as free, unbound, point particles when struck by a high- Q^2 probe. It is for this reason alone that perturbation theory can be applied to QCD in the high-energy limit.

1.1.2 Parton Distribution Functions

The Parton Distribution Functions (PDFs) describe the longitudinal structure of hadrons within the context of the Quark-Parton Model (QPM). A PDF is interpreted as the probability to find a quark, of a particular charge and momentum fraction, in a hadron in inelastic interactions. The PDFs, specified for each hadron and quark flavor, are functions of only a single variable, x . In this context, x is interpreted as the momentum fraction of the quark, i.e. the fractional contribution of the quark to the longitudinal momentum of the hadron. It is also identified with the Bjorken scaling variable, $x_{Bj} = \frac{Q^2}{2M\nu}$, which describes the properties of the exchanged virtual photon in deeply-inelastic lepton scattering experiments. This association is only exact in the “infinite momentum frame”, where $|\vec{p}_{\text{hadron}}| \gg M_{\text{hadron}}$. In this frame, the assumptions of quark freedom and longitudinal-only hadron structure become exact, and the impulse approximation unifies the Bjorken scaling variable and the parton momentum fraction. In the absence of QCD-induced scaling violations, the PDFs are defined by:

$$2xF_1(x) = F_2(x) = \sum_f e_f^2 xq_f(x) \tag{1.2}$$

where the sum over f is a sum over parton species:

$$\begin{aligned}
q_f(x) & \quad \text{PDF for quark of flavor } f \in (u, d, s, c) \\
\bar{q}_f(x) & \quad \text{PDF for antiquark of flavor } f \in (u, d, s, c) \\
g(x) & \quad \text{PDF for gluon}
\end{aligned}$$

The hadron PDFs, in addition to serving as a container for acquired experimental knowledge about the longitudinal momentum structure of hadrons, provide a simple way to write down various expectations of hadron structure, usually written down as “sum rules”. The valence “sum rule”:

$$\int_0^1 dx [q_f(x) - \bar{q}_f(x)] = N_f^v \tag{1.3}$$

constrains the quark anti-quark distributions such that the net number of quarks matches the naive, valence quark composition. The sum of the first moments of all partons within a hadron should yield the total longitudinal momentum of the hadron:

$$\sum_f \left[\int_0^1 dx [xq_f(x)] + \int_0^1 dx [x\bar{q}_f(x)] \right] + \int_0^1 dx [xg(x)] = 1 \tag{1.4}$$

This is simply the expectation of (longitudinal) momentum conservation in the QPM written in terms of PDFs.

Of particular interest to the study of the nucleon sea is the Gottfried Sum Rule [16]. The Gottfried Sum quantifies the longitudinal quark structure differences between the proton and the neutron.

$$S_G = \int_0^1 dx \sum_f e_f^2 [q_f^p(x) + \bar{q}_f^p(x) - q_f^n(x) - \bar{q}_f^n(x)] \tag{1.5}$$

Under the assumption that the proton and neutron are perfect isospin conjugates, i.e. $q_u^n = q_d^p$, $\bar{q}_u^n = \bar{q}_d^p$, etc. the Gottfried Sum becomes

$$S_G = \int_0^1 dx \left(\frac{1}{3} \right) [q_u(x) + \bar{q}_u(x) - q_d(x) - \bar{q}_d(x)] \tag{1.6}$$

where we have dropped the superscripts since all PDFs now refer to the proton. Plugging in the assumption

that the valence sum rule holds for the net up and net down quarks in the proton:

$$S_G = \frac{1}{3} - \int_0^1 dx \left(\frac{2}{3} \right) [\bar{q}_d(x) - \bar{q}_u(x)] \quad (1.7)$$

Assuming light-quark nucleon sea symmetry, $S_G = \frac{1}{3}$, which is the Gottfried Sum Rule. The experimental violation of this prediction came as a surprise, and is described in Section 1.3.

Loosening the aforementioned restrictions of the co-linear quark-parton model, it is possible to write down additional leading-twist PDFs which are dependent on the transverse momentum of the interacting parton. Some of the so-called TMDs result from initial-state and/or final state interactions of the partons and represent a slight breakdown of the impulse approximation picture. The TMDs are the first experimental hook into the measurement of the effects of orbital angular momentum (OAM) of the partons within hadrons. The experimental determination and theoretical understanding of the TMDs are currently topics of intense interest in the field of hadronic structure.

1.2 The Drell-Yan Process

First proposed by Sidney Drell and Tung-Mow Yan [17] to explain the experimentally observed continuum of dilepton pairs produced in hadron collisions, the Drell-Yan process has become a mainstay of the study of hadron structure and is complementary to DIS. The Drell-Yan process has been a crucially important second electromagnetic window on the study of nucleon structure and QCD. The lepton pair is produced by the decay of a massive virtual photon, which was produced by the electromagnetic annihilation of a quark/anti-quark pair. In order to produce a massive virtual photon, the quark and anti-quark must come from different hadrons. By conservation of energy and momentum, the lepton pair carries information about the momenta of the annihilating quarks. Since leptons are not color-charged, the Drell-Yan process can act as a clean, electromagnetic probe of hadron structure.

1.2.1 Kinematics

In the hadron-hadron center of mass frame, the hadron 4-momenta are $P_1 = (P, 0, 0, P)$ and $P_2 = (P, 0, 0, -P)$. The center-of-mass energy of the hadron-hadron system is $\sqrt{s} = 2\sqrt{P * P} = 2P$. With the assumption of negligible parton transverse momentum ($k_T = 0$), and negligible quark mass, the annihilating parton momenta are:

$$p_1 = x_1 P_1 = (x_1 P, 0, 0, x_1 P) \quad p_2 = x_2 P_2 = (x_2 P, 0, 0, x_2 P) \quad (1.8)$$




Figure 1.4: Leading Order Drell-Yan Process Diagram.

where $0 < x_i < 1$ is the fraction of the hadron momentum carried by the i th parton. The square of the invariant mass of the virtual photon is then given by $M_\gamma^2 = Q^2 = \hat{s} = x_1 x_2 s$. The longitudinal momentum fraction of the virtual photon (x -Feynman) is $x_F = 2 \frac{p_1 - p_2}{\sqrt{s}} = x_1 - x_2$.

In general, the virtual photon can have a non-zero transverse momentum, p_T . In Q^2 regimes where the parton transverse momentum (k_T) is non-negligible, the vector sum of the k_T of the annihilating quark and anti-quark contributes to the photon p_T . Effects from higher-order α_s processes (such as initial state gluon emission) can also make significant contributions to p_T . When the photon p_T is non-zero, there is a lab-frame observable, ϕ_γ , which describes the azimuthal direction of the photon p_T around the lab-frame z -axis (the beam axis).

The last two observables require a choice of axes in the photon rest frame. In all the commonly used photon rest frame conventions, the y -axis is chosen as parallel to $\vec{P}_1 \times \vec{P}_2$. This document will use the Collins-Soper (CS) frame throughout, except where otherwise specified. In the CS frame, the z -axis is chosen parallel to the bisector of \vec{P}_1 and $-\vec{P}_2$. Figure 1.9 depicts this inertial frame. θ_{CS} and ϕ_{CS} are defined by this frame. θ_{CS} is the polar angle between the positive muon momentum and the z -axis. ϕ_{CS} is the azimuthal angle of the positive muon around the z -axis.

So, the full set of observables derivable from the measurement of dilepton momenta in Drell-Yan experiments is $(M_\gamma, x_F, x_1, x_2, p_T, \phi_\gamma, \theta_{CS}, \phi_{CS})$. As shown above, the pairs of observables $(M_\gamma)_{,x_F}$ and (x_1, x_2) are not independent, being related through \sqrt{s} . Therefore, the full set of observable kinematics is calculable from just six experimentally measured quantities: the three-momenta of both muons $(p_{1x}, p_{1y}, p_{1z}, p_{2x}, p_{2y}, p_{2z})$.

1.2.2 Dynamics

The dynamics of the naive Drell-Yan process are derived by embedding the standard $e^+e^- \rightarrow \mu^+\mu^-$ result in a hadron-hadron collision system, replacing the electrons with quarks, and folding in the quark and antiquark

PDFs.

The naive Drell-Yan model, first worked out by Drell and Yan in 1970 [17], relies on the impulse approximation to simplify the derivation of the cross section for the $q\bar{q} \rightarrow \mu^+\mu^-$ process. The impulse approximation assumes that the partons within the colliding hadrons can be treated as asymptotically free. That is, in the interaction with the virtual photon, the quarks act as free, unbound particles. With this assumption, the Drell-Yan cross section can be factorized into two independent parts: a hard-scattering annihilation of a free quark with a free anti-quark, and the momentum distributions of the quark and anti-quark within their respective hadrons. The impulse approximation additionally asserts that, on the time-scale of the annihilation process, the annihilating quark and anti-quark are uncorrelated with the other partons in their respective hadrons.

The calculation of the hard-scattering annihilation process is based on the standard calculation of the QED annihilation diagram, using spin- $\frac{1}{2}$ fermions and a spin-1 virtual photon. The cross-section for $e^+e^- \rightarrow \mu^+\mu^-$ is:

$$\sigma(e^+e^- \rightarrow \mu^+\mu^-) = \frac{4\pi\alpha^2}{3M^2} \quad (1.9)$$

Substituting the electrons for quarks requires multiplying by a color factor, and replacing the electron charge, -1 , with the quark charge, e_f . For up quarks: $e_u^2 = \frac{4}{9}$. For down quarks: $e_d^2 = \frac{1}{9}$. Only colorless pairings of quark-anti-quark pairs can undergo QED annihilation. Since there are three colors (red, green, and blue) there are 9 total pairings and 3 colorless pairings. Therefore, the color factor is $c_f = \frac{1}{3}$. The hard-scattering process cross-section for the $q_f\bar{q}_f \rightarrow \mu^+\mu^-$ process is:

$$\sigma(q_f\bar{q}_f \rightarrow \mu^+\mu^-) = \left(\frac{1}{3}\right) \left(\frac{4\pi\alpha^2}{3M^2}\right) e_f^2 \quad (1.10)$$

Finally, we must embed the annihilating quarks in their parent hadrons. Under the assumption of asymptotic freedom, the partons in each hadron are uncorrelated with each other, so we can simply assign a probability density to find each quark (antiquark) of a particular flavor and momentum fraction. These probability densities are exactly given by the parton distribution functions (PDFs), q_f and \bar{q}_f . The probability for finding a quark of flavor f , with a momentum fraction of x in a hadron is $q_f(x)dx$. For anti-quarks, it's $\bar{q}_f(x)dx$. Therefore the combined probabilities for finding the partons required to generate Drell-Yan dileptons are:

$$q_f(x_1)\bar{q}_f(x_2)dx_1dx_2 \quad \text{For } q \text{ from } h_1 \text{ and } \bar{q} \text{ from } h_2 \quad (1.11)$$

$$\bar{q}_f(x_1)q_f(x_2)dx_1dx_2 \quad \text{For } q \text{ from } h_2 \text{ and } \bar{q} \text{ from } h_1 \quad (1.12)$$

Since the identities of the partons contributed by each hadron is unobservable, the total probability required for Drell-Yan is the sum of the above products. Combining with Eq. 1.10, and summing over quark flavors, we arrive at the final result for the differential Drell-Yan cross section in hadron-hadron collisions:

$$\frac{d\sigma}{dx_1 dx_2} = \left(\frac{1}{3}\right) \left(\frac{4\pi\alpha^2}{3M^2}\right) \sum_f e_f^2 [q_f(x_1)\bar{q}_f(x_2) + \bar{q}_f(x_1)q_f(x_2)] \quad (1.13)$$

1.2.3 Universality

A key feature of the impulse approximation is universality: the universal applicability of PDFs to all hard-scattering quark interactions, regardless of the specific process. This feature ultimately derives from factorizability. If the Drell-Yan cross-section is factorizable into two independent parts, the hard process and the PDFs, and if the DIS (SIDIS)² cross-section is factorizable into two (three) independent parts, the hard process and the PDFs (and the fragmentation functions), then the PDFs measured by both processes should be identical.

Indeed, in the perturbative regime, quark distributions measured by DIS and Drell-Yan experiments are largely consistent. It is only in kinematic regions where non-perturbative effects are large that complicating effects (e.g. TMDs) are expected to violate universality. Experimental verification of universality-violating distribution functions (e.g. the Sivers function) remains unseen.

1.2.4 QCD-Improved Drell-Yan

With the advent of QCD, the Drell-Yan mechanism transformed from a model for dilepton production through parton-parton annihilation into a natural consequence of high-energy collisions of QCD bound-states through quark, anti-quark, and gluon annihilation [3]. Furthermore, QCD improves many theoretical predictions of the Drell-Yan mechanism, providing natural explanations for the K-factor, large p_T , logarithmic Q^2 dependence, and the angular distributions.

Figure 1.5 lists the diagrams contributing to lepton pair production in hadron collisions at $\mathcal{O}(\alpha_s)$. Diagrams in (a) and (b) are natural QCD extensions to naive Drell-Yan. In (a), the $q\bar{q}$ -annihilation diagram is subject to all possible single gluon loop corrections. In (b), the quark or the anti-quark emits a gluon before annihilating with the other into the Drell-Yan photon. The diagrams in (c) are quite different. This is the so-called “quark-bremsstrahlung” process, in which a qg initial state fuses, yielding a quark and the “Drell-Yan” photon in the final state.

²SIDIS (Semi-Inclusive Deep-Inelastic Scattering) is DIS where one or more final state hadrons are detected in addition to the scattered electron.




Figure 1.5: $\mathcal{O}(\alpha_s)$ diagrams for QCD improved Drell-Yan. (a) One-loop corrections to the LO diagram. Only the interference terms contribute at $\mathcal{O}(\alpha_s)$. (b) Gluon emission diagrams, $q\bar{q}$ -diagrams (c) QCD “bremsstrahlung” diagrams, qg -process [3].

The K-factor was one of the major challenges in verifying the naive Drell-Yan picture with experimental data. Experimental efforts measured the total Drell-Yan cross-section to be ~ 2 times larger than predicted by the naive Drell-Yan model. Since the model described the *shape* of the Drell-Yan mass distribution quite well, the K-factor was adopted, a fudge factor to match the predicted total cross-section to the experimentally measured values. Resummation techniques for all orders of α_s result in a modification of the total cross-section consistent with the experimentally measured K-factor of ~ 2 .

QCD resummation also accounts for observed logarithmic Q^2 scaling violations. In the modern language, the parton distribution functions themselves are modified by QCD evolution and generate a logarithmic dependence on Q^2 .

In the naive picture, only the vector sum of the intrinsic k_T of the annihilating quark and anti-quark can contribute to the p_T of the photon. Therefore, it was quite a surprise when experiments measured $\langle p_T \rangle$ much larger than the 300–400 MeV expected from the Fermi motion of the partons. However, gluon emission (quark bremsstrahlung) diagrams in Figure 1.5 predict a much larger p_T , due to momentum conservation with the off-axis final state gluon (quark). The experimental data (Figure 1.6) show a linear dependence of the average transverse momentum on \sqrt{s} . The extrapolated value at $\sqrt{s} = 0$ agrees quite well with the predicted intrinsic k_T contribution. Further, the intercept for pion-induced Drell-Yan is larger than that for




Figure 1.6: Experimental Results for $\langle p_T \rangle$ vs \sqrt{s} for proton- and pion-induced Drell-Yan. The linear dependence is due to QCD corrections, with the $\sqrt{s} = 0$ intercept verifying the intrinsic k_T prediction.

proton-induced Drell-Yan, as expected for Fermi motion in the smaller volume of the pion.

The angular distributions of naive, co-linear Drell-Yan exactly match the expectation from the QED subprocess $q\bar{q} \rightarrow \gamma \rightarrow l^+l^-$, which is $\frac{d\sigma}{d\Omega} \propto 1 + \cos^2 \theta$. With the addition of the small intrinsic k_T , the angular distributions can vary from this simple prediction, but only slightly. Large deviations, such as those observed at high- p_T in recent measurements by CMS, are evidence that the QCD-improved diagrams play a major role. In fact, the angular distributions can be used to disentangle the contributions of the $q\bar{q}$ and gq diagrams to the lepton pair production cross-section. This idea is explored further in Section 5.3.

1.3 The Nucleon Light-Flavor Sea Asymmetry

One feature of the Drell-Yan process, which is uniquely important to the study of hadron structure, is the requirement of an anti-quark in the initial state. This feature means that Drell-Yan experiments can open a window into the structure of the hadron sea, which is largely inaccessible to DIS, and even SIDIS.

The NMC experiment, using DIS, measured the proton and neutron structure functions using H_2 and D_2 targets. [18] They then extracted $\bar{u} - \bar{d}$, and published the first observation of a significantly flavor-asymmetric nucleon sea. The NA51 experiment first measured the $\frac{\bar{d}}{\bar{u}}$ asymmetry with the Drell-Yan process in 1991 at $\langle x_2 \rangle = 0.17$. [19] These results were quite unexpected. Until this result, it was assumed that the nucleon sea-quarks originated from gluon splitting, and, having very similar mass, the \bar{u} - and \bar{d} -quarks should be produced in similar proportions. The significant \bar{d} dominance prompted large changes to PDF fits




Figure 1.7: NA51 and E866 data, with various PDF fits. [20]

for anti-quark distributions, and many theoretical proposals for the origin of the asymmetry.

Another surprise came in 1998, when the E866 experiment at Fermilab improved upon NA51, measuring the x_2 -dependence of the asymmetry in the range $0.015 < x_2 < 0.35$. [20] In addition to confirming the NA51 result at $x_2 = 0.17$, the E866 data showed that the asymmetry appeared to drop back toward symmetry at the upper end of the measured x_2 -range, contrary to the prevailing assumptions. The true trend of asymmetry is difficult to extrapolate, as the uncertainties in the higher x_2 bins get quite large. Nevertheless, this result again had a profound impact (Figure 1.7) on anti-quark PDF fits and theoretical attempts to explain the light-quark sea asymmetry. In particular, the apparent turnover of the asymmetry at $x_2 = 0.27$ is especially theoretically challenging.

To further investigate the high- x_2 trend of \bar{d}/\bar{u} , the SeaQuest Collaboration proposed [21] to repeat the E866 measurement using the 120 GeV beam provided by the FNAL Main Injector. The lower beam energy and higher intensity of the MI beam will allow SeaQuest to measure the light sea flavor asymmetry up to $x_2 = 0.53$. The SeaQuest experiment is fully described in Chapter 2. SeaQuest is currently taking data, and preliminary results have been presented at international conferences.

1.4 Drell-Yan Angular Distributions

The angular distributions of Drell-Yan dimuons have proven to be a fantastically useful tool in the study of hadron structure and, especially, the details of QCD-improved Drell-Yan. Through the polar angle dis-




Figure 1.8: Photon center of mass frame for naive Drell-Yan

tribution, Drell-Yan experiments have verified the spin- $\frac{1}{2}$ nature of the quarks, and the expected transverse polarization of the virtual photon [22]. Through the observation of Lam-Tung violation, Drell-Yan angular analyses have demonstrated the breakdown of the naive Drell-Yan picture. In small- p_T data, these breakdowns are variously interpreted as indications of higher-twist effects and/or Transverse Momentum Dependent Parton Distribution Functions. Lam-Tung violation in large- p_T data is addressed in Chapter 5.

1.4.1 Photon Rest Frames

Consideration of Drell-Yan angular distributions is always done in a frame in which the virtual photon is at rest, with the virtual photon at the origin. In such a frame, the quark and anti-quark approach the origin head-on, with exactly opposite momenta. Likewise, the leptons leave the origin back-to-back, with exactly opposite momenta. In the naive Drell-Yan model, the quarks are co-linear with their parent hadrons, and so it is natural to align the z-axis with the quark axis. Figure 1.8 depicts the photon rest frame for naive DY. Note that, due to $p_T = 0$ GeV, the azimuthal symmetry is unbroken, and the x-axis, y-axis, hadron plane, and ϕ -angle is undefined. In general, however, the quarks are not co-linear with the hadrons, and it is experimentally impossible to define a reference frame using the quark axis since it is unobservable.

There are three common rotation choices for the axes of the photon rest frame: Collins-Soper (CS), Gottfried-Jackson (GJ), and u-channel. In all three frames, the y-axis points along the direction of $\vec{p}_{h1} \times \vec{p}_{h2}$, perpendicular to the hadron plane. This direction is well-defined as long as $p_T > 0$ GeV. Therefore, all three common frames only differ via a rotation in the xz -plane.




Figure 1.9: Collins-Soper Frame

- The Gottfried-Jackson frame defines the z-axis as parallel to the beam hadron's momentum: \vec{p}_{h_1} .
- The u-channel frame defines the z-axis as anti-parallel to the target hadron's momentum: \vec{p}_{h_2} .
- The Collins-Soper frame splits the difference, defining the z-axis as parallel to the bisector of \vec{p}_{h_1} and \vec{p}_{h_2} .

This document will use the CS-frame, except where otherwise specified.

1.4.2 General Form of Drell-Yan Angular Distributions

The general form of the unpolarized Drell-Yan angular distribution is given by:

$$\frac{d\sigma}{d\Omega} \propto 1 + \lambda \cos^2 \theta + \mu \sin 2\theta \cos \phi + \frac{\nu}{2} \sin^2 \theta \cos 2\phi \quad (1.14)$$

where λ , μ , and ν are parameters which completely describe the angular distribution, and depend on the choice of photon rest frame. An intuitive kinematic derivation for this expression is given in Section 5.1.1.

1.4.3 The Lam-Tung Relation

The Lam-Tung relation [23] is a remarkable, non-trivial equality relating the λ and ν parameters of the general Drell-Yan angular distribution.

$$\lambda = 1 - 2\nu \quad (1.15)$$

The Lam-Tung relation is trivially obeyed in naive Drell-Yan, where $\lambda = 1$ and $\nu = 0$. It is also expected to hold exactly for QCD-improved Drell-Yan, $\mathcal{O}(\alpha_s)$. At $\mathcal{O}(\alpha_s^2)$, some Lam-Tung violation is expected. However, it is expected to remain small [24].

The Lam-Tung relation is the Drell-Yan counterpart to the Callan-Gross relation [25] in DIS. Equivalently to the Callan-Gross relation, it is a direct consequence of the spin- $\frac{1}{2}$ nature of the annihilating quarks and the produced leptons. In the high-energy limit, the quark and lepton masses are negligible, and helicity conservation requires a transverse mediating photon. A kinematic derivation of the Lam-Tung relation is detailed in Section 5.2.1.

Despite the strong theoretical expectation, significant Lam-Tung violation has been observed in various Drell-Yan experiments. Interpretation of these observed violations typically cites one of two sources, which dominate in different kinematic regimes. When $p_T \sim k_T$, the intrinsic parton transverse momentum contributes significantly to the observed p_T , and a non-zero Boer-Mulders function is expected to induce Lam-Tung violation. In contrast, when $p_T \gg k_T$, the intrinsic parton transverse momentum is negligible, and any Boer-Mulders function should have no impact on the Lam-Tung relation. In this regime, often probed by high-energy collider experiments, any Lam-Tung violation must be interpreted as evidence of significant contributions from higher-order diagrams.

1.4.4 The Boer-Mulders Function

The Transverse Momentum Dependent distribution functions (TMDs) [26] are a novel type of PDF that describes non-perturbative correlations between various combinations of quark-spin, quark-transverse momentum, and hadron spin.

The Boer-Mulders function, h_1^\perp , is the TMD associated with a correlation of the transverse spin and transverse momentum of quarks within unpolarized nuclei. As such, it is the only TMD accessible to unpolarized Drell-Yan.

In unpolarized Drell-Yan, only a measurement of the convolution of the two h_1^\perp functions, one for the beam and one for the target, is possible. The effect of non-zero Boer-Mulders functions is a large, positive deviation of the naive $\nu = 0$ expectation, increasing with increased p_T . Therefore, it has been invoked [27] to explain the experimental observations of large Lam-Tung violation seen in pion-induced Drell-Yan experiments.

The most striking violation of the Lam-Tung relation was observed in two independent pion-induced Drell-Yan experiments, CERN NA10 [28, 29] and FNAL E615 [30, 31]. Both experiments measured significantly non-zero values of ν , increasing with p_T . The Drell-Yan cross section in these pion-induced fixed target experiments is dominated by “double-valence” annihilation, $q_v \bar{q}_v$. Therefore, the measured Lam-




Figure 1.10: Boer model fit to NA10 data.[27]

Tung violation, if caused entirely by the Boer-Mulders mechanism, would arise due to the convolution of two *valence* Boer-Mulders functions. Interpreting the shape of ν as entirely due to non-zero pion and proton valence Boer-Mulders functions, Boer derived an expression for the shape of the p_T -dependence of ν :

$$\nu = 16\kappa_1 \frac{p_T^2 M_C^2}{(p_T^2 + 4M_C^2)^2} \quad (1.16)$$

where κ_1 and M_C are the fitting parameters.

Boer found that the NA10 data are well described by the Boer-Mulders function functional form. Boer's result, fitted to the NA10 data, is shown in Figure 1.10.

The E866 experiment at FNAL was the first to measure Drell-Yan angular distributions with proton-induced Drell-Yan. Results on λ , μ , and ν were reported as a function of p_T , and compared with the pion-induced results. In contrast to the pion-induced data, the antiquark in the proton-induced data must originate from the sea, as neither the target nor beam hadron contains valence antiquarks. In this way, the angular distributions of proton-induced Drell-Yan are uniquely suited to directly access the Boer-Mulders function for sea-quarks. Fascinatingly, the E866 data show[32, 33] that ν is consistent with zero throughout the measured p_T range. Coupled with the measurement of relatively minor deviations from $\lambda = 1$, E866 finds no significant Lam-Tung violation. The E866 $p+d$ data are shown in Figure 1.11, plotted with the NA10 and E615 data, and fit with Boer's functional form. The fits performed by Zhu et al.[32] find $\kappa_1 = 0.93 \pm 0.10$ for the E615 data and $\kappa_1 = 0.11 \pm 0.04$, using $M_C = 2.4 \text{ GeV}/c^2$ from the original fit to the NA10 data. If the




Figure 1.11: E866 data on ν vs p_T compared with NA10 and E615 results. The curves are fits using Boer’s functional form[32].

large Lam-Tung violations of the pion-induced data are indeed caused by a valence Boer-Mulders function, the E866 results strongly suggest that the Boer-Mulders function for sea-quarks is much smaller, possibly zero.

The Boer-Mulders function is one of only two T-odd TMDs. As such, it is predicted [34] to have opposite sign when measured in Drell-Yan vs DIS interactions. The Boer-Mulders functions for valence u and d quarks in the proton have been measured by SIDIS experiments[35] and found to be negative. In proton-induced Drell-Yan, experiments are sensitive only to the product of the valence and sea Boer-Mulders functions via the coefficient on the $\cos 2\phi$ term, ν , in the angular distribution. The measurement of positive values for ν by Zhu et al. [32], shown in Figure 1.12, suggest that the u and \bar{u} Boer-Mulders functions in DY have the same sign³. Therefore, the TMD “sign-flip” predicts that the DIS Boer-Mulders function for \bar{u} should be negative. This expectation remains to be tested experimentally.

The SeaQuest experiment is currently taking 120 GeV, proton-induced Drell-Yan data at FNAL. These data should permit an analysis of the Drell-Yan angular distributions similar to the E866 results. An angular analysis and Boer-Mulders extraction should confirm the positive-valued measurement of ν from E866. Reporting progress toward a full angular analysis of SeaQuest’s data is one purpose of this dissertation.

1.4.5 The Sivers Function

Since the proton is a spin- $\frac{1}{2}$ particle, and it is composed of three spin- $\frac{1}{2}$ valence quarks, it was initially assumed that the spins of the valence quarks, in their lowest-energy configuration, ($\uparrow\uparrow\downarrow$), must constitute the entirety of the proton’s spin. In this naive picture, any contribution from the sea quarks and gluons

³assuming u -quark dominance




Figure 1.12: E866 data on ν in $p + p$ and $p + d$. [33].

must average out to contribute nothing to the overall proton spin.

An experimental test of this assumption was first reported in 1988 by the EMC collaboration [36]. The experiment used a longitudinally-polarized muon beam, incident on a longitudinally-polarized proton target, to extract the structure of the spins of the quarks within the proton. They found that the quark spins constitute only a small part of the proton spin. This finding was dubbed the “Proton Spin Puzzle”, and led to a surge of experimental and theoretical interest in the study of the spin-structure of the proton.

Parton Orbital Angular Momentum (OAM) could play a significant role in the full accounting of the proton’s spin structure. One path to accessing the parton OAM is through Transverse Momentum Dependent PDFs (TMDs). Specifically, one TMD of great interest to current and future Drell-Yan experiments is the Sivers function, f_{1T}^\perp [37]. The Sivers function captures the correlation between the transverse spin of the proton and the transverse momentum of the interacting quark. Therefore, it is only accessible in experiments with polarized hadrons.

Similarly to the Boer-Mulders function, the expected non-universal behavior of the Sivers function is of great experimental interest. Due to its naive-T-odd nature, it is expected that the sign of the Sivers function will differ between DIS and Drell-Yan measurements [38]. This profound prediction provides an excellent concrete test of the validity of describing the proton spin structure in terms of non-perturbative, time-reversal-breaking TMDs.

The Sivers asymmetry for pion and kaon lepto-production with polarized protons has been measured and found to be non-zero [39, 40]. A Drell-Yan measurement has not been performed yet. Two experimental efforts have been proposed and have begun work on experimentally verifying the sign-flip prediction: COMPASS-II at CERN, and upgraded SeaQuest (E1039) at FNAL.

Chapter 2

The SeaQuest Experiment

The SeaQuest experiment, designated Experiment 906 (E906), is a fixed-target dimuon experiment currently running at the Fermi National Accelerator Laboratory (FNAL). SeaQuest utilizes the high-intensity, 120 GeV proton beam delivered by the FNAL Main Injector (MI) to measure Drell-Yan dimuon production off of various nuclear targets. The experiment consists of relatively few main components, including a target table, two dipole magnets, eight scintillating hodoscope planes, eighteen wire-chamber planes, and four proportional tube planes. This chapter describes the purpose, configuration, and operation of each of the major experimental components.

2.1 Overview

SeaQuest is the latest in a long history of fixed-target dimuon experiments operated at FNAL. It is the first to use the 120 GeV proton beam provided by the Main Injector. All previous FNAL dimuon experiments used the 800 GeV proton beam provided by the Main Ring. The lower beam energy and higher intensity of the MI beam provide the opportunity to extend previous measurements into previously inaccessible kinematic regions. Previous experiments published measurements of the light sea flavor asymmetry [20], Drell-Yan angular distributions [32, 33], quarkonia angular distributions [41, 42], nuclear dependence effects in Drell-Yan [43] and charmonium production [44], and absolute Drell-Yan cross-sections [45]. SeaQuest aims to repeat all of these measurements at a lower \sqrt{s} (15 GeV vs 39 GeV), pushing sensitivity into heretofore unexplored kinematic territory.

The 120 GeV protons delivered by the FNAL accelerator complex interact in one of the targets or in the first magnet, FMAG, which also serves as a beam dump since it is built around a solid iron core. All long-lived particles generated by the proton interactions are absorbed by the 5 m long iron core of FMAG, with the exception of muons. High-energy muons, resulting from pion decay, charmonium decay, and the Drell-Yan process, easily penetrate the iron, suffering relatively minor multiple-scattering, and enter the detector region of the spectrometer.




Figure 2.1: SeaQuest spectrometer schematic. The Main Injector proton beam enters on the left.

The high-momentum muons pass through four tracking stations, an air-core dipole magnet (K MAG), and an additional one-meter-thick iron wall. The layout of the SeaQuest spectrometer is depicted in Figure 2.1. One tracking station sits between the two dipole magnets, allowing a measurement of muon momentum from the deflection angle in K MAG. The iron wall sits between the final two tracking stations, providing additional particle identification discrimination power.

The SeaQuest coordinate system is defined by the incoming protons traveling in the positive z -direction, and gravity pointing in the negative y -direction. Therefore, the fields of FMAG and K MAG both point in either the $+y$ or $-y$ direction, and the “ p_T -kick” is either in the $+x$ or $-x$ directions. The origin of the coordinate system is the intersection of the beam axis and the front-face of FMAG. The terms *upstream* and *downstream* refer to the $-z$ and $+z$ directions, respectively. The terms *Top* and *Bottom* refer to the $+y$ and $-y$ halves of the spectrometer, respectively. The terms *bend-plane* and *non-bend-plane* refer to the zx -plane and zy -plane, respectively.

2.2 The Proton Beam

2.2.1 The Main Injector

The high-intensity proton beam delivered to many experimental areas throughout the FNAL accelerator complex is first accelerated by a series of machines. The chain begins with the Ion Source, which generates

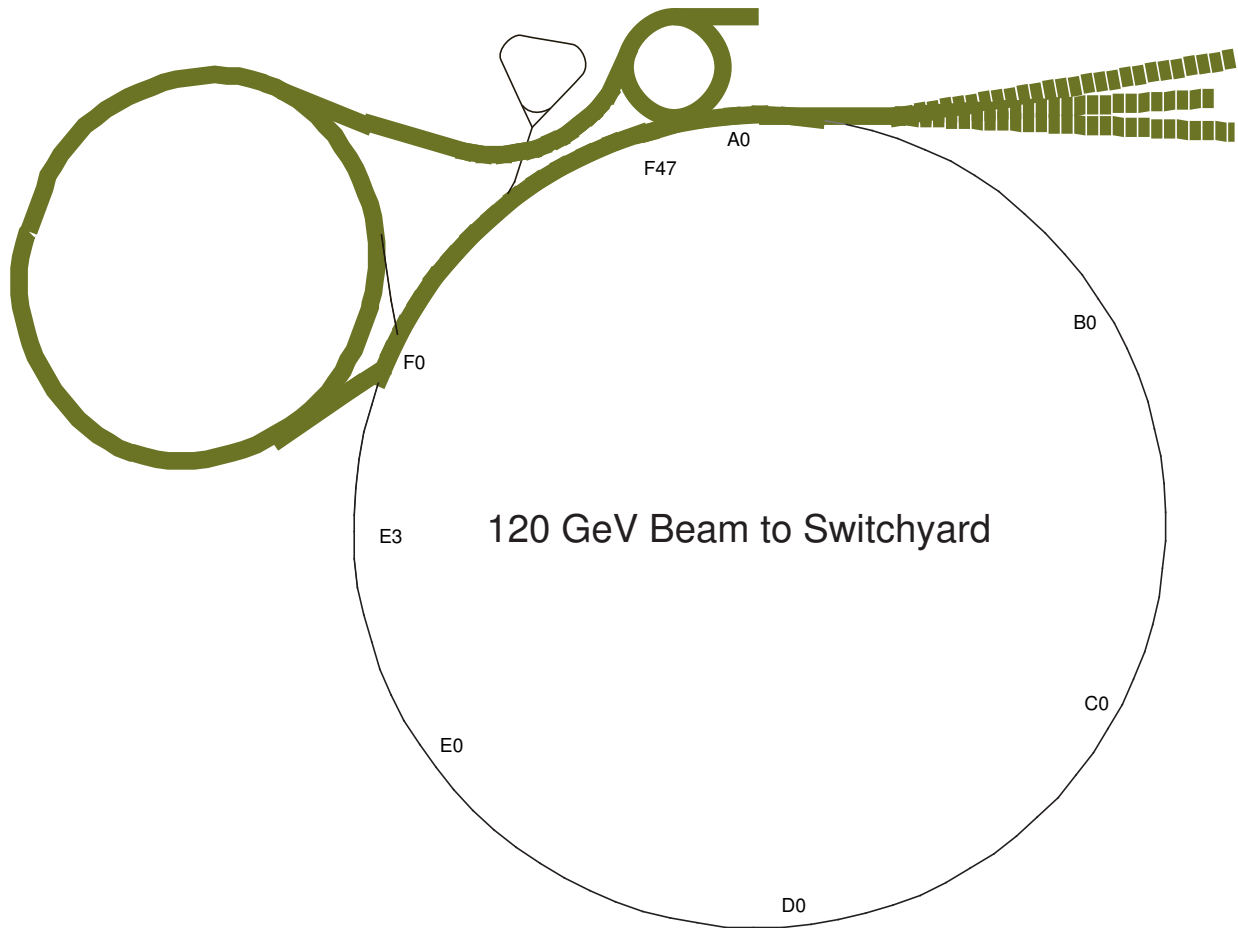


Figure 2.2: FNAL accelerator complex schematic. High-lighted from start to finish: Linac, Booster, Main Injector, Tevatron segment, Switchyard.

a continuous 35 keV proton beam by accelerating protons from ionized hydrogen with a static electric field.

Next, the Radio-Frequency Quadrupole (RFQ) accelerates the protons to 750 keV, and bunches them into packets separated by 18.9 ns. These packets are colloquially called “buckets”. The bunching is a natural consequence of the 53 MHz operating frequency of the RFQ and is required for efficient through-put in the next stage. Next, the proton beam enters the “Linear Accelerator” (Linac), a linear chain of superconducting RF cavities. The linac accelerates the buckets of protons up to 400 MeV.

The Booster is a circular accelerator ring which, through multiple orbits around its 460 meter circumference, brings the protons up to 8 GeV. Each full circle of 8 GeV protons, grouped into 84 buckets, is injected into the Recycler. The recycler can hold 588 bunches in its full circumference. Typically, six Booster ‘batches’

are injected end-to-end, leaving the seventh vacant to serve as the “abort-gap”. In case of an emergency, the “abort-gap” is used to ramp up kicker magnets which will safely dump the stored beam into a beam-dump. Due to the abort-gap and small gaps between Booster batches, there are typically only 498 filled buckets.

Additionally, the Booster batches can be “slip-stacked”, that is, injected on top of one another to increase the instantaneous intensity of the beam. Once the recycler is full, it feeds the entire stored beam into the Main Injector in one turn. Maintaining the beam structure from the Recycler, the Main Injector accelerates the protons up to 120 GeV.

2.2.2 Slow-Spill Extraction

For the high-intensity neutrino experiments, the entire stored beam is delivered in one turn and the entire process is repeated as quickly as possible. However, for beam delivery to SeaQuest, the contents of the Main Injector is delivered slowly over five seconds with a technique called “slow-spill extraction”. To achieve this, an electro-magnetic septum is used to carefully scrape off a small fraction of the circulating beam. The scraped protons are transported to the SeaQuest experimental hall through a series of vacuum pipes, and dipole and quadrupole magnets. Over the course of the 5-second “spill”, which takes 369,000 “turns” of the Main Injector, the electromagnetic septum slowly extracts all of the protons.

On average, the slow-spill extracted beam delivered to SeaQuest has $\sim 10^4$ protons in each bucket. Therefore, the average instantaneous intensity of the beam is $\sim 2 \times 10^{12}$ protons per second, for a total of $\sim 1 \times 10^{13}$ protons per spill.

SeaQuest receives a 5-second spill approximately once per minute. In addition to allowing beam delivery to the other experiments operating at FNAL, the 55-second break between spills allows for target table motion, scaler readout, DAQ recovery, heat transport in the targets and beam-dump, and dissipation of beam-induced radiation in the hall.

2.2.3 Beam-Intensity Monitor

Unfortunately, the slow-spill extraction technique does not generate equal-occupancy buckets. In fact, the distribution of protons among the buckets is highly non-uniform. Most buckets contain fewer than 10^4 protons, and some buckets contain more than 10^5 protons. The non-uniformity of the beam poses a significant challenge to the trigger, DAQ, and reconstruction.

Due to the relatively poor position resolution of the hodoscopes, the trigger is easily fooled by high-occupancy hit patterns generated by combinatoric superposition of multiple uncorrelated tracks. As such, the trigger rate is highly sensitive to the instantaneous beam intensity, and the triggered events are dominated

by combinatorics even at modest intensities. As the beam intensity increases, the trigger rate increases like I^2 , or even faster.

The DAQ readout system is composed of many zero-suppressed TDCs, and the readout time is dependent on the number of recorded hits in each event. Since the DAQ is “dead” while reading out, the triggered high-intensity events can cause a significant reduction in the rate of data collection.

Finally, even if a triggered and recorded high-intensity, high-occupancy event contains a good Drell-Yan dimuon, it is very unlikely that the reconstruction algorithms can find it. The event reconstruction efficiency is highly dependent on beam intensity, dropping to $\sim 20\%$ efficiency at $\sim 10^5$ protons per bucket.

For all of the above reasons, SeaQuest designed, constructed and installed a Beam-Intensity Monitor (BIM) in the vacuum pipe just upstream of the targets. The BIM is able to resolve individual beam buckets, and it issues an inhibit to the DAQ when the beam intensity crosses a programmable threshold. When the intensity falls below the threshold, the inhibit is released, and the DAQ resumes accepting triggers. Therefore, the BIM allows SeaQuest to avoid the costly loss of live-time caused by reading out very high-intensity events.

The BIM consists of a Photo-Multiplier Tube (PMT), which measures the Čerenkov radiation induced by the beam protons passing through a short volume of atmospheric-pressure gas installed within the vacuum beam-pipe. Figure 2.3 depicts the BIM apparatus. Čerenkov radiation produced in the 80% Argon, 20% CO₂ gas mixture between the baffle and the mirror is reflected by the mirror into the PMT. Neutral-Density Filters installed in front of the PMT ensure a linear response across the required range of beam intensities, from 30 protons per bucket up to 3×10^6 protons per bucket.

The readout electronics for the BIM consists of a “QIE” (Charge Integrator and Encoder) circuit board, custom-designed by FNAL for use in the CMS experiment at CERN. The QIE digitizes and integrates the PMT signals, generating bucket-by-bucket intensities as well as the total intensity of the entire spill. For each spill, the spill-integrated intensity is used to calibrate the single-bucket intensities, using the more precise spill-integrated absolute intensity provided by a Secondary Emission Monitor (SEM) just upstream of the BIM. The QIE module also receives information about the DAQ deadtime, through the DAQ “busy” signal. Therefore, the QIE is able to provide:

- Spill-integrated intensity
- Integrated intensity while the inhibit is asserted
- Integrated intensity while the DAQ is dead (excluding buckets that are also inhibited)
- Intensities for each of the 33 buckets adjacent to, and including, the triggering bucket of each event




Figure 2.3: Čerenkov instrumentation package.

- A complete record of individual bucket intensities for the entire spill

Using the information recorded by the QIE, SeaQuest is able to calculate the number of “live” protons in each spill. This number is crucial for any absolute cross-section analyses and any rate-dependence corrections. Equation 2.1 demonstrates the live-proton calculation.

$$\text{liveProton} = I_{\text{SEM}} \left[\frac{I_{\text{BIM}} - I_{\text{inhibited}} - I_{\text{dead}}}{I_{\text{BIM}}} \right] \quad (2.1)$$

where I_{SEM} is the spill-integrated intensity from the SEM, I_{BIM} is the spill-integrated intensity from the BIM, $I_{\text{inhibited}}$ is the spill-integrated intensity for inhibited buckets measured by the BIM, and I_{dead} is the spill-integrated intensity for non-inhibited DAQ dead buckets measured by the BIM.

Position ID	Descriptor	Density [g/cm ³]	Length	Interaction Length	Spills/Cycle (% of spills)
1	H ₂	0.07065	50.8	0.06902	10 (43%)
2	Empty	–	–	0.0016	2 (9%)
3	D ₂	0.1617	50.8	0.1144	5 (22%)
4	None	–	–	0.0	2 (9%)
5	Iron	7.874	1.905	0.1135	1 (4%)
6	Carbon	1.802	3.322	0.0697	2 (9%)
7	Tungsten	19.30	0.953	0.0958	1 (4%)

Table 2.1: The seven SeaQuest targets. “Spill/Cycle” is a typical configuration. The exact spill cycle varied somewhat throughout SeaQuest’s data-taking periods.

2.3 Targets

The SeaQuest physics program requires the measurement of proton-induced Drell-Yan on targets with a wide range of atomic mass numbers. The flagship analysis, which will extract the proton sea’s \bar{d}/\bar{u} ratio, requires two nearly identical cryogenic liquid targets, one filled with deuterium (D_2), and the other filled with hydrogen (H_2). Various effects of Drell-Yan cross-section modification in heavy nuclei will utilize the D_2 target along with three solid targets, composed of carbon (C), iron (Fe), and tungsten (W). In addition to these five “physics” targets, there are two calibration targets. The first is an empty cryogenic flask, identical to the D_2 and H_2 flasks, but filled with vacuum. This target is referred to as “Empty Flask”, or just “Empty”. The second is a set of solid target holder clips, with no target disks installed. This is referred to as the “None” target. These two calibration targets are used to subtract the contribution of events generated by proton interactions in the flask metal, air, and any other materials in the target cave.

As shown in Figure 2.4, these seven targets are mounted on a motorized and remotely-controllable target table, which is able to accurately position one target directly in the beam line for each spill. The target table motion is fast enough to move from any target position to any other within the 55-second downtime between each spill. The targets are rotated continuously throughout data taking. Therefore, any potential systematic effects caused by slow, time-dependent drifting of beam intensity, beam quality, and detector response should cancel out in the analysis of target cross-section ratios.

2.4 Magnets

The SeaQuest spectrometer relies on two large dipole magnets, FMAG and KMAG, to enhance the acceptance of the physics events of interest, and to measure the momentum of detected muon tracks, respectively.

FMAG, depicted in Figure 2.5, is the more upstream of the two magnets, and consists of an aluminum “bedstead” coil embedded in a large stack of extremely high-purity iron blocks. The magnet has no air




Figure 2.4: Target Table schematic, top view.

gaps, and the 2000 A excitation current is able to maintain a nearly uniform 1.9 Tesla field within the fully-saturated iron. This field delivers an average p_T -kick of 3.07 GeV to the high-momentum muons which pass through the 160 cm by 43 cm central region.

FMAG serves four main functions: focusing magnet, sweeper magnet, beam-dump, and hadron absorber. The FMAG field focuses high- p_T muon tracks into the detector stations, thereby increasing the geometric acceptance for high-mass, opposite-charge muon pairs. The FMAG field also “sweeps” lower-momentum muons away from the detectors, lowering the acceptance of low-momentum backgrounds. The iron core of FMAG serves as a beam-dump, absorbing the 90% of protons that do not interact in the targets. These dump-interaction protons can also generate Drell-Yan, J/Ψ , and Ψ' dimuons within the spectrometer acceptance. The large recorded sample of “dump” dimuons will allow high-statistics analyses of p+Fe interactions. The iron of FMAG also serves as a secondary hadron absorber, stopping most secondary hadrons generated by beam-proton interactions within the target and dump materials. FMAG is able to absorb most pions before they generate muons via weak decay, significantly reducing the single-muon background.

At the intersection of the beam axis and the front face of FMAG, there is 5 cm diameter, 25 cm deep hole drilled into the iron. This hole moves the very radioactive dump interaction point downstream, away from the target cave. This helps reduce radiation in the target cave, into which SeaQuest personnel occasionally must make accesses for target maintenance.




Figure 2.5: Left: FMAG. Right: KMAG.

The downstream magnet, KMAG, is constructed of a set of coils with an open air core. Running with a current of 1600 A, KMAG produces a central field of 0.4 Tesla. This field gives an average momentum kick of 0.41 GeV to traversing high-momentum muons.

Situated between the first two tracking detector stations, KMAG allows for the determination of muon momenta by measuring the track deflection. KMAG's field is aligned parallel to FMAG's, so its relatively small p_T -kick slightly enhances the focusing and sweeping power of FMAG.

2.5 Tracking Detectors

SeaQuest employs three types of tracking detectors, used for recording hit positions along the paths of charged particles (muons) traversing the spectrometer. The hit positions are later reconstructed into particle trajectories. The trajectories are projected back to the target area, and for dimuon events, the projected momenta of the two muon tracks allow calculation of many properties of the interacting partons which generated the muon pair. The three types of detectors are hodoscopes, for fast triggering; wire chambers, for precise hit positions; and proportional tubes, for muon identification.

Detector Name	Paddle Width	Overlap [cm]	# paddles	Width×Height [cm×cm]	z-position [cm]
H1T	7.32	0.3175	23	162×70	666
H1B	7.32	0.3175	23	162×70	666
H1L	7.32	0.3175	20	79×140	653
H1R	7.32	0.3175	20	79×140	653
H2T	13.00	0.3175	16	203×152	1421
H2B	13.00	0.3175	16	203×152	1421
H2L	13.00	0.3175	19	132×241	1400
H2R	13.00	0.3175	19	132×241	1406
H3T	14.59	0.3175	16	228×168	1959
H3B	14.59	0.3175	16	228×168	1959
H4T	19.65	0.3175	16	305×183	2234
H4B	19.65	0.3175	16	305×183	2251
H4Y1L	23.48	0.3175	16	152×366	2130
H4Y1R	23.48	0.3175	16	152×366	2146
H4Y2L	23.48	0.3175	16	152×366	2200
H4Y2R	23.48	0.3175	16	152×366	2217

Table 2.2: Hodoscope plane specifications.

2.5.1 Hodoscopes

SeaQuest uses scintillating hodoscope paddle detectors in each of its four tracking stations. The physical specifications of the scintillator paddles are given in Table 2.2. The paddles are arranged into “picket fence”-style planes, with slight overlaps between adjacent paddles, to prevent small gaps in acceptance. Each paddle is wrapped with aluminum foil (as a reflective spacer), and an opaque Tyvek outer wrap. Each paddle is instrumented with a single PMT (except for Station 4, in which both ends of each hodoscope are instrumented), connected to the paddle by a plastic light-guide.

Ionizing particles passing through the paddles excite molecules within the material, which in turn release photons. These photons are mostly totally internally reflected within the paddle, and make their way to the light-guide, where they are transmitted to the PMT for detection. Within the PMT, the photons liberate electrons via the photo-electric effect [46] on the cathode. The liberated electrons are accelerated with static electric fields, and are multiplied via secondary emission through a series of ‘dynodes’. The voltage on the dynodes is maintained by a voltage-divider, housed in the base attached to the PMT. The large pulse of electrons resulting from incident photons at the cathode is transmitted to CAMAC discriminators, which apply a threshold to the analog pulse and generate a logical output for threshold-crossing pulses. The logical signals are then digitized by TDCs and used as inputs to both the FPGA- and NIM-based trigger systems.

The Station 1 and 2 hodoscopes use scintillator paddles and PMTs recycled from the HERMES experiment at DESY. Stations 1 and 2 each consists of two planes: the x-measuring plane, with vertically-oriented paddles; and the y-measuring plane, with horizontally oriented paddles. Each plane is split in the middle,

and the PMTs are attached at the outer end of each paddle.

The Station 1 and 2 hodoscopes must operate at very high intensity since many out-of-acceptance tracks, which will miss Stations 3 and 4, still pass through Stations 1 and 2. During the 2012 Commissioning Run, a “sag” in the hit-rates of PMTs exposed to a high-intensity of particle tracks suggested that the passive voltage dividers used in the Station 1 and 2 PMT bases were not capable of maintaining the correct dynode voltages. Therefore, new “transistorized” bases were designed, tested, constructed, and installed before the next data-taking period. The transistorized PMT base upgrade is described in detail in the Ph.D. dissertation of Bryan Dannowitz (Pending).

Stations 3 and 4 use new Eljen EJ-200 scintillator paddles, and a combination of recycled PMTs and new Hamamatsu PMTs. Due to their significant length (> 150 cm), the Station 4 hodoscopes are instrumented with a PMT on each end. This effectively halves the maximum light-travel time within the paddles. The signals from the two PMTs for each paddle were originally combined using CAMAC “mean-timer” modules. These modules detect time-separated signals in two input channels, and output a single signal with a timing exactly halfway between the timing of the two inputs. However, during the commissioning run, it was discovered that the mean-timers could erroneously combine signals from two different beam buckets, generating spurious output signals which correspond to no real hit. Therefore, the hodoscope signals now bypass the mean-timers and are combined with a simple logical “OR” in the first stage of the FPGA-based trigger.

2.5.2 Wire Chambers

SeaQuest uses Drift Chambers to precisely measure the positions of muons in each tracking station as they travel the length of the spectrometer. The Drift Chambers each consist of a contained volume of specialized gas, through which six planes of wires are held taut. The wires are organized three “views”, with one view referring to two planes with the same orientation and number of wires, but slightly separated in z , and offset transversely by one-half the wire spacing. The “primed” planes help the reconstruction to resolve the left-right ambiguity in the track position relative to a fired wire. The three views use the standard “X, U, V” configuration, where the “X” wires are vertical, and the “U” and “V” wires are tilted away from the y -axis, remaining in the xy -plane, by -14° and 14° respectively. Therefore, each chamber consists of six planes, labeled X, X', U, U', V, and V'.

Table 2.3 lists some details of the chambers used by SeaQuest. The chambers are all labeled DN , where N specifies the tracking station to which the chamber belongs. Since Station 3 uses two separate chambers, one for the top-half of the spectrometer, and one for the bottom-half, these chambers are further

Chamber	View	# wires	cell width [cm]	Width×Height [cm×cm]	z-position [cm]
D1.1	X	160	0.64	102×122	617
	U, V	201	0.64	101×122	±20
D1.2	X	320	0.50	153×137	617
	U, V	384	0.50	153×137	±1.2
D2	X	112	2.1	233×264	1347
	U, V	128	2.0	233×264	±25
D3p	X	116	2.0	232×166	1931
	U, V	134	2.0	268×166	±6
D3m.1	X	176	1.0	179×168	1879
	U, V	208	1.0	171×163	±19
D3m.2	X	116	2.0	232×166	1895
	U, V	134	2.0	268×166	±6

Table 2.3: Wire Chamber specifications. The “primed” and unprimed planes in each view are slightly separated in z-position. U- and V-view z-positions are given relative to the associated z-X-view.

designated as “plus” (D3p), and “minus” (D3m). Upgraded chambers are labeled with an additional index after a decimal point. Chambers *D1.1*, *D2*, and *D3m.1* were built for previous Drell-Yan experiments at FNAL (E605 and E866) and are being re-used for SeaQuest. Significant work was required to return these chambers to working order, including restringing $\sim 30\%$ of sense wires. *D3p* was built for SeaQuest and was installed before data-taking began. *D1.1* and *D3m.2* were also built new for SeaQuest, and were installed after SeaQuest received first beam. Table 2.4 lists the time periods of use of each different configuration of chambers. The upgraded chambers are all wider than their refurbished stand-ins, meaning that SeaQuest’s sensitivity to the highly-desired high- x_2 kinematic region has increased with each chamber upgrade.

For all analyzable data, the gas mixture used for most chambers is Argon:Methane:CF₄ (88% : 8% : 4%). This gas has a drift velocity of about $20\mu\text{m}/\text{ns}$. The exception is D1.2, which uses a special “fast” gas mixture: Argon:CF₄:Isobutane:Methylal (68% : 16% : 13% : 3%). The fast gas has drift velocity of about $50\mu\text{m}/\text{ns}$. The fast gas is required in D1.2 for optimal performance due to the high hit-rate at Station 1.

The momentum resolution of reconstructed muons is limited by multiple scattering in FMAG. To keep the contribution of the position resolution to less than 10% of the total width of the dimuon invariant mass resolution, the chambers are required to achieve at least $400\mu\text{m}$ position resolution. This contributes a

Run	Dates	St. 1	St. 2	St. 3 Top	St. 3 Bottom
1	2012 Mar. – 2012 Apr.	D1.1	D.2	D3p	D3m.1
2	2013 Nov. – 2014 Aug.	D1.1	D.2	D3p	D3m.2
3	2014 Nov. – 2015 May	D1.1	D.2	D3p	D3m.2
3	2015 Jun. – 2015 Jul.	D1.2	D.2	D3p	D3m.2
4	2015 Sep. –	D1.2	D.2	D3p	D3m.2

Table 2.4: Run history and installed chamber combinations.




Figure 2.6: Left: xz-plane view of prop tubes. Right: yz-plane view of prop tubes.

momentum resolution of $\Delta p/p = 0.03$.

The wire chambers are all instrumented with custom amplifier-discriminator integrated circuit boards called “ASDQ cards”. These cards were designed at University of Pennsylvania for the CDF experiment at CERN. The ASDQ cards perform the required amplification and discrimination on the analog signals from the sense wires. The differential outputs of the ASDQ cards are fed into “Level-Shifter Boards” (LSBs), which convert the signals to standard LVDS. These LVDS signals are then sent to the custom SeaQuest TDC modules for digitization and readout by the DAQ.

2.5.3 Proportional Tubes

In lieu of a wire chamber, Station 4 utilizes four planes of proportional tubes (prop tubes) to measure precise track positions. Each plane is composed of 9 proptube modules, and each module contains 16 individual prop tubes arranged in two staggered rows of 8 tubes. Each prop tube plane acts similarly to the primed-unprimed plane pairs in the wire chambers, with the staggered second row resolving the left-right ambiguity in track positions. The prop tubes are oriented similarly to the hodoscopes, with two x-measuring planes and two y-measuring planes. Figure 2.6 depicts a diagram of the prop tube arrangement.

Each individual tube is composed of a 2 inch diameter, hollow aluminum tube with a wall thickness of

Plane Name	# Modules	# tubes per module	Width×Height [cm×cm]	z-position (second layer) [cm]
P1H	9	16	368×368	2099 (+4)
P1V	9	16	368×368	2175 (+4)
P2H	9	16	368×368	2367 (+4)
P2V	9	16	368×368	2389 (+4)

Table 2.5: Proportional tube detector plane specifications.

1/16 inch. An gold-plated, $40\ \mu\text{m}$ tungsten anode wire is held suspended taut down the center of each tube. The prop tubes use the same “slow” gas mixture as most of the chambers.

The reconstructed “tracklets” in the prop tubes serve as muon identification. Station 4, including the prop tubes, are installed behind a 1m -thick iron wall. By matching a track formed from hits in the Station 2 and Station 3 chambers to a “tracklet” formed from the prop tube hits, we can guarantee that the track was caused by a high-energy muon. Hadrons would be absorbed in the iron wall, and high-energy electrons would suffer much greater scattering, observable as a large difference in track angle before and after the iron wall.

2.6 Trigger

My major hardware contribution to SeaQuest was my work on the trigger system. After assisting with the initial installation and commissioning, I took over full responsibility for the FPGA-based trigger system in late 2012. Since that time, I have made major hardware and firmware upgrades to improve the effectiveness and reliability of the trigger. I have also written peripheral software to perform maintenance and commissioning tasks and monitor the trigger performance. Chapter 3 describes the details of SeaQuest’s novel trigger system in great detail.

2.7 Data Acquisition

SeaQuest employs four separate Data Acquisition (DAQ) systems, to reliably record four categories of experimentally-relevant data. These four DAQ systems are called “Main DAQ”, “Scaler DAQ”, “Beam DAQ”, and “Slow Control”. Main DAQ records information from the tracking detectors and FPGA trigger system, at the event level. Scaler DAQ records the four scaler readouts of detector and diagnostic systems, three once per spill, and one on a 7.5kHz clock. Beam DAQ records data and diagnostic information from the Čerenkov BIM electronics. Slow Control records all other diagnostic readouts, such as target position and magnet currents, that only need to be recorded once per spill.

2.7.1 Main DAQ

SeaQuest’s Main DAQ is built around the CODA (CEBAF¹ On-line Data Acquisition) software, developed at Jefferson Lab. CODA runs on a Scientific Linux PC in the SeaQuest control room, and communicates with the numerous “front end” VME crates installed in the SeaQuest experimental hall via private network.

¹Continuous Electron Beam Accelerator Facility

Each VME crate houses a single board processor (also known as a “Read-Out Controller”, or “ROC”), a trigger interface card, and many TDCs. One special VME crate, the “Trigger Supervisor”, orchestrates the starting, stopping, and reading out of all other VME crates, communicating through their trigger interface cards. The TDCs record hit and timing information from the tracking detectors. The ROCs handle the transmission of the recorded information back to CODA when a trigger is received.

The Trigger Supervisor (TS) can monitor 12 trigger signal inputs. The first four are prescalable by up to 24 bits. The next four are prescalable up to 16 bits, and the final four cannot be prescaled. The 12 inputs are assigned to: (Matrix 1–5, NIM 1–4, flush, BOS, EOS). When the TS accepts an input trigger, it sends an “accepted trigger” signal to the other VME crates, through their trigger interface cards. The VMEs will immediately “stop” their TDCs, and begin processing the collected data.

The TS also handles the BIM inhibit functionality. If the BIM detects a bucket over the inhibit threshold, it raises the “inhibit” signal, which is sent directly to the TS. While the inhibit signal is raised, the TS will accept no triggers. The BIM lowers the inhibit signal after a (remotely adjustable) duration has elapsed since the latest over-threshold bucket.

The deadtime of the Main DAQ is considerable, and its management is one of the primary considerations in the optimization of beam intensity, inhibit threshold, and trigger acceptance. Each TDC requires $10\ \mu\text{s}$ to “stop” after receiving the “trigger accepted” signal. Furthermore, each TDC requires a fixed, $32\ \mu\text{s}$ copy-in-progress time, even if the TDC has no stored hits. While all the VME crates can process and read-out in parallel, the TDCs within each VME must be read-out by the ROC in series. The slowest VME, which determines the DAQ deadtime, holds 7 TDCs and needs $150\ \mu\text{s}$, on average, between receiving the “trigger accepted” signal and completing read-out.

2.7.2 Scaler DAQ

The Scaler DAQ records information used for identifying and diagnosing large unexpected changes in trigger and detector performance. The Scaler DAQ consists of one VME crate, with one MVME5500 CPU and four scaler cards, and is controlled and readout by a CODA system which is independent of the Main DAQ CODA.

One scaler is triggered by the coincidence of a 7.5 kHz gate clock and the beam spill gate. This scaler samples rates of discriminated signals from the hodoscope planes at 7.5 kHz. One channel of this scaler is assigned to each hodoscope half-plane.

The other three scaler cards are triggered by the Beginning Of Spill (BOS) and End Of Spill (EOS) signals, and therefore integrate counts over each spill. One card assigns channels to each Main DAQ trigger

signal, including the Raw, Prescaled, and Accepted state of each trigger. Another card records hit counts from the various Beam Intensity instrumentation packages upstream of SeaQuest (but not SeaQuest’s BIM). The last scaler card records the spill-integrated hit rates in the hodoscope half-planes.

Custom software analyzes the Scaler DAQ data in real time, and displays the results on a large monitor mounted in the SeaQuest control room.

2.7.3 Beam DAQ

The Beam DAQ is a third independent system responsible for recording data collected by the Čerenkov BIM monitor, and its QIE board. The Beam DAQ is operated by custom C++ software which handles the control and readout of the BIM. Due to the single-bucket resolution of the BIM, the Beam DAQ can calculate the 53 MHz duty factor:

$$DF_{53} = \frac{\left(\frac{1}{N} \sum_i I\right)^2}{\frac{1}{N} \sum_i (I^2)} = \frac{\langle I \rangle^2}{\langle I^2 \rangle} \quad (2.2)$$

where i enumerates all the beam buckets in a spill. DF_{53} is the primary parameter used by the accelerator division when tuning the beam to SeaQuest.

As described in Section 2.2.3, the BIM records the intensity of each beam bucket, sum of intensity blocked by the inhibit, sum of intensity blocked by DAQ deadtime (and not inhibited), and $\langle I \rangle$ and $\langle I^2 \rangle$. The QIE board records and stores the “raw” data throughout each spill. The data is read-out at the end of each spill. The data volume for a single spill is quite large, $\sim 300MB$, so special effort must be made to read out the data in the 55 seconds between spills. The Beam DAQ readout three threads to read data from three ethernet chips mounted on the QIE board. Additionally, eight threads are used for live analysis of the raw data. Analyzed results are displayed in the SeaQuest control room alongside the Scaler DAQ information.

The Beam DAQ also outputs the individual-bucket intensity measurements surrounding each triggered event to the Main DAQ, via a latch TDC.

2.7.4 Slow Control

‘Slow Control’ is the general term applied to a host of auxillary scripts recording information from all parts of the spectrometer. These scripts are all semi-autonomous, and are coordinated by interfacing with an Experimental Physics and Industrial Control System (EPICS) server running on the SeaQuest target computer. The Slow Control system serves three purposes: synchronization, auxillary data recording, and status monitoring.

The independent nature of the SeaQuest DAQ systems necessitates a robust synchronization scheme, to

ensure that offline analysis can aggregate data from all DAQ systems correctly. This is achieved by assigning a unique, eternally-incrementing identification number to each spill (spillID). The spillID is managed by the EPICS server and all other DAQs reference the master spillID when recording their own data.

The Slow Control system also records various categories of auxiliary experimental information. From the Fermilab Accelerator Controls Network (ACNET), the Slow Control records beam intensity and quality parameters, accelerator configuration, and the status of SeaQuest’s magnets. From the SeaQuest target control system, the Slow Control records the target position and rotation pattern, and the temperatures and pressures of the cryogenic targets. A Keithley monitor provides the Slow Control system with data from temperature, pressure, and humidity sensors placed throughout the SeaQuest experimental hall. The Slow Control system submits these collected data to both the Main DAQ and Scaler DAQ systems.

The Slow Control system also performs continuous monitoring of recorded values which are critical for successful data taking. An alarm system alerts the shift crew if any of the run-critical values fail the checks. Monitored values include accelerator beam delivery status, whether or not the four DAQ systems are live and responding, and the available disk space.

2.8 Data Curation

The CODA-based mainDAQ outputs raw hexadecimal datafiles, one for each “run”. MainDAQ runs are typically allowed to accumulate ~ 1 GB of raw data, with a duration of ~ 1 hr. These raw data files are stored on one of SeaQuest’s in-house servers and are immediately backed-up using the tape storage service provided by FNAL Computing.

Custom C++ software was written to “decode” the raw hexadecimal CODA output. The MainDAQ data consists almost exclusively of TDC data from the readouts of the hodoscopes, chambers, and prop tubes. (This includes the duplicate read-out of the x-measuring hodoscopes through the trigger system.) The only additional data is the bucket-by-bucket intensity information from the BIM, which is inserted into the MainDAQ CODA file for each event. The TDC hits from the various readouts belong to “events”; one event is created each time the DAQ accepts a trigger. The events are independently associated with a “run” and a “spill”. Runs are ~ 1 hr long periods of data taking, during which the DAQ ran uninterrupted. Spills are defined by the ‘spillID’, which increments at each slow spill extraction.

The decoder is responsible for reading the raw CODA files and filling MySQL database tables with the decoded information. A new schema is used for each run. The decoder also references static information associated with each detector/channel from text files. This information consists of:

- Detector names and element IDs of each hodoscope, chamber, and prop tube channel
- In-time windows for each channel
- Drift-velocity information for each chamber and prop tube channel
- List of “roads” loaded in the v1495 trigger

The raw data only contains Spill ID, Event ID, ROC ID, channel ID, and TDC time for each hit. Into the database, the decoder injects these quantities: Run ID, Spill ID, Event ID, Detector Name, Element ID. The decoder then calculates a number of additional quantities: In-Time Flag, Hodoscope Masked Flag, Drift Time, Drift Distance, and Data Quality.

In-Time Flag An in-time flag is calculated, based on the in-time window for each channel. A more complicated in-time flag calculation is required for perfect simulation of the internal trigger in-time window. This calculation is described in Section 3.5.2.

Hodoscope Masking Flag Due to the relatively slow detection time of the drift chambers, hits cannot be accurately identified with a single beam bucket. Some adjacent-bucket hits can be removed from consideration in the analysis by exploiting the much finer time resolution of the hodoscopes, by assuming that an in-time muon track would fire both the chamber wire and a hodoscope paddle as it passed through the detector station. Hits from muons in adjacent buckets would not necessarily have an associated in-time hodoscope hit. Each in-time chamber is checked to see if it lands “behind” and in-time hodoscope hit. Those that do are flagged as “masked”, and will be considered by the reconstruction and analysis.

Drift Time and Drift Distance For chamber hits, the difference between the TDC time of the hit and the time of the triggered beam bucket (t_0), is the Drift Time. Assuming the detected charge was deposited in the gas at the time of the arrival of the triggering beam bucket, the drift time is simply the time that transpired before the charge was collected on the sense wires. Using a parameterization of the known drift velocity properties of the chambers, the drift time is used to calculate a drift distance. This drift distance then represents the distance between the muon path and the fired wire at the point of closest approach.

Data Quality The data quality field is a bit pattern used to store information regarding various errors that can occur in the raw CODA data and decoding calculations. If a hit suffers none of the known issues, the data quality is zero.

In addition to the expanded and calculated information for each hit, the decoder creates additional tables to store:

- Spill-Level information about the target, scalers, and beam monitors

- “Roads” loaded into the trigger during each run
- “Fired-Roads” for each event, calculated from the hit patterns in the v1495 TDC readout
- Overall occupancy of each detector station in each event

The MySQL scheme has worked magnificently for storage of decoded data, storage of reconstructed data, and ease of access to the reconstructed data by analyzers. The reconstructed data is duplicated across all SeaQuest servers and is available to all analyzers at all times from multiple possible sources. Analyzers can interact with the databases with nearly any favorite programming language, as most modern languages have well-developed MySQL APIs.

Chapter 3

The SeaQuest Trigger System

3.1 Overview

The SeaQuest experiment uses the high-intensity 120 GeV proton beam from Fermilab's Main Injector to measure muon pairs generated through the relatively rare Drell-Yan process. The Drell-Yan process is the dominant contributor to opposite-sign muon pairs in the 4-10 GeV mass range. However, considering muons individually, the Drell-Yan process is only a very small contributor to the total muon flux through the experiment. Most muons are generated through pion decay and are the results of other, higher cross-section hadronic processes. Furthermore, the primary SeaQuest physics goals will use data from proton beam interactions in the targets, which are all $\sim 10\%$ targets. The remaining $\sim 90\%$ of the beam interacts in the solid iron core of the sweeper magnet, FMAG, which also acts as a beam dump. Any high-momentum muons produced in the beam's interaction in the dump, whether generated by Drell-Yan, quarkonia, or single-muon backgrounds, will also pass through the detectors.

For this reason, SeaQuest requires a carefully tuned, live trigger system which can preferentially select those events containing a true, Drell-Yan, target dimuon. Four planes of hodoscopes, consisting of plastic scintillator paddles and PMTs, record the coarse track position. The trigger system takes the digitized hodoscope hits as inputs. The precise track position is measured by three stations of drift chambers and one station of proportional tubes. The hodoscope+trigger system must be fast enough to generate an output signal, for a particular event, before the drift chambers read out the hits corresponding to that bunch of beam protons. This constraint is $\sim 1\mu s$. The paddle-width, which determines the hodoscope position resolution, is limited by the number of input channels available on the chosen hardware, the CAEN v1495 VME module. Therefore, the hodoscopes provide only a very coarse measurement of the track 3-momenta.

Due to the coarse-grained track momentum resolution and the high rate of single-muons, the trigger-fired events are dominated by spurious combinations of hits that do not correspond to a true dimuon. The DAQ deadtime is non-negligible; therefore, maximization of "good data on tape" requires that we carefully balance the beam-intensity, trigger acceptance, and DAQ livetime. A critical new tool in this balancing act was the

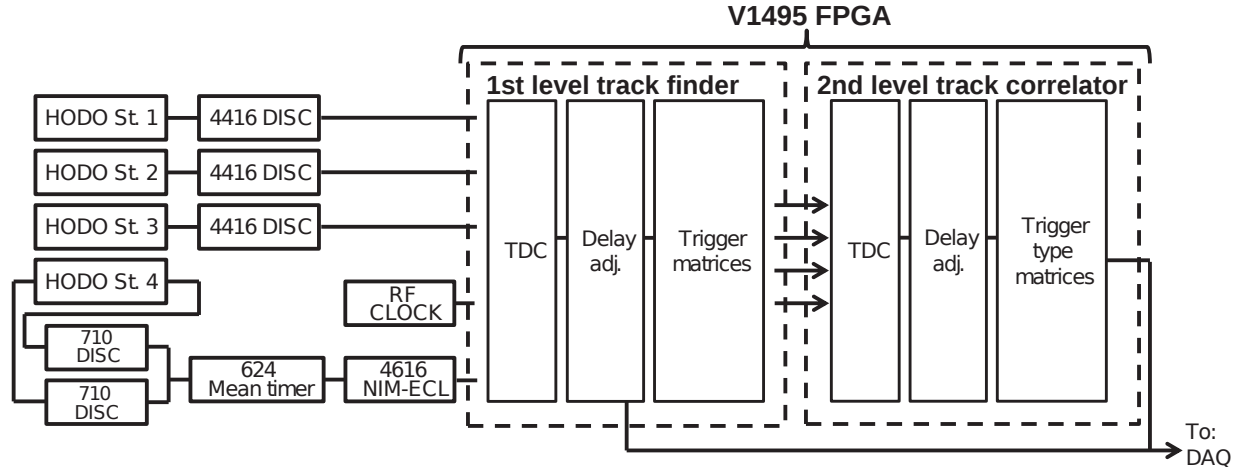


Figure 3.1: SeaQuest v1495 Trigger System initial design schematic [47]

addition of the “Beam Monitor Veto”, which allows the experiment to veto triggers with single beam-pulse granularity based on an upstream measurement of the beam intensity alone.

This chapter is organized as follows. First, I will explain the details of the initial trigger firmware design. In the next section, I describe the extensive and crucial series of improvements implemented during my tenure as trigger system leader. Then I describe the various peripheral software tools required for repeatable and smooth operation of the trigger system. Next, I evaluate the performance of the trigger system thus far in terms of signal acceptance, background rejection, and internal consistency. Relatedly, I describe the few trigger-related issues that significantly affected data-taking. Finally, I describe a number of ideas for additional trigger system upgrades which have not (yet) been implemented.

3.2 Initial Design and Implementation

The contents of this section have been published, with more detail, by Shiu et al. [47].

3.2.1 Overall Structure

The original design of the SeaQuest trigger system consisted of two “levels” of v1495 modules. In Level 1, four v1495 modules take inputs from the upper x-measuring, lower x-measuring, upper y-measuring, and lower y-measuring hodoscopes, respectively. These four independent groups of hodoscope elements are referred to as the four “quadrants” of the trigger input. The firmware-implemented TDC, delay adjustment, and “trigger matrix” decision logic in Level 1 are responsible for identifying candidate tracks from the hodoscope



Figure 3.2: Station 3 x-measuring hodoscopes, in front of the iron wall [47]

hit patterns in each event, separately in each “quadrant”. Information about the found track candidates is encoded and output from each Level 1 module on a 32-channel LVDS cable. Level 2 takes the 4×32 channels of Level 1 output as input, and uses the same firmware-implemented TDC, delay adjustment, and “trigger matrix” machinery to form pairs of individual track candidates that satisfy certain “di-track” criteria. A schematic of the high-level design of the original FPGA-based trigger implementation is shown in Figure 3.1.

Although the hardware for the y-measuring (non-bend plane) part of the trigger system has been installed since the beginning, SeaQuest has not implemented an “analyzable physics” trigger which utilizes the non-bend plane hodoscope information in the trigger decision. The St. 3 x-measuring hodoscopes (H3T and H3B), hanging in front of the iron wall, are depicted in Figure 3.2.

The module chosen for the SeaQuest trigger is the CAEN v1495 6U VMEbus module. The heart of this module is the Altera EP1C20F400C6 Field-Programmable Gate Array (FPGA). Custom firmware, described in Shiu et al.[47], implements the TDC, delay adjustment, and “trigger matrix” features required for SeaQuest. Figure 3.3 is an annotated photo of one v1495 module. The v1495s are controlled and read-out through the VME interface. The Bridge FPGA uses fixed firmware to operate all of the other on-




Figure 3.3: CAEN v1495 VME module [47]

board components, primarily communication through the VME interface. The custom-designed SeaQuest trigger firmware is loaded into the User FPGA, which performs the three main functions described above. ‘Mezzanine’ cards are additional small Integrated Circuit Boards which add extra input and/or output channels to the v1495s.

The TDC is responsible for measuring and digitizing the leading edge timing of discriminated hodoscope signals. Then, the delay adjustment step aligns the up-to-96 input channels in each v1495, such that any known timing differences between hodoscopes are corrected. Finally, the digitized, aligned hits are sent to the trigger logic pipeline, where the “look-up-table” trigger matrix is applied, generating outputs for every matrix-matching input pattern.

This design is common to the Level 1 and Level 2 firmware designs. The only difference between Level 1 and Level 2 are the internal arrangement of input channels and the choice of trigger matrix logic. The following three subsections describe each of the major components of the firmware in more detail.

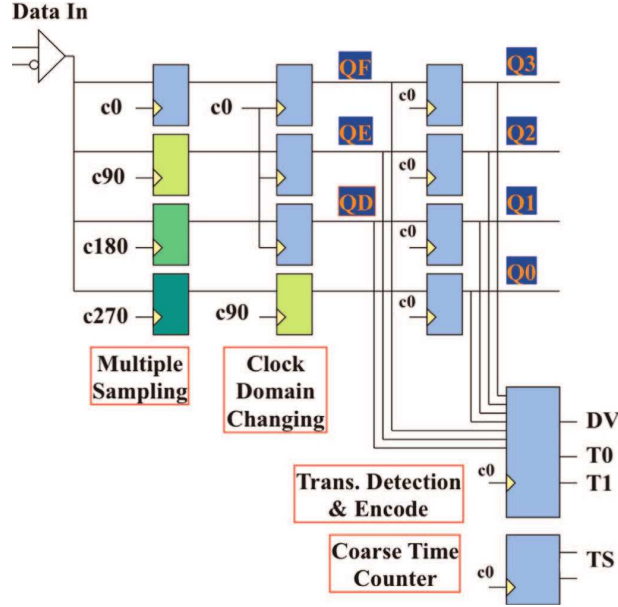


Figure 3.4: Four-phase sampling schematic in v1495 firmware TDC implementation [47]

3.2.2 TDC Block

The first part of the custom v1495 firmware is the Time-to-Digital Converter (TDC) block. The TDC block is identical in the Level 1 and Level 2 firmware designs. The purpose of the TDC is to measure, digitize and record the time of the leading edge of discriminated hodoscope signals. The on-board clock is 40 MHz, and a PLL is employed to generate a 250 MHz fast clock signal, but Seaquest requires ~ 1 ns resolution on the hodoscope timing. An effective 1 GHz clock is achieved using a 4-phase sampling technique, shown schematically in Figure 3.4. The 250 MHz clock is used to generate four 250 MHz clocks, with equally-spaced phases. Thus, these clocks are called c0, c90, c180, and c270, for the phase, in degrees, relative to the original clock. Using these four clocks to sample the input channels, an effective 1 GHz sampling is achieved, and the 1 ns TDC resolution requirement is met.

3.2.3 Delay Adjustment Pipeline

After digitization in the TDC, the hits are pipelined through blocks of RAM. The adjustable, channel-by-channel delays are applied to the data in the pipeline. Furthermore, the pipelined hits can be read-out when the DAQ is triggered.

Figure 3.5 schematically depicts an example of hits in the pipeline. The hits are binned into 16 ns bins, and the fine timing (1 ns resolution) for each hit is also stored in memory. Note that each 16 ns bin can only

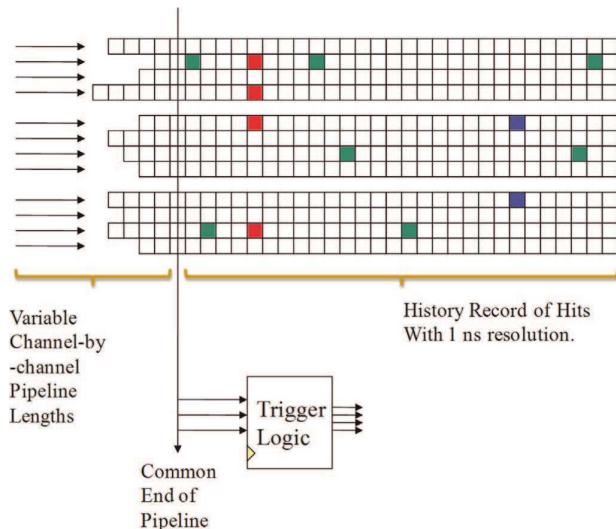


Figure 3.5: Digitized hit TDC pipeline [47]

hold at most one hit. If there are multiple hits on a single channel within one 16 ns bin, the pipeline keeps only the latest hit.

The pipeline also acts as hit storage. If the Trigger Supervisor accepts a trigger, it sends a “stop” signal to all the TDCs, including the trigger system v1495s. When a hit on the “stop” channel is digitized, the pipeline is halted, and a section of the pipeline is read out via the VME interface, to be recorded by the DAQ. Both the position and width of the read-out pipeline section are user-selectable. These settings are stored in control registers, and updated through the VME interface whenever the DAQ begins a run. The readout is zero-suppressed, and the zero-suppressed data is stored in a buffer that can hold 256 hits. Therefore, for very high-occupancy events, it is possible to overflow the buffer during zero-suppression, and some hits from the pipeline are lost. In this case, only the 256 latest hits are read out. For reasonably tight choices of readout width, any events that overflow the buffer are certainly not analyzable.

3.2.4 Trigger Matrix

The trigger matrix part of the SeaQuest v1495 trigger firmware design differs somewhat between Level 1 and Level 2. While both utilize a look-up-table, applied to the pattern of hits associated with each beam bucket, the details of the look-up-tables differ significantly.

One of the 96 internal, pipelined channels is reserved for the “RF-Clock”, the signal distributed by the Main Injector to indicate the timing of each beam bucket. A window, of adjustable width, is applied to the pipelined hodoscope signal channels for each hit in the RF-Clock channel. All hodoscope hits within each “RF-window” are considered to have come from the same beam bucket, the same “event”, and are grouped

together when applied to the look-up-table.

The look-up-table logic itself is pipelined into many steps, nine in Level 1 and five in Level 2. At each stage of this pipeline, gate elements combine up to 4 inputs each with simple logical operators. An effectively arbitrary number of gate elements can participate in each stage. By “AND”-ing groups of four outputs from each stage together into an input in the next stage, the “bandwidth” of the pipeline is reduced by a factor of four at each step. At the end of the pipeline, the results of the matrix logic are reduced enough for transmission via the 32-channel output mezzanine board. This scheme can, in principle, accommodate very complex logical manipulations in the pipelined matrix. However, the logic used for the SeaQuest Level 1 and Level 2 matrices is actually quite simple.

The Level 1 logic is tasked with identifying candidate tracks from combinations of four hodoscope hits, one from each detector station. One Level 1 board is assigned to the upper-half bend plane hodoscopes; another is assigned to the lower-half bend plane hodoscopes. The non-bend plane hodoscopes are not (yet) used in the trigger decision. The Level 1 look-up-tables are loaded with “trigger roads”. These roads are combinations of four hodoscope hits, one from each detector station, which have been identified as likely combinations to fire due to the passage of a muon which belongs to a Drell-Yan muon pair. One such road might look like: (8, 6, 5, 6). This road fires when a single beam bucket fires paddle eight in H1, paddle six in H2, paddle five in H3, and paddle six in H4. In each Level 1 board, the roads are organized into two distinct groups: those identified (by the Monte Carlo) as positive muon roads, and those identified as negative muon roads.

The Level 1 pipeline logic begins by “AND”-ing the hodoscope inputs into all loaded road combinations. The remaining eight pipeline steps simply combine all of the fired road combinations into a smaller number of output bits, to be sent to Level 2. Throughout the pipeline, the positive and negative roads are kept distinct, so that the Level 2 logic can distinguish the charge of each fired road.

The roads are also binned by average p_X . This was intended to allow Level 2 to make cuts on the rough mass of candidate dimuon pairs. However, the p_X resolution of the hodoscopes is too poor, and the Level 2 mass cut could not be applied effectively, so this scheme was not implemented for the analyzable data.

The Level 2 logic receives inputs specifying fired candidate roads, specifying the charge and approximate p_X of each road, along with whether the road was on the upper or lower half of the spectrometer. The Level 2 logic combines these inputs in pairs in the first step of its logic pipeline. Valid pairs for firing the main dimuon trigger must have opposite sign and come from different vertical “sides” of the spectrometer. That is, one from the upper half and one from the lower half. Table 3.1 lists the requirements for the five Level 2 output triggers used for collecting all of SeaQuest’s analyzable data.

Name	Side	Charge	p_X Requirements	Notes
Matrix 1	TB/BT	+/-/+	None	Main physics trigger
Matrix 2	TT/BB	+/-/+	None	Same-Side trigger (unused)
Matrix 3	TB/BT	++/- -	None	Like-Charge trigger (background estimation)
Matrix 4	T/B	+/-	None	All singles trigger
Matrix 5	T/B	+/-	$p_X > 3 \text{ GeV}$	High- p_T singles trigger

Table 3.1: The five triggers output by the v1495 trigger system. Matrix 2 was turned off (with a high prescale factor) because it was overwhelmingly dominated by background. Matrix 3 is used to estimate the rate of combinatoric backgrounds in the Drell-Yan sample.

The details of the generation of sets of trigger roads from Monte Carlo are described in more detail in Section 3.4.2.

3.3 Trigger System Upgrades

During the long beam shutdown period in 2012-2013, we took the opportunity to greatly improve the SeaQuest Trigger System. We added four CAEN v1495s, dubbed “Level 0”, to facilitate run-time monitoring and rigorous *in-situ* performance testing. We also added clip-lines to all hodoscope PMTs and removed the Station 4 Hodoscope Mean-Timers from the trigger-signal path. We added firmware design features required for the special functions of Level 0. Rigorous bench and *in-situ* testing revealed a number of bugs, which we then fixed. We changed the clock driving all logic operations to the external MI RF-Clock, from individual, internal, on-board clocks. We developed peripheral software, used for generating firmware, fine-tuning the timing, and evaluating performance.

3.3.1 Level 0 v1495 Modules

The trigger system is composed of hodoscope arrays, associated electronics, and nine v1495 boards. The upgraded trigger system schematic is depicted in Figure 3.6. Shiuan-Hal Shiu has described the preliminary hardware setup of the trigger system in a previous publication [47]. The remainder of this subsection describes changes to the hardware configuration after the 2012 commissioning run.

The largest hardware change during the long 2012-2013 shutdown was the addition of four new CAEN v1495 modules. During the 2012 commissioning run, we were frustrated by two competing, incompatible desires for the behavior of the trigger system. One, we wanted to read-out the on-board TDCs for every recorded event, in order to carefully monitor the behavior of the trigger logic. Two, we wanted to record the total number of times the v1495 trigger system fired each output trigger during each spill. Due to the basic design of the firmware, the trigger system must stop processing inputs while the on-board TDC is being

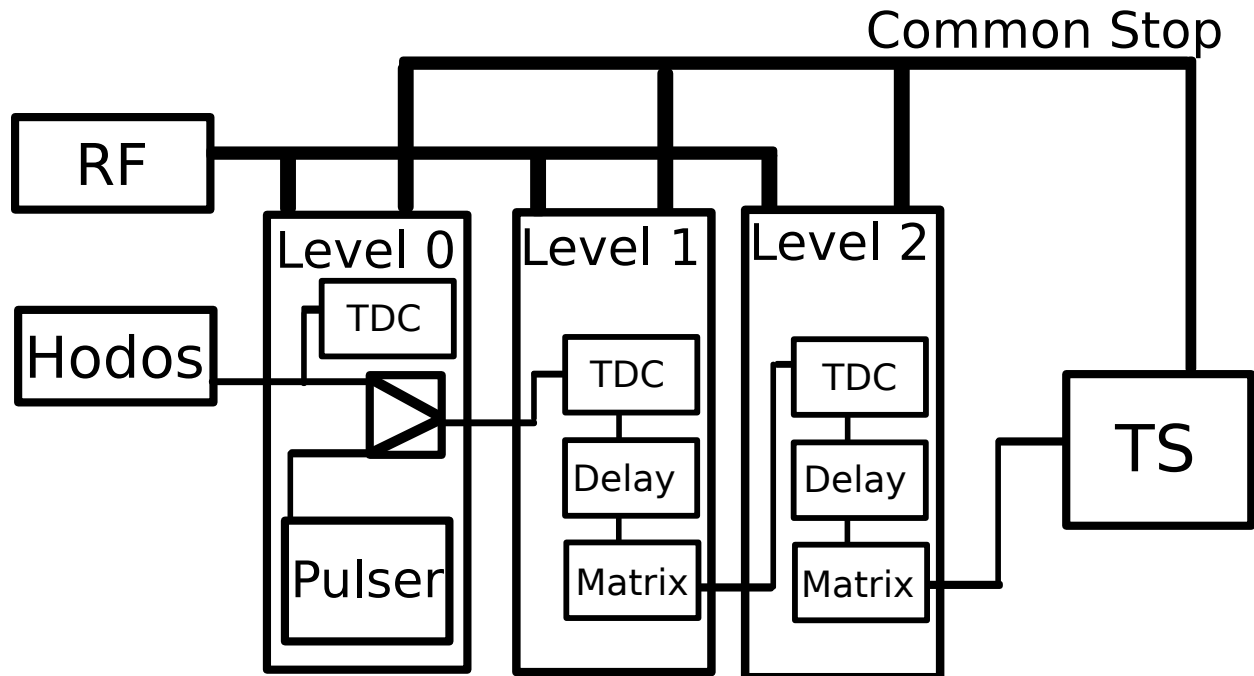


Figure 3.6: Trigger System Schematic with addition of Level 0. Level 0 output can choose between pass-through hodoscope signals or pulser signals. Level 0 copies the on-board TDC functionality of Level 1 and Level 2.

read-out by the DAQ. Therefore, if the on-board TDCs are being read-out for each event, the total number of output triggers is severely reduced, due to the DAQ deadtime. To solve this issue for future data-taking, we added four new v1495 modules and dubbed them “Level 0”. Each new module receives the input signals that one Level 1 module previously received from the discriminators. Because the Level 0 modules are doing no trigger logic processing, they are able to simultaneously read out TDC data for each event, **and** pass the input signals directly on to the Level 1 modules. Therefore, the TDCs on the Level 1 modules can remain disabled, allowing the Level 1 logic to continue generating outputs even when the DAQ has stopped to read out.

To serve their main purpose, the Level 0 modules require firmware that is only slightly modified from the Level 1 version. The input signals are copied straight into the outputs. Thus, the signals are propagated through to Level 1 with minimal latency. The operation of the onboard TDCs remains exactly as designed for Level 1. The trigger logic block was removed, as this function will never be needed in the Level 0 modules.

In addition to serving as the trigger input TDC readout, the Level 0 modules can also facilitate crucial diagnostic tests on the trigger system when the experiment is not actively taking data. With some additional changes to the Level 0 firmware, we were able to set up a “pulser mode” in addition to the normal “run mode” described above. The purpose of the pulser test mode is to send pre-defined signal patterns through

the Level 1 and Level 2 modules. By comparing the inputs “pulsed” by Level 0 and the outputs of Level 2, we are able to rigorously check that the entire trigger system is performing exactly as expected.

The Level 0 pulser mode uses six RAM blocks to hold pulse patterns that are read in through the VME interface from user-created text files. Each RAM block drives half of one of the three 32-channel LVDS outputs. The pulser logic is driven by the beam-cycle RF-clock, which comes from the Main Injector to each v1495 module through a LEMO cable connected to the G0 port. The 53.1 MHz RF clock is processed by a PLL (Phase-Lock Loop) to generate a 212 MHz clock. Therefore, the pulser is able to split each RF-clock cycle into four independent sections. For a typical pulser test, a pulse signal will remain on for three of the four sections, simulating the pulse length normally sent by the discriminators in run mode. The pulser logic is able to choose whether to send one single pulse at a time (each time a “go” signal is received on the G1 input port), or to continuously pulse, stepping through the entire series of pulse patterns stored in the RAM blocks.

Control of the Level 0 pulser mode logic is achieved through a few new special functions added to the Level 2 logic and some additional NIM hardware to facilitate centralized control of the Level 0 and Level 1 board from Level 2. During pulser tests, starting and stopping the pulsers and TDCs for each board directly through the VXWORKS commands would make it nearly impossible to achieve the precise timing needed to match the timing precision of the trigger logic. Instead, VXWORKS commands are only sent to the Level 2 module. Level 2 then sends precisely timed commands to itself and the Level 0 and Level 1 modules, using the 212 MHz clock (generated by PLL from the 53.1 MHz beam-cycle clock). The new PLL and RAM blocks in the Level 2 firmware are controlled by bit registers which can be directly edited with VXWORKS console commands. Carefully timed NIM electronics are responsible for relaying the start/stop commands from Level 2 to Level 0 and Level 1, for the pulser and TDCs, respectively. With this setup, it is possible to use a single VXWORKS command to start the Level 1 and Level 2 TDCs, fire the Level 0 pulsers, and stop the Level 1 and Level 2 TDCs. The contents of the Level 1 and Level 2 TDCs are then compared against expectations given a certain firmware and a certain pulsed pattern.

3.3.2 Trigger System Testing

The ‘pulser mode’ of the Level 0 firmware allows for rigorous testing of the Level 1 and Level 2 logic. After the development of the Level 0 firmware, and before data-taking resumed, we utilized Level 0’s pulser capabilities to identify and characterize many bugs in the Level 1 and Level 2 firmware. We used two broad categories of pulser test. First, a simpler 3-board “bench-test” setup proved useful for checking internal timing and consistency in the Level 1 and 2 firmware. Second, a full *in-situ* pulser test, which uses the

entire trigger system as installed in the experiment, was also crucial for establishing internal consistency in the as-built trigger system. Lastly, the *in-situ* pulser test can also be used to quickly check for egregious internal timing errors in newly-compiled Level 1 firmware.

The 3-board, bench pulser test uses one Level 0 v1495, feeding pulser signals into one Level 1 board, feeding its output into one Level 2 board. Hit-patterns are pulsed from Level 0, Level 1 carries out its normal trigger logic functions, and Level 2 records the outputs of Level 1 using the onboard TDC. Synchronization of the pulser test timing is achieved through the new control functions added to the Level 2 firmware design. Specialized driver functions were written to carry out pulser tests and compare the Level 1 output (recorded by the Level 2 TDC) with the expectation derived from the known input pulser pattern. There are three main categories of bench pulser test: purity, clean efficiency, and messy efficiency. In a purity-style pulser test, the pulser patterns are arbitrary, except that they must not satisfy the road-conditions required by the Level 1 firmware. Sending a large number of (almost)-random pulse patterns through Level 1, we expect that Level 1 will never give any output. In a clean efficiency pulser test, the pulser RAM is populated with patterns that exactly match the hits required to fire a single Level 1 trigger road. Since the Level 1 output bins the roads in $\langle p_X \rangle$, each road pattern should fire one particular Level 1 output channel. Lastly, in a messy efficiency pulser test, the pulser RAM is filled with patterns that contain all hits to fire a particular road, *plus* some number of extra hits. The trigger system is designed to be a veto-free logic pipeline, which will fire on any input pattern containing hits that correspond to a pre-defined road. Therefore, the Level 1 output is expected to be independent of the addition of an arbitrary number of extra hits. The extra hits may indeed satisfy additional roads and could possibly cause additional output channels to fire. However, extra hits should never cause the expected output bit to fail to fire.

Through bench pulser testing, we found and fixed multiple issues with the trigger system. The *in-situ* pulser test uses the experimental configuration of the nine v1495 boards. Pulser signals from Level 0 are sent through the rest of the trigger and data acquisition systems. We used two different schemes for conducting *in-situ* pulser tests: an “automatic” pulser test scheme and a “manual” pulser test scheme. The automatic pulser test scheme functions identically to the bench pulser test except that all four Level 0 boards and all four Level 1 boards are involved. Since this pulser test configuration uses the contents of the Level 2 TDC to verify the outputs of the Level 1 boards, the four Level 0/Level 1 pairs are effectively independently tested. One drawback of this pulser test scheme is that the Level 2 firmware logic is not tested. The “manual” pulser test remedies this by checking the output of Level 2, after pulsing a particular set of patterns from the Level 0 boards. For the manual pulser test, the Level 1 and Level 2 TDCs are disabled (so that the Level 1 and Level 2 logic never needs to stop processing), and the “start pulse” signals are sent from a

manually-controlled gate-generator, rather than through the Level 2 board. By simply checking the count of Level 2 output signals on a visual scaler, we can check that the Level 1/Level 2 combined system fired every time a valid input pattern was pulsed from Level 0. Similarly to the bench pulser test, the manual pulser test can check the purity, clean efficiency, and messy efficiency of the entire Level 1/Level 2 system. The only difference is that the manual pulser test scheme can only check the total number of outputs against expectation, rather than individually record the output for each pulsed input.

The *in-situ* pulser test was used to verify the expected behavior of the full trigger setup. It is also useful as a quick diagnostic tool to check for internal critical timing errors in the Level 1 firmware that occasionally arise without causing compile-time errors.

3.3.3 Synchronization with the MI Beam RF-Clock

During the Main Injector shutdown between SeaQuest’s Run 2 and Run 3 data-taking periods, we made another significant improvement to the trigger system: clock synchronization.

In the Run 2 trigger, all logic processed by the v1495 modules was driven by an internally-generated 40 MHz clock. A PLL ramped this 40 MHz clock up to a 250 MHz clock, which then drove all of the fast-timing logic. The 250 MHz clock was further split into four phases, yielding an effective 1 GHz clock, which was used only for achieving 1 ns time resolution in the timing alignment and TDC blocks. All of the trigger decision pipeline logic used the 250 MHz clock, and therefore had a time resolution of 4 ns. Because each v1495 module used its own on-board 40 MHz clock, the boards’ clocks were not synchronized. In fact, the rates of the clocks were not even exactly identical, causing the phase-difference between each pair of boards to slowly vary.

We identified a number of issues that were directly caused by the non-synchronized clocks. Firstly, the output of the Level 2 trigger jittered around the RF-clock signal by ± 2 ns. Although the output is explicitly aligned to the RF-Clock in the last step within Level 2, the time resolution at this stage is 4 ns. Such a jitter in the Level 2 output, and therefore in the Common Stop signal, caused a number of difficulties, including alignment of NIM and FPGA triggers and applying in-time windows to recorded hits.

Furthermore, the asynchronous trigger system allowed for an intensity-dependent inefficiency. Each trigger module used an internal 250 MHz clock to do all logic processing, however, the experiment runs on the 53.1 MHz MI beam clock. Therefore, the trigger’s input and output must represent an 18.9 ns binning. The TDC and pipeline logic stored hits with a 16 ns binning, keeping only one hit per channel in each 16 ns block. If a single channel is hit twice within one of these 16 ns blocks, the earlier hit is discarded. The input hits are also only accepted in a 16 ns ‘in time’ window, defined by the timing of the RF-Clock. Because the

internal clocks were asynchronous with the RF-Clock, the RF-Clock-defined ‘in time’ window often spans two different internal 16 ns blocks. Therefore, in a high-rate hit environment, it is quite possible for a hit from one beam pulse to ‘delete’ a hit from a previous beam pulse in the same channel. If the later hit does not fall within the same RF-Clock-defined ‘in time’ window, the trigger logic might miss a hit from an otherwise good track.

Lastly, the firmware is limited to an ‘in time’ window width of $16\times$ the time resolution. For the Run 2 firmware, this means the ‘in time’ window could not be longer than 16 ns. Signal timing jitter caused by scintillator paddle length and pulse-height dependent PMT rise-times is expected to approach ± 8 ns for the largest paddles. If the hodoscope channels were not all *perfectly* timed-in, it is possible that a small number of hits might be lost at the edges. Therefore, we decided that we needed to somehow synchronize all of the v1495s for Run 3 and beyond.

We decided that the simplest way to address all of the above issues at once was to directly use the MI RF-Clock to drive all logic in the v1495s. Replacing the 40 MHz internal clock with the 53.1 MHz RF-Clock, and changing the PLL ratio from $\frac{25}{4}$ to $\frac{16}{4}$, the fast-clock became 212.4 MHz, and the TDC time resolution became 1.177 ns. All logic operations are completely synchronized with the experimentally-relevant MI RF-Clock. The output jitter is completely fixed. The internal 16 ns binning is now an 18.9 ns binning and is always synchronized with the experimental 18.9 ns pulse separation. The RF-defined ‘in time’ window now never spans two different internal bins. Hits can no longer be ‘deleted’ by hits from later beam pulses. The adjustable ‘in time’ window can now extend all the way up to the full 18.9 ns beam pulse separation, and is adjustable in 1.177 ns steps. We have chosen to use a 15 clock-tick wide ‘in time’ window (17.723 ns). This leaves a gap of one clock-tick (1.177 ns) between RF buckets.

There were two disadvantages to the change to RF-Clock-based v1495 logic clocking. First, the v1495 TDCs no longer have a nice 1 ns binning, and the input alignment minimum step has also increased to 1.177 ns. This change in binning is easily handled by small changes to the relevant data-handling software. Secondly, due to the slightly slower logic processing, the entire trigger system takes 100 ns longer to generate an output. Therefore, the timing of the NIM triggers and non-v1495 TDCs needed to be redone, but these systems were able to adjust with relative ease.

3.4 Peripheral Software

To ensure robust and repeatable operation of the trigger system, peripheral software was developed to automate many tasks which must be repeated whenever the experimental configuration changes. The two

most critical repeated “re-commissioning” tasks are input signal timing adjustment and trigger matrix roadset generation. The timing of the input signals to the trigger system was recalculated after any extended downtime and after any significant adjustment to the high-voltage settings of the hodoscope PMTs. A new trigger roadset was generated after large changes to the spectrometer detector stations. Small adjustments to the trigger roadsets (hot road re-evaluation) were made after each extended downtime and after any possible change in the hodoscope positions.

3.4.1 Timing Adjustment

The purpose of the trigger system is to examine hits on the hodoscopes from each beam pulse and decide whether or not an ‘event of interest’ occurred. Therefore, it is important that the timing of the inputs to the trigger system are carefully aligned. We must ensure that all hits from a single beam pulse are considered together and that no hits from adjacent beam pulses have leaked in. The 1 ns timing adjustment and TDC resolution, together with the 19 ns spacing of beam pulses, make satisfactory timing alignment achievable.

The internal delay adjustment feature of the firmware allows the trigger system to shift the timing of all input signals before applying the ‘in time’ window. Each input channel can be adjusted individually. In triggered data, a TDC spectrum for a single channel will show peaks at each beam pulse, ~ 19 ns apart. The peak corresponding to the triggered beam pulse will be much higher since it must contain a hit. Taking a large sample of triggered data, it is possible to identify the timing of the “in time peak” for each channel. The ‘in time peak’ timing for each channel informs a shift in the delay adjustment settings, aligning the timing of all input channels. The timing alignment procedure is iterative, beginning with data triggered by a roughly-timed NIM-based trigger, and improving until FPGA-triggered data shows all in-time peaks nicely aligned.

Although identifying the in-time peaks is usually simple to do by eye, automation requires careful consideration for some special cases. At first, we identified in-time peaks simply by finding the most-full bin in each TDC histogram. However, asymmetric peaks and random statistical fluctuations often cause the most-full bin to shift away from the center of the desired in-time window. Fitting a gaussian to the entire in-time peak does an excellent job of eliminating offsets due to statistical fluctuations. Although a gaussian shape cannot perfectly fit an asymmetric in-time peak, the centroid of the gaussian fit is an excellent estimator for the actual center of the desired in-time window position.

Since the delay adjustment feature is internal to the v1495 trigger system, other hodoscope readouts do not benefit from its careful alignment. In particular, all hodoscope channels are also recorded by the so-called “Taiwan TDCs”. These TDCs are identical to the ones used for the wire-chamber readout. Before

Run 2, we spent some time transferring the internal delay adjustments to physical cable delays upstream of the discriminators. We began by optimizing the internal delay settings as described above. Then, we simultaneously flattened the internal delays and added or subtracted the appropriate cable delays in each channel’s signal path. This procedure also required an iterative process of adjusting and remeasuring to achieve good timing alignment. Once the cable delays had brought all channels to within ± 2 ns alignment, we recalculated the internal delay settings, and corrected for the small remaining offsets withing the trigger system.

The hodoscope relative timings have changed significantly since the careful alignment with the cable delays. However, due to slow PMT performance drift or changes to PMT high-voltage settings, the internal timing delays are periodically recalculated to adjust for small shifts in the relative timing between hodoscope channels. Retiming is typically done any time there is a long down-time or a large change to the experimental apparatus. Since the timing differences are very small, retiming is quite easy, and usually converges in just one iteration.

3.4.2 Trigger Matrix Generation

The heart of SeaQuest’s trigger system is the Level 1 logic that defines when the trigger will fire, typically referred to as the “Trigger Matrix”. This logic consists of a look-up table and subsequent pipeline. The look-up table defines what combinations of individual hits on the four x-measuring hodoscope planes will fire the Level 1 “track finder”. The pipeline combines these look-up table matches, binning roughly in p_x , into the 24 bits that are sent to Level 2.

The look-up table, or “Trigger Matrix”, is generated through analysis of Geant4-based Monte Carlo (GMC), and ‘random background’ measured in the real apparatus. Two GMC productions are used to simulate the signal dimuons. The first is a production of target dimuons, used to determine the signal that the trigger is trying to identify. The second is a production of dimuons from the beam dump. Dimuons from the beam dump will often look quite similar to the target signal hit patterns. Since $\sim 90\%$ of the beam interacts in the dump, we must be careful to suppress the dump dimuons at the trigger level to reduce DAQ deadtime. To estimate the background, we use real data fired by the ‘Random RF’ trigger. The Random RF trigger is a NIM-based trigger designed to fire on a pseudo-random selection of beam buckets.

Trigger Matrix generation amounts to an optimization problem: achieve the highest possible rate of analyzable, high-mass, target-generated, Drell-Yan physics events by balancing signal acceptance against background-induced DAQ deadtime. The first step is to apply event-level cuts to both dimuon simulations. The event-level cuts restrict the kinematic coverage of the dimuons used to form the roads. This is necessary

because the GMC’s coverage extends beyond the experimental kinematic “region of interest”.

After the event-level cuts, the GMC events are grouped together into “roads”, and some average and cumulative information is calculated for each road. A separate step simulates applying this new set of trigger roads to the randomRF run, and the total number of randomRF events that would have satisfied each trigger road is stored along with the other ‘road-level’ information. We then apply “road-level” cuts, to remove roads which are expected to have a low signal-to-background ratio. The typical three road-level cuts are a cut on the number of target GMC events, a cut on the number of randomRF events, and a cut on the minimum GMC p_x .

The final step in the trigger road selection procedure is hot road removal. After generating and testing v1495 firmware for the new roadset, we can take a few test runs in real beam conditions. Online analysis of these test runs generates a list of the “hottest” roads from each quadrant. We typically remove any outlier hot roads, after checking that each one is not expected to make a significant dent in the high-mass, target Drell-Yan acceptance. After each group of outlier hot roads is removed, a new firmware must be compiled, tested, and rechecked for new outlier hot roads. This procedure typically converges after just a couple iterations, with usually fewer than ten roads removed in total.

3.5 Performance of the Trigger System

The SeaQuest Trigger System has performed admirably through the 4 years of SeaQuest’s data taking. Coupled with the Beam Monitor Inhibit system, the trigger is able to preferentially select candidate dimuons from the very high-rate environment. Analysis of the collected data shows that the experiment is capable of collecting signal dimuon events from Drell-Yan, J/Ψ , and Ψ' decays with a modest combinatorial background. Mass fits, using mass distribution shapes from Monte Carlo for signals, and mixed-singles from data for background, are able to successfully describe the mass distribution of the collected data, as seen in Figure 3.7. Continuous improvements to the trigger system have resulted in increased signal acceptance, increased background rejection, improved self-monitoring, and improved fine-timing alignment.

3.5.1 Signal Acceptance and Background Rejection

The primary purpose of the trigger is to select events which are likely to contain a high-mass dimuon. Therefore, it is important that the set of trigger roads provides adequate acceptance, at the trigger level, in the kinematic regions which are most important to the intended physics analyses. Here I present the kinematic acceptance distributions for the sets of trigger roads used in SeaQuest data-taking periods.

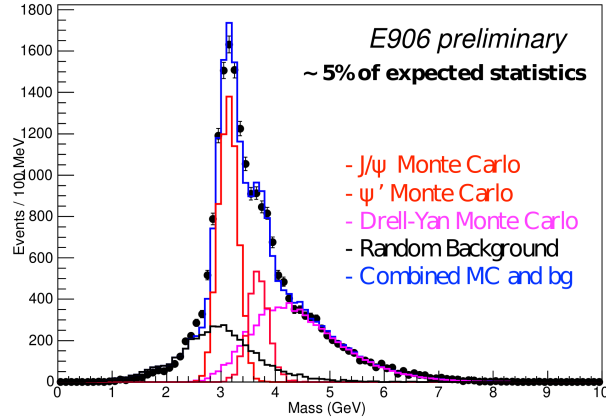


Figure 3.7: Signal+Background Mass Spectrum Fit to the small subset of data first released at international meetings.

The mass acceptance is heavily biased toward the high-mass region. The reason for this is three-fold. First, the flagship analysis of SeaQuest is the measurement of $\frac{D_2}{H_2}$ at $0.3 < x_2 < 0.5$. Due to the geometric acceptance of the spectrometer, high- x_2 events also tend to have high-mass. Even though these events are much rarer than the lower-mass events, they are much more valuable to the flagship physics. Second, the charmonium (J/Ψ and Ψ') cross-sections are much larger than that for Drell-Yan. Therefore, the trigger acceptance must heavily suppress the charmonium mass region to prevent charmonium decay events from dominating the DAQ live time. Third, the trigger rate is dominated by combinatoric background events consisting of two single muons that together conspire to satisfy a dimuon hodoscope hit-pattern. These single muons mostly originate from pion decays, and, relative to the Drell-Yan generated muons, tend to have lower single-track p_T . Therefore, fake dimuons built from combinatoric pairings of these low- p_T singles typically have quite low mass. So the trigger must strongly suppress the low-mass region in order to prevent the combinatoric background from completely dominating the DAQ live time.

The Run 3 trigger roadsets picked up some additional acceptance at the high-mass edge relative to the previous three roadsets. Due to a change of the orientation of the two magnetic fields at the end of the roadset 62 data taking, it was decided that the trigger roadsets should be symmetric with respect to positively and negatively charged tracks. Enforcing the charge symmetry caused some rare, very high k_T roads to be added to the positive or negative set of trigger roads. This then allows some additional, exceedingly rare, very high mass dimuons to enter the acceptance. This small acceptance boost is expected to have no effect on the collected data since the expected rate of dimuons in that very high mass region is vanishingly small.

Due to improvements in the beam quality, we were able to keep a bit more acceptance in the Run 3 trigger roadsets, relative to Run 2. The trigger rate, and therefore the DAQ livetime, is very sensitive to

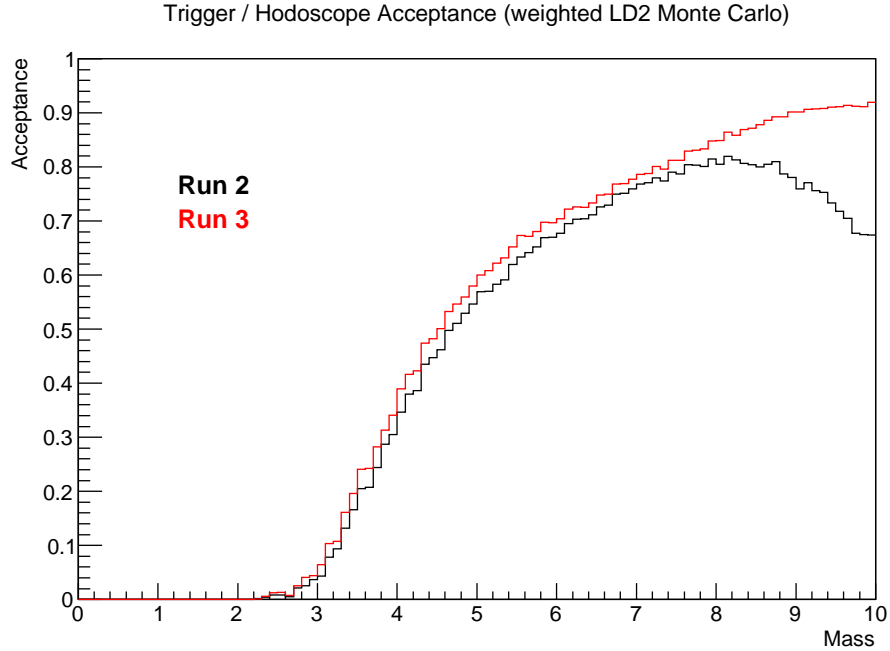


Figure 3.8: Comparison of Mass Acceptance for Run 2/Run 3 Trigger

beam quality.

Since the flagship physics analysis will be presented as a function of x_2 , it is helpful to avoid a trigger acceptance which is a strong function of x_2 .

An important consideration when determining a set of trigger roads is the expected rate of combinatoric background fired triggers. Estimating the combinatoric background rate is not as straightforward as determining the signal acceptance of the trigger, as the combinatoric background is much more difficult to reproduce with pure simulation. The difficulty in reproducing the background comes largely from the highly irregular beam structure observed by the real data.

Luckily, one of SeaQuest’s calibrations triggers (NIM3) was designed specifically to fire on random beam pulses, with no relation at all to beam-monitor readouts or detector signals. Using dedicated runs fired by only NIM3, we collect enough random-data statistics to make a reasonable estimate of the fraction of those random events that would have fired a particular trigger roadset. Since the readout of the v1495 L1 TDCs perfectly records the inputs to the trigger system, background rate analysis of a trigger roadset does not even require firmware generation and taking a special run. We simply simulate the trigger decision for a particular combination of L1 roadsets and L2 logic on an appropriate NIM3 run, and compare the results with other considered roadsets.




Figure 3.9: Comparison of x_2 Acceptance for Run 2/Run 3 Trigger

3.5.2 Input/Output Consistency

We have coined the terms “Logic Efficiency” and “Logic Purity” to refer to the efficiency and purity of the trigger systems output, relative to the recorded trigger system inputs. To make clear, this is separate from the efficiency and purity of a particular roadset, in it’s ability to pick out good candidate Drell-Yan events from the background. The design of the SeaQuest trigger system is intended to have perfect logic efficiency and logic purity. If the v1495 TDCs record a road-matching hit pattern, then the appropriate output trigger should always fire. Inversely, if the TDCs record a non-matching hit pattern, the trigger should never fire. Unfortunately, this expectation was not strictly true until the “RF-Clocking” upgrade between Run 2 and

Run Period	Roadset	Signal Acc.	Bkg. Acc. Run 3 (Run 4)	Signal Amp. Run 3 (Run 4)
Run 2	57	59.76%	0.1833% (0.2461%)	326.0 (242.8)
	59	59.76%	0.1833% (0.2461%)	326.0 (242.8)
	62	59.76%	0.1833% (0.2461%)	326.0 (242.8)
Run 3	67	62.95%	0.1852% (0.2694%)	339.9 (233.7)
	70	62.45%	0.1775% (0.2642%)	351.8 (236.4)
Run 4	77	90.77%	0.6021% (0.7486%)	150.8 (121.3)

Table 3.2: Trigger roadset signal acceptance and “Random RF” background acceptance estimation. Signal acceptance is calculated from DY GMC and represents the Matrix 1 acceptance vs a hypothetical “all roads” trigger. The DY analysis mass cut, Mass > 4.2 GeV, is applied to the GMC. Background acceptance is also Matrix 1 relative to ‘all roads’, but using Random RF data from Run 3 (Run 4). The “Signal Amplification” is the Signal Acceptance divided by the Background Acceptance.

Run 3.

A new method of flagging in-time vs not-in-time hits from the v1495 TDCs was required to ensure accurate measurement of the trigger logic performance. The so-called “RF-based in-time flagging” method exploits the fact that the v1495 hit acceptance windows are set relative to the input RF-Clock received from the Fermilab Main Injector. Using a small, fixed window applied to the RF-clock readout channels, the exact timing of the “in-time” RF-Clock signal is identified. Then, when decoding the data from the raw CODA file into the database, the in-time flags for all v1495 hits are set according to their relationship to the timing of the in-time RF-Clock hit on the appropriate v1495 board. The width of this in-time window is user-selectable on a run-by-run basis. For Run 2, the window width was set to 15 clock cycles (15.0 ns). For Run 3, the window width was set to 15 clock cycles (17.7 ns).

With the internal timing of the trigger logic synched to the MI RF-Clock, and the decoded data using the RF-based in-time flagging method, the DAQ records a fired dimuon trigger every time a matching hit pattern is seen in the hodoscopes, and the hit pattern in the hodoscopes matches a pair of trigger roads for every event in which the trigger system fired the dimuon trigger.

3.6 Issues in the Trigger System

Although the v1495 modules and firmware have performed as designed throughout SeaQuest’s data taking periods, some peripheral issues related to the trigger system have adversely affected recorded data in some small chunks of data.

Chief among these issues are unpredictable timing shifts due to bizarre failure modes of LeCroy Fan-In/Fan-Out modules. These Fan-In/Fan-Out modules (LeCroy 429A) are used in many places in the SeaQuest DAQ chain to distribute various NIM-logic signals. These modules are used to duplicate the “Common Stop” signal, as it is sent from the Trigger Supervisor out to all the TDCs. A few modules are also used in the propagation and distribution of the RF-Clock signal, received from the Main Injector, and sent into the v1495 and NIM trigger systems. SeaQuest acquires these modules in good working order from FNAL PREP. However, they seem particularly prone to creative and often surreptitious failures.

The most common failure mode of the LeCroy 429A modules, as observed by SeaQuest, is a sudden shift in the output timing of some or all output channels, relative to the input timing. We have observed these timing shifts of various Δt values: 2 ns, 4 ns, and even 16 ns. Even relatively small timing shifts, of a few nanoseconds, can be detrimental to recorded data if they are not identified and fixed quickly. While diagnostic timing information is recorded with every event, the possibility of such timing shifts necessitates

constant vigilance in reviewing the recorded timing information. Ideally, an automatic monitoring alarm system would monitor the timing of recorded channels, and signal an un-ignorable alarm immediately upon observation of a timing shift. However, such a sophisticated monitoring system is not trivial to implement and has not been completed for SeaQuest yet.

Another, less common, failure mode of the LeCroy Fan-In/Fan-Out modules is a failure of the “ $4 \times 4/8 \times 2/16 \times 1$ ” selection feature. On one occasion, one Fan-In/Fan-Out, which was set to “ 8×2 ” mode, suddenly began fanning *all* 16 input channels to *all* 16 output channels, effectively operating in “ 16×1 ” mode. This failure mode is particularly nefarious and was quite a challenge to track down and diagnose.

3.7 Future Upgrades

Although the Trigger System will remain mostly unchanged through the end of the first incarnation of SeaQuest, future extensions to the SeaQuest science program will likely re-use and upgrade the existing trigger.

The simplest extension is to include the information from the y-measuring hodoscopes in the trigger decision. The hardware is already in place, and the firmware is already designed and ready to use. All that remains is to decide on an effective selection logic and modify the Level 2 logic to include the non-bend-plane information in the final trigger decision. One possibility is simply defining “roads” for the non-bend direction, in a similar fashion as the standard, bend-plane roads. The addition of non-bend roads should somewhat cut down on the number of triggers fired by combinatorial additions of hits from more than two tracks. The non-bend roads are not expected to reduce the rate of triggers fired by the coincidence of two real, unassociated tracks. However, requiring a non-bend di-track to be coincident with a bend-plane di-track significantly increases the trigger system’s sensitivity to any hodoscope inefficiency. The bend-plane di-track alone required the coincidence of 8 independent hodoscope channels. Implementing the non-bend tracks into the trigger decision increases this coincidence requirement to 16 independent hodoscope channels.

It would also be relatively straightforward to transform part of the p_X -binning machinery into a scheme for removing high-background “di-roads”, without completely removing both of the individual roads. This scheme would work as follows. Each road in the Level 1 firmware can be tagged as either “normal” or “hot”. Roads which are completely dominated by background would still be removed from the roadset. However, roads with a high background rate, but which also contribute significantly to the signal rate would be flagged as “hot” rather than completely removed. Then, the Level 2 logic checks all pairs of roads, and fires the trigger on the combinations “normal-normal”, “normal-hot”, and “hot-normal”. However, it does not fire on “hot-

hot” road combinations. Therefore, the high trigger rate caused by hot roads may be significantly reduced, without completely removing the signal events which traverse the hot roads. A significant disadvantage to this proposal is handling the trigger acceptance effects in data analysis. Normalizing the acceptance of the “single-hot”-fired events against the “double-normal”-fired events would not be trivial. In light of the significant rate dependence observed in the SeaQuest data, this signal-boosting scheme might turn out to cause much more analysis difficulty than the extra signal events are worth.

A similar intriguing possibility is to attempt to split different categories of events into independent roadsets, which could all run simultaneously in the v1495 trigger system. Again, repurposing the p_X -binning machinery makes this relatively straightforward. Each trigger “subset” is assigned a Level 1 output bit. The Level 2 trigger then only accepts events in which at least one dimuon was triggered in a single subset. With this scheme, roads could even be repeated within the various subsets. Additionally, since each subset could be assigned to a different Level 2 output, each subset could be prescaled independently. Such a system could provide a fantastic degree of control over the relative live-time fractions of different subsets. This could be especially useful for the “parasitic” analyses, for which some trigger fine tuning might make the difference between a possible and an impossible measurement. In the current trigger logic design, any roadset changes that attempt to alter the acceptance for a fringe analysis are very likely to also directly affect the flagship measurement acceptance.

A relatively more difficult possible upgrade is the implementation of machine learning techniques in designing the pipeline logic. SeaQuest collaborators at University of Michigan have begun detailed proof-of-concept studies for using Neural-Networks or Boosted Decision Trees to design the Level 1 and Level 2 trigger decision logic.

For the purpose of significantly improving the trigger’s sensitivity to dark photon events, a significant hardware upgrade has been proposed. New, fine-grained hodoscopes would be installed just upstream of KMAG, and would feed into a completely overhauled v1495 trigger system. The additional hodoscope channels added by such an upgrade cannot be accommodated in the current trigger setup. The hardware limits each v1495 module to a maximum of $32 \cdot 4 = 128$ signal input channels. The firmware limits each v1495 to 96 total internal channels. Two of these are used for the RF-Clock and the STOP signal. The existing bend-plane hodoscopes occupy $23 + 16 + 16 + (2 \cdot 16) = 87$ channels on each bend-plane v1495.

Chapter 4

Data Analysis

For this analysis, the p+d and p+p data from two run periods will be used. The two run periods, called “Roadset 62” and “Roadset 67”, are the latest, largest, and highest quality sets of SeaQuest data currently available. Table 4.1 lists some information about these run periods. Table 4.2 lists the total and high-mass numbers of valid target dimuons for each roadset and target.

4.1 Track Reconstruction

The primary reconstruction program at SeaQuest is called ‘kTracker’. This program was developed in-house by Kun Liu and has been optimized for the collected SeaQuest data.

4.1.1 Pre-Tracking Cuts

Due to the high rate of background present in SeaQuest’s triggered data, a number of pre-tracking cuts are applied before any track finding is attempted. These cuts serve both to remove events which are very likely to be untrackable, to remove extraneous hits from trackable events, and to speed up the runtime of the tracking algorithm. There are four hit removal cuts and many multiplicity-based event removal cuts.

The simplest hit removal cut is the in-time cut. Based on the timing of the trigger signal, all chamber hits which could not have been caused by a particle from the triggering beam pulse are removed. Since the drift time in the chambers is relatively slow as compared with the beam pulse frequency, the in-time window of the chambers spans multiple beam pulses. The chambers with the largest cell width have the widest in-time windows, up to ± 15 beam pulses (~ 600 ns) wide.

The next hit removal cut is “afterpulse” removal. Due to the tendency of the wire chamber channels to ‘ring’ after being hit, and thereby causing multiple hits on the same channel after the true hit, a cut is applied to select only the first hit. If there are multiple hits on a single channel after the in-time cut is applied, then only the first in-time hit is kept. All later hits (in that one event) on that channel are removed, and not considered by the tracking.

Table 4.1: Analyzed Datasets

Roadset	Begin Date	End Date	# p+d Spills	# p+p Spills
62	2014-11-25	2015-01-14	9350	19751
67	2015-01-25	2015-06-19	26114	38218

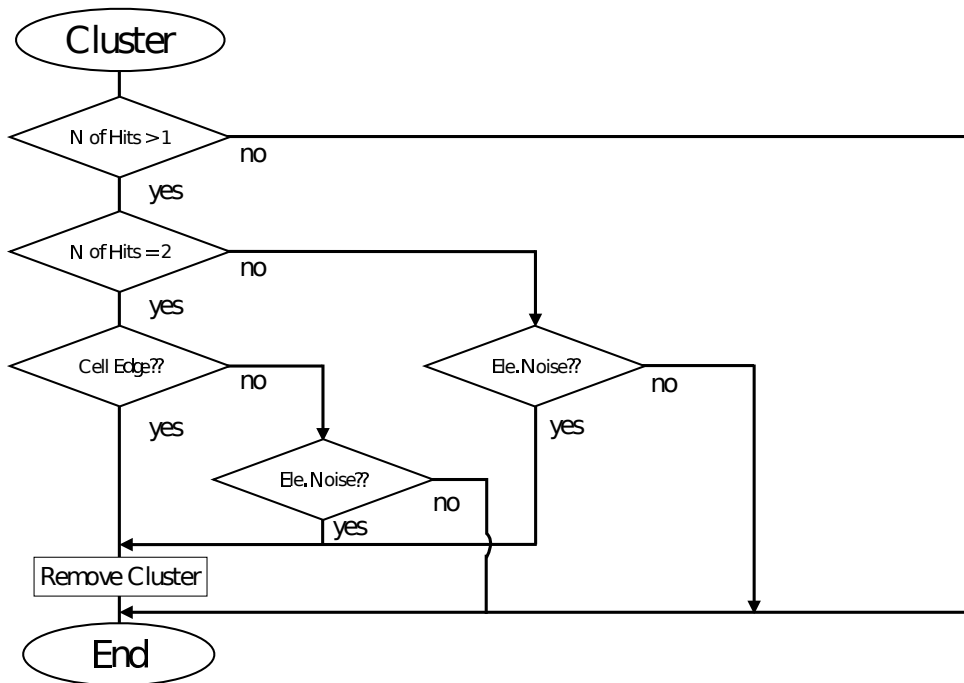


Figure 4.1: Flow chart of Hit Cluster removal algorithm

The most complex hit removal cut is the cluster-removal. There are three identified types of hit removal cuts: electronic noise, cell-edge hits, and delta-rays. Any grouping of two or more adjacent hits in a wire plane is considered a cluster.

The chambers sometimes experience spurious signals in the readouts due to electronic noise. These noise hits are not associated with any real signal and should be completely removed. If a cluster of hits all have similar timing (within 10 ns), they are identified as an electronic noise cluster and are removed. For Stations 1, 2, and 4, this cut is only applied to clusters with size > 3 . In the D3p chamber, which suffers from

Table 4.2: Dimuon Counts

Roadset	Target	Total	$M > 5 \text{ GeV}$
62	D_2	17335	1815
62	H_2	17997	1893
67	D_2	64891	6414
67	H_2	61577	5843

higher-than-usual electronic noise, this cut is also applied to 2-hit clusters with a window width of 8 ns.

Cell-edge hits are defined as pairs of adjacent hits in a single wire plane where the drift-distance is greater than 90% of half the cell width. In these cases, it is most likely that a single muon passed through the boundary between the two cells and fired both channels. To speed up the tracking, only the hit with the smaller drift distance is kept.

'Delta-rays' are low-energy, knock-on electrons that can scatter away from the muon's path at high angles. Some of these electrons travel parallel to the wire plane, and can fire many wires in a row. The drift speed in the chamber gas is low enough that delta-ray induced clusters will not be removed by the electronic noise cuts. To reduce the load on the tracking algorithm, the delta-ray hits should be removed. However, the 'real' hit left by the muon which induced the delta-ray needs to be kept. Since the real hit will usually be located at the end of a delta-ray induced hit cluster, the two edges of these clusters are kept, but all the hits between them are dropped.

The final hit removal cut to be applied is the Trigger Hodoscope Masking. In normal hodoscope masking, chamber hits are dropped if they are not 'behind' a fired hodoscope paddle. Since the hodoscopes are very fast (single-bucket timing resolution), this cut can dramatically reduce the number of chamber hits that the tracker must consider. Assuming very high hodoscope efficiency, this cut should remove almost no good chamber hits. It is possible to improve the hodoscope masking by considering the road requirements of the trigger. First, all in-time hodoscope hits in an event are combined to make all possible road combinations, like (H1,H2,H3,H4). Then these possible roads are checked against the list of actual active trigger roads loaded into the v1495 trigger system. Road combinations not active in the trigger are removed. Finally, the list of all hodoscope paddles in the remaining road combinations is used to do hodoscope masking, as described above. Therefore, only chamber hits which are masked by hodoscope hits which could have actually fired the trigger are used in track reconstruction.

The 'event removal' pre-tracking cuts serve to remove only those events which would have been untrackable anyway. Events with extremely high multiplicity can cause the tracking algorithm to spend a very long time attempting to identify tracks from the huge number of hit combinations. Often for these extreme events, the tracking algorithm hangs indefinitely. The definitions of these cuts, usually called "Multiplicity Cuts", are given in Table 4.3.

There is one additional multiplicity cut on the number of matched plus or minus trigger roads. If the fired in-time hodoscope pattern in an event matches five or more positive trigger roads or five or more negative trigger roads, then tracking is not attempted.

Table 4.3: Multiplicity Cuts

Detector(s)	Hit Limit
D1	250
D2	200
D3p	150
D3m	120
Prop Tubes	250
H1T+H1B	10
H2T+H2B	10
H3T+H3B	10
H4T+H4B	10

4.1.2 Track Finding

The first part of the true track reconstruction program is the track finding. In the track finding stage, combinations of chamber hits are whittled down to produce a list of plausible track candidates. There are four basic steps: local triplet reconstruction, station 2-3 track construction, projection to station 1, and global track construction.

The first step of the track finding routine is local triplet reconstruction within the St. 2 and St. 3 chambers. The goal of this step is to construct track ‘stubs’, formed from hits in all three views, X, V, and U. Figure 4.2 depicts the triplet reconstruction. First, hits in the primed/unprimed pairs are combined in order to reduce the combinatoric load. For each X-view hit, all overlapping U-view hits are collected. The X-U combination defines a very narrow range of possible V-view hits. All collected X-U-V combinations are fit to generate triplets with well-defined position and pointing direction. At this stage, the use of single hits in multiple triplets is permitted.

With the collection of triplets in St. 2 and St. 3, the tracker then attempts to combine them into straight tracks between the two stations. In addition to checking that the triplets point at each other, combinations are removed if they do not point generally toward the target, if they do not point at fired hodoscopes in H2, H3, and H4, or if they do not point at fired ‘tracklets’ in the prop tubes. The check for corresponding prop tube tracklets ensures good muon identification, due to the iron wall upstream of St. 4.

With a handful of good St. 2-St. 3 tracks candidates in hand, the track finding algorithm projects each one back to St. 1, and checks for corresponding St. 1 triplets. Depicted in Figure 4.3, the St. 1 projection step uses the ‘sagitta ratio’ method to limit the search area on St. 1. The magnetic fields are treated as single-impulse p_T -kicks. The effective bend angles in the magnets vary only with p_z , and constant factors such as the field strength. Therefore, the ratio of s_1 and s_2 is approximately constant for reasonably high-momentum tracks. Figure 4.4 demonstrates that the sagitta ratio is nicely gaussian-distributed around, and




Figure 4.2: Triplet reconstruction in a drift chamber

localized enough to be of use for projection of tracks forward to St. 1.

After projecting the track to St. 1, a ± 5 cm window is applied around the expected hit position. All hits within this window are used to form St. 1 triplets. The St. 1 triplets are matched up with the projected track, forming a candidate global track.

The last step of the track finding is global track construction. The purpose of the global track construction is to clean up the candidate tracks and to prepare them for fitting. First, bad hits are iteratively removed; hits with residuals greater than three times the resolution are removed one at a time, refitting the track between each. When all hits are within 3σ , the global track is subjected to quality cuts. Global tracks are kept if: the appropriate hodoscope paddle fired in each station, the momentum is within 5 – 100 GeV, the track has at least four hits in each station, the track has at least one hit in each “view”, and the multiple scattering angle in the iron wall is not too large. Tracks passing this stage are passed along to the Kalman Filter based Track Fitting algorithm.

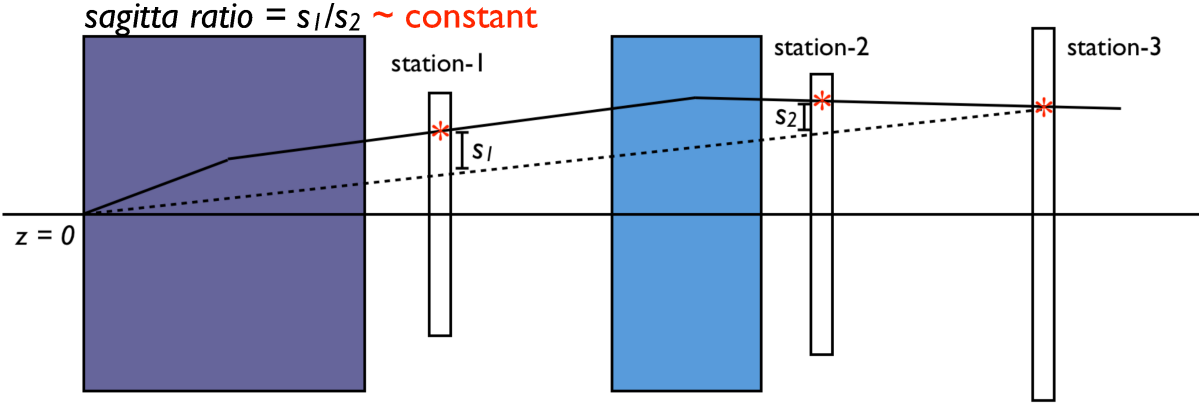


Figure 4.3: Sagitta-Ratio Projection to Station 1




Figure 4.4: Monte Carlo Sagitta-Ratio Distribution

4.1.3 Track Fitting

The kTracker reconstruction program uses a Kalman Filter based algorithm to fit the collection of hits associated with a found track, extracting the track parameters.

The Kalman Filter [48] is a widely used recursive algorithm which can efficiently estimate the state of a dynamic system by sequentially applying knowledge gained from imperfect measurements. To apply a Kalman Filter, the details of the physical model and the appropriate state vector must first be defined. The model is used to propagate the state of the system from each measurement to the next, generating an estimate of the observable parameters prior to the application of each measurement. At each measurement, the measured parameter values (and associated uncertainties) are combined with the estimate of the state vector (and its uncertainties) in a weighted average. The weighting is used to scale the measurements by their expected accuracy, and is typically representative of detector resolution. After the weighted average is performed, the physical model propagates the new state vector to the next measurement. Thusly, the Kalman Filter propagates the state of a dynamic system through a series of noisy measurements, generating

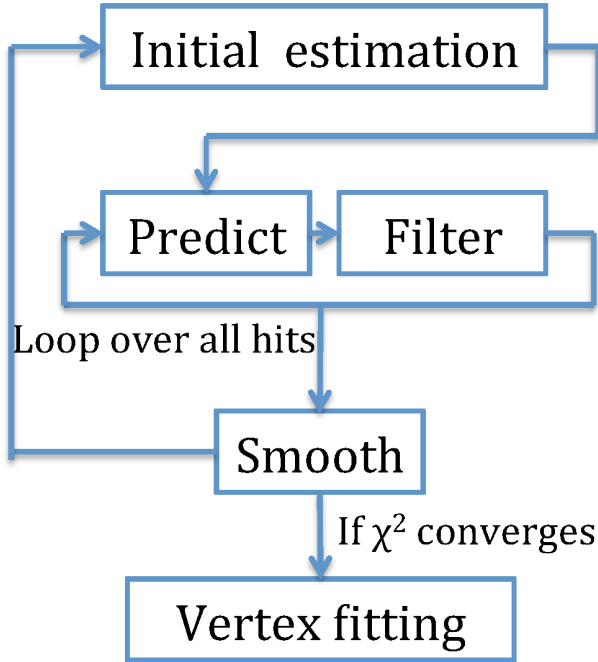


Figure 4.5: kTracker’s track fitting algorithm flow

a final state vector with properly defined covariance.

For implementation in the track fitting step of kTracker, the general principles of the Kalman Filter are somewhat modified. Firstly, instead of using an analytical physical model, kTracker implements Geant4 to perform the dynamic state propagation. Geant4 is widely used and well tested, and can simulate a comprehensive variety of interactions between relativistic particles and various materials. In particular for SeaQuest, Geant4 provides detailed simulation of muon propagation through the materials and magnetic fields present in the spectrometer.

Secondly, kTracker applies the Kalman Filter *backwards* in time, propagating the muons from the downstream end of experiment back to the target/beam dump vertex position. The filter is applied backwards for two main reasons. First, the downstream end of the experiment is the cleanest, in terms of hit multiplicities. Therefore, kTracker need only attempt to track relatively few tracks per events, as opposed to starting at the very noisy upstream end of the spectrometer. Secondly, due to the relatively large number of detectors at station 4 (three hodoscope planes and four prop tube plains), and due to the lack of any magnetic fields, the ‘tracklets’ are quite well-determined. Starting with a good estimate of the state vector helps reduce the required number of iterations for track convergence.

In the kTracker implementation, the state vector contain the three-momentum, (p_x, p_y, p_z) , and the position, (x, y, z) . As the muon is propagated backward through the spectrometer, the measurements provide up-




Figure 4.6: Kalman Filter reverse propagation in kTracker track fitting




Figure 4.7: Single slice of reverse FMAG ‘swim’ in kTracker

dated knowledge about the observable parameters, (x, y, z) . Uncertainties on the measurements are derived from known wire chamber resolutions. Geant4 provides all propagation, including the final measurement-free propagation between Station 1 and the target area. After the backward-in-time propagation, the muon is propagated back through the spectrometer, forward in time. The wire chamber measurements are re-applied using the Kalman Filter algorithm. This second pass through the measurements helps to smooth the variations in the state vector along the track path. If the χ^2 of the fitted track converges, the track is sent to the vertex fitting stage of kTracker.

4.1.4 Vertex Fitting

After tracks have been found, and individually fit to convergence, they must be combined into pairs and the best estimate for the dimuon vertex must be extracted. The kTracker algorithm again uses a Kalman Filter based approach for this step. The Kalman Filter vertex fitting approach is based on a similar implementation [49] at ALICE by Gorbunov and Kisel.

Before fitting the dimuon vertex, the found tracks must be projected back through FMAG to the target region. Each track ‘swims’ backwards through FMAG to find the point of closest approach to the z-axis.

For the swim, FMAG is cut into many z-slices, and an energy loss (energy gain) and p_T -kick are applied in each slice. After the swim, each track has a projected straight segment passing through the target region. The parameters of these straight segments are the inputs to the vertex fitting step.

For the vertex fitting step, the state vector is the dimuon vertex position (x_d, y_d, z_d) , and each track is treated as a measurement. The initial guess for the dimuon vertex position is $(0, 0, z_{\text{mean}})$, where z_{mean} is the average of each track’s z-of-closest-approach. Since there are only two ‘measurements’ in this scheme, the entire Kalman Filter step is iterated. The output vertex position is fed back into the Kalman Filter, with the covariance matrix re-initialized, and the ‘measurements’ of the tracks are applied again. When the fitted vertex position converges, the dimuon reconstruction is completed and the results are stored for analysis.

4.2 Rate Dependence

The reconstructed SeaQuest dimuon data unfortunately suffers from a severe rate-dependence. The rate-dependence ultimately stems from the non-uniformity of the slow-spill extracted Main Injector beam. Large variations in incident proton intensity cause large variations in hit multiplicity. The efficiency of the track reconstruction is sensitive to the hit multiplicity, and therefore the number of reconstructed dimuons, scaled by same-pulse incident protons, is a strong function of the beam intensity.

Correcting the rate-dependence is non-trivial. The rate-dependence is correlated with some of the important kinematic variables; the spatial distribution of hits across the spectrometer planes is non-uniform for both the intensity-induced multiplicities, and various kinematic variables. Therefore, all analyses must satisfactorily correct the rate-dependence before extracting physics results.

4.2.1 Rate-Dependent Reconstruction Efficiency

The strategy for the rate-dependence correction rests on the simple idea of simulating and extracting the rate-dependent reconstruction efficiency using Monte Carlo. Simulating the nonuniform, intensity dependent background hits is very difficult, so we have used real beam-generated backgrounds using our “Random RF” triggered event sample. Since the Drell-Yan cross-section is relatively small, the likelihood of finding a real DY dimuon in a randomly-chosen beam pulse is vanishingly small.

An apt question concerning this approach is: “Do the intensities and multiplicities of the Random RF events accurately reflect the backgrounds in our recorded ‘dimuon’ data?”. The answer is yes under the assumption that the Drell-Yan generation probability is independent of the background-generating hard-

scattering probability. The multiplicity of background hits found in each bucket is simply a statistical probability function of one variable: the intensity of the bucket. Likewise, the probability of a Drell-Yan dimuon pair is a linear function of the intensity. Since the Drell-Yan cross-section is small, the additional background generated by the beam proton remnant is negligible. Therefore, it is safe to assume that the background multiplicity of each event is a function of only the intensity, and is uncorrelated with the production of Drell-Yan dimuons. It is then clear that the ‘‘Random RF’’ events must accurately reflect the background hits ‘‘under’’ the true Drell-Yan dimuon signal.

However, our reconstructed data sample is not composed entirely of true Drell-Yan dimuons. A significant fraction of the reconstructed dimuons are ‘‘Combinatorial Background’’, i.e. random pairings of uncorrelated single tracks that spuriously successfully reconstruct to a ‘dimuon’.¹ These fake dimuons should rise with I^2 (or higher) and are probably highly correlated with the background multiplicity. Therefore, the extracted rate-dependence reconstruction efficiency cannot be safely applied to the combinatorial background events.

To isolate the reconstruction efficiency caused only by the intensity-induced background multiplicity, we define the kEfficiency as:

$$e_k = \frac{e_{\text{messy}}}{e_{\text{clean}}} \quad (4.1)$$

where e_{messy} and e_{clean} are the binomial efficiencies of successfully reconstructing the thrown Monte Carlo dimuon within the messy and clean samples, respectively. The clean sample is simply the clean, thrown Monte Carlo dimuons. The messy sample is the Monte Carlo dimuons, merged at the hit-level with Random RF events, one Random RF event for each Monte Carlo dimuon. Thus, the clean sample is the ‘no-background’ reference and is used to cancel out any reconstruction efficiencies not caused by the background hits.

Using the reconstructed messy and clean samples, the accuracy of the dimuon reconstruction at different levels of background can be studied. For this purpose, the events are split into three categories: overlap, clean-only, and messy-only. The overlap sample is all events in which a good dimuon was found in both the messy and clean sample. The clean-only and messy-only samples are self-explanatory. Figure 4.8 shows the mass residuals for these three samples; the overlap residual is calculated for the messy and clean reconstructions separately. Clearly, the addition of background hits has no significant effect on the accuracy of the reconstruction. Therefore, it is approximately accurate to consider the good dimuons found in the messy sample to be ‘the same’ as the thrown dimuons. Furthermore, the residuals demonstrate that the ‘efficiency enhancement’ effect is finding ‘the same’ thrown dimuons.

The reconstruction efficiency curves are calculated separately for each roadset and target. Figure 4.9

¹(There is also some contribution from ‘higher-order’, background, where hits from more than two uncorrelated tracks conspire to reconstruct a dimuon. For example, one single track and a St. 1 segment of a second track, together with a St2-3-4 segment of a third track.)

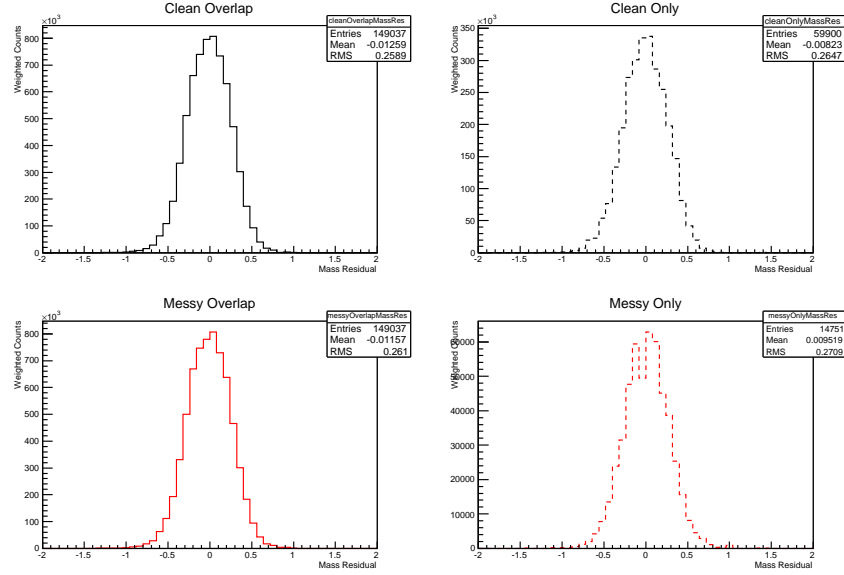


Figure 4.8: Mass residuals for Messy, Clean, and Overlap components of the GMC+NIM3 embedded data




Figure 4.9: kTracker Rate-dependent Efficiency for Roadset 67, Deuterium target

shows the reconstruction efficiency as a function of intensity for the deuterium target for roaset 67. The fit is an exponential, with the y-intercept fixed at 1. Therefore, each fit is characterized by just one parameter, a .

$$e_k(I_C) = e^{-aI_C} \quad (4.2)$$

The shape of the rate-dependence is well described by the constrained exponential fit. The efficiencies are used to calculate a rate-dependent weight, w .

$$w = \frac{1}{e_k(I_C)} = \frac{1}{e^{-aI_C}} = e^{aI_C} \quad (4.3)$$

Table 4.4: Rate-Dependent Reconstruction Efficiency Parameters, “a”, θ -bins

θ bin	67 Deuterium	67 Hydrogen	62 Deuterium	62 Hydrogen
1.071–1.171	$2.250(150) \times 10^{-5}$	$2.250(144) \times 10^{-5}$	$2.360(266) \times 10^{-5}$	$2.130(188) \times 10^{-5}$
1.171–1.271	$2.040(135) \times 10^{-5}$	$1.680(113) \times 10^{-5}$	$2.010(241) \times 10^{-5}$	$1.860(171) \times 10^{-5}$
1.271–1.371	$1.910(126) \times 10^{-5}$	$1.790(118) \times 10^{-5}$	$1.790(204) \times 10^{-5}$	$1.680(158) \times 10^{-5}$
1.371–1.471	$1.770(119) \times 10^{-5}$	$1.640(110) \times 10^{-5}$	$1.820(211) \times 10^{-5}$	$1.440(137) \times 10^{-5}$
1.471–1.571	$1.710(117) \times 10^{-5}$	$1.530(106) \times 10^{-5}$	$1.580(193) \times 10^{-5}$	$1.440(142) \times 10^{-5}$
1.571–1.671	$1.610(109) \times 10^{-5}$	$1.520(102) \times 10^{-5}$	$1.540(192) \times 10^{-5}$	$1.380(138) \times 10^{-5}$
1.671–1.771	$1.670(114) \times 10^{-5}$	$1.590(107) \times 10^{-5}$	$1.850(216) \times 10^{-5}$	$1.430(141) \times 10^{-5}$
1.771–1.871	$1.710(118) \times 10^{-5}$	$1.510(100) \times 10^{-5}$	$1.650(183) \times 10^{-5}$	$1.500(145) \times 10^{-5}$
1.871–1.971	$1.970(130) \times 10^{-5}$	$1.910(121) \times 10^{-5}$	$1.800(221) \times 10^{-5}$	$1.330(132) \times 10^{-5}$
1.971–2.071	$2.170(137) \times 10^{-5}$	$1.980(129) \times 10^{-5}$	$1.800(215) \times 10^{-5}$	$2.050(180) \times 10^{-5}$

Using the fit parameters, the reconstruction rate dependence can be corrected event-by-event, weighting each dimuon according to its event’s measured beam intensity, as in Equation 4.3.

The reconstruction rate dependence correction weight, w , is calculated independently for each roadset, target, and in each kinematic bin. There is no measurable difference in the fit parameters between different roadsets. Due to the relatively large differences in target ‘thickness’, the reconstruction rate dependence is significantly stronger for Deuterium than Hydrogen. This is expected since the thicker target will generate more background single muons, and the resulting higher occupancy in the chambers will adversely affect the reconstruction efficiency, even for the same beam intensity. The fits to the reconstruction rate dependence for θ and ϕ bins, for the four analyzed data sets, are shown in Figures 4.10 4.11 4.12 4.13 4.10 4.11 4.12 4.13.

The reconstruction rate dependence also varies with kinematic variables. Figure 4.18 shows the dependence of the rate dependence on the various dimuon kinematic variables. There is only a very weak, if any, dependence on most of the variables. The exceptions are x_1 and x_2 , which have a significant dependence, and x_F which has a very strong dependence. Recall that the pairs (x_1, x_2) and (x_F, M) are not independent. Therefore, to a good approximation, the dependence on x_1 and x_2 is not independent of the dependence on x_F . Since the dependence on M is not significant, it appears that the kinematic dependence of the rate-dependence is almost exclusively due to x_F . Therefore, in any analysis which requires the rate-dependence correction, it is advisable to simultaneously bin in x_F and the chosen analysis kinematics, in order to most accurately apply the correction.

4.2.2 Chamber and Hodoscope Rate-Dependent Efficiencies

In principle, the data can also suffer from rate-dependent effects caused by detector hardware inefficiencies. It is expected that the hodoscopes and wire chambers have some non-zero inefficiency, for which non-ratio analyses must make a correction. The ratio analyses can largely ignore these static inefficiencies since the




Figure 4.10: The reconstruction rate-dependent efficiency fits for Roadset 67, Deuterium target, in 10 bins of θ . Fit parameters are listed in Table 4.4




Figure 4.11: The reconstruction rate-dependent efficiency fits for Roadset 67, Hydrogen target, in 10 bins of θ . Fit parameters are listed in Table 4.4




Figure 4.12: The reconstruction rate-dependent efficiency fits for Roadset 62, Deuterium target, in 10 bins of θ . Fit parameters are listed in Table 4.4




Figure 4.13: The reconstruction rate-dependent efficiency fits for Roadset 62, Hydrogen target, in 10 bins of θ . Fit parameters are listed in Table 4.4




Figure 4.14: The reconstruction rate-dependent efficiency fits for Roadset 67, Deuterium target, in 10 bins of ϕ . Fit parameters are listed in Table 4.5




Figure 4.15: The reconstruction rate-dependent efficiency fits for Roadset 67, Hydrogen target, in 10 bins of ϕ . Fit parameters are listed in Table 4.5




Figure 4.16: The reconstruction rate-dependent efficiency fits for Roadset 62, Deuterium target, in 10 bins of ϕ . Fit parameters are listed in Table 4.5

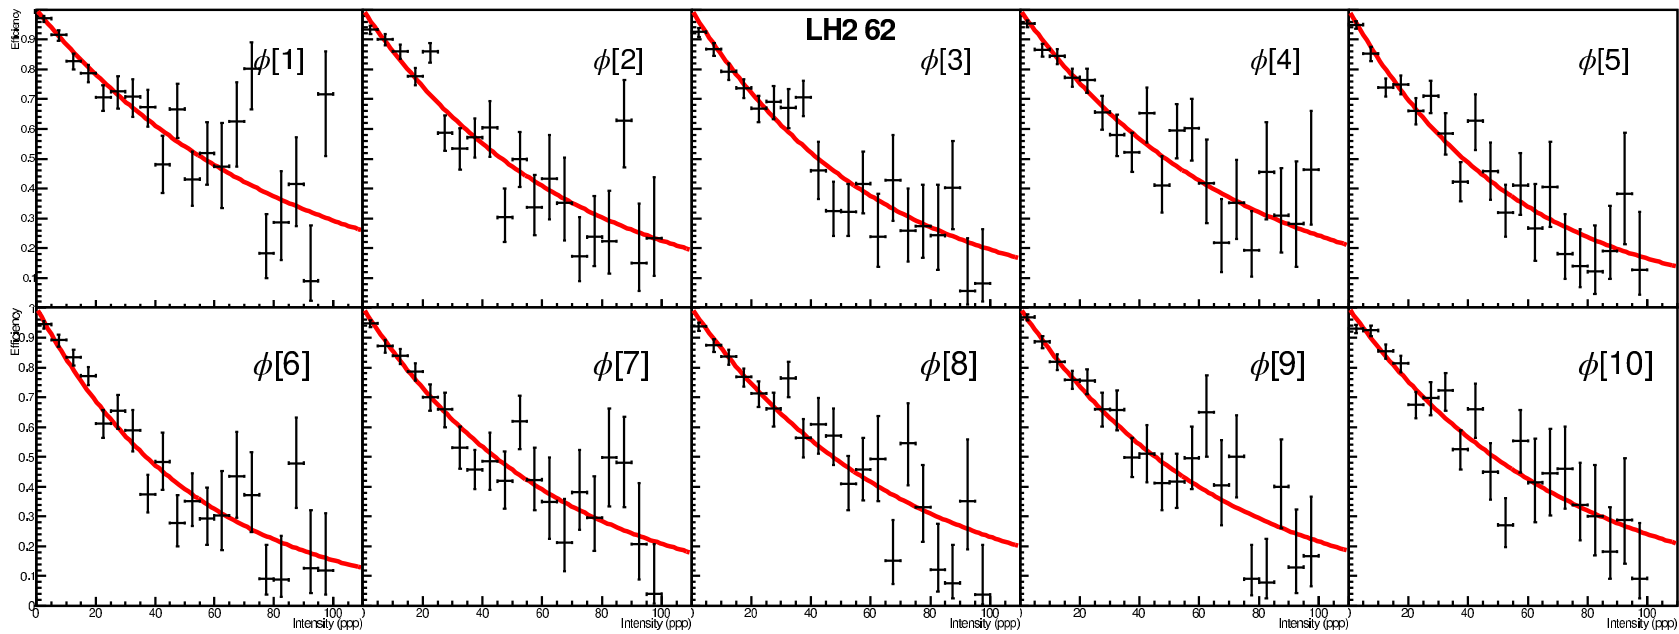


Figure 4.17: The reconstruction rate-dependent efficiency fits for Roadset 62, Hydrogen target, in 10 bins of ϕ . Fit parameters are listed in Table 4.5




Figure 4.18: The reconstruction rate-dependent efficiency fits for Roadset 67, Deuterium. Three equal-statistics bins are plotted for each kinematic variable. The kinematic dependence of the efficiency of the reconstruction is likely linked dominantly to x_F .

Table 4.5: Rate-Dependent Reconstruction Efficiency Parameters, “a”, ϕ -bins

ϕ bin	67 Deuterium	67 Hydrogen	62 Deuterium	62 Hydrogen
0.000–0.628	$1.690(78) \times 10^{-5}$	$1.440(67) \times 10^{-5}$	$1.40(13) \times 10^{-5}$	$1.23(9) \times 10^{-5}$
0.628–1.257	$1.650(74) \times 10^{-5}$	$1.480(68) \times 10^{-5}$	$1.71(15) \times 10^{-5}$	$1.49(10) \times 10^{-5}$
1.257–1.885	$1.840(83) \times 10^{-5}$	$1.570(70) \times 10^{-5}$	$1.43(13) \times 10^{-5}$	$1.63(10) \times 10^{-5}$
1.885–2.513	$1.890(84) \times 10^{-5}$	$1.750(78) \times 10^{-5}$	$1.82(16) \times 10^{-5}$	$1.42(9) \times 10^{-5}$
2.513–3.142	$1.980(88) \times 10^{-5}$	$1.760(77) \times 10^{-5}$	$1.84(15) \times 10^{-5}$	$1.80(11) \times 10^{-5}$
3.142–3.770	$2.090(92) \times 10^{-5}$	$1.910(85) \times 10^{-5}$	$2.37(19) \times 10^{-5}$	$1.87(12) \times 10^{-5}$
3.770–4.398	$1.850(82) \times 10^{-5}$	$1.750(78) \times 10^{-5}$	$1.72(15) \times 10^{-5}$	$1.57(11) \times 10^{-5}$
4.398–5.027	$1.810(81) \times 10^{-5}$	$1.760(78) \times 10^{-5}$	$1.74(15) \times 10^{-5}$	$1.46(9) \times 10^{-5}$
5.027–5.655	$1.680(76) \times 10^{-5}$	$1.660(74) \times 10^{-5}$	$1.69(14) \times 10^{-5}$	$1.53(10) \times 10^{-5}$
5.655–6.283	$1.740(79) \times 10^{-5}$	$1.550(69) \times 10^{-5}$	$1.73(15) \times 10^{-5}$	$1.42(10) \times 10^{-5}$

acceptance effects they induce should be independent of the target material. However, since the targets do not all have equal interaction lengths, rate-dependent hardware efficiencies can easily cause systematic biases.

The possibility of rate-dependence in the hodoscopes was the impetus for a significant upgrade to the Station 1 and 2 hodoscopes. A transistorized base was designed for the Station 1 and 2 PMTs, and these new bases replaced all the old resistor-capacitor bases. A comparison of count rate before and after this upgrade demonstrates that the new bases alleviated the rate-induced “sag” in the hodoscopes.

Determination of the hodoscope efficiencies is especially important for non-ratio analyses because any inefficiency will affect the trigger, and may affect the effective trigger acceptance. Unfortunately, the absolute efficiencies for the hodoscopes during each period of data-taking have not yet been well measured. An effort is currently underway to extract the hodoscope efficiencies using calibration data which has been taken alongside the physics data throughout SeaQuest’s operation. The final product of the hodoscope efficiency analysis should be a set of rate-dependent, time-dependent efficiencies which will be used to modify the trigger acceptance correction.

Due to independence from the trigger decision, the chamber efficiencies have a much weaker effect on the acceptance. The reconstruction algorithms are flexible enough that dimuons can be successfully reconstructed even if many hits are missing. The extracted reconstruction efficiencies in figure 4.19 show that the rate-dependence of the reconstruction efficiency is independent of an artificially-applied rate-dependent chamber efficiency. In one sample, Monte Carlo chamber hits were randomly dropped without regard to the intensity, resulting in a 94%, rate-independent chamber efficiency. In the other sample, a very strongly rate-dependent artificial chamber efficiency was applied, dropping from 95% at zero intensity down to 75% at $I_p = 100,000$ ppp. There is no significant difference between these two samples. Therefore, even a very severe rate-dependent chamber inefficiency does not adversely affect the successful reconstruction of dimuons.




Figure 4.19: The reconstruction rate-dependent efficiency is independent of the chamber efficiency

The rate-dependent chamber efficiencies have been checked. The efficiencies are well within the acceptable bound, and should not affect the reconstructed dimuon sample at all. A correction for the chamber efficiencies is not applied in any current SeaQuest analyses.

4.3 Drell-Yan Dimuon Selection

4.3.1 Quality Cuts

Selecting a sample of high-quality dimuons requires cutting away a large fraction of the reconstructed dimuons. Any rate-dependent inefficiency for the DY signal caused by these cuts is corrected for by the rate-dependent efficiency correction calculated in Section 4.2.1. A large number of quality cuts have been collected and are applied to the reconstructed data prior to any physics analysis.

Spill Cuts

Table 4.6 lists the spill-level cuts applied to the data before physics analysis. In addition to these, all listed quantities must exist exactly once in each spill. Any spills with duplicated or missing entries are not considered for analysis. Table 4.7 lists four excluded spill ranges and the reason that these spills are excluded from the analysis.

Event Cuts

The only event cut is the requirement: $MATRIX1=1$. That is, only events in which the Trigger Supervisor registered a hit from the MATRIX1 trigger (the physics dimuon trigger) are considered for analysis.

Table 4.6: Spill-Level Cuts

Variable	Requirement
TargetPos	(spill table)=(target table)
TargetPos	[1-7]
TSGo	[100-6000]
AcceptedFPGA1	[100-6000]
AfterInhFPGA1	[100-10000]
AcceptedFPGA1/AfterInhFPGA1	[0.2-1.05]
G2SEM	[2e12-1e13]
QIESum	[4e10-1e12]
Inhibit	[4e9-2e11]
Busy	[4e9-1e11]
Duty Factor	[10-60]

Table 4.7: Bad Spill Ranges

Range	Description
416207-424180	Manual Target Motion
371870-376533	Trigger Timing Shift
378366-379333	Trigger Timing Shift
482574-484924	Reversed FMag Field

Track Cuts

Table 4.8 lists the track-level quality cuts which are applied to reconstructed data (and reconstructed Monte Carlo) before analysis. For the tracks, N_{dof} is numHits-5.

Dimuon Cuts

Table 4.9 lists the reconstructed dimuon quality cuts which must be applied to reconstructed data (and reconstructed Monte Carlo) before any physics analysis.

Target/Dump Selection

To facilitate efficient target/dump separation in analysis, the re-tracking step in the event reconstruction is run two separate times. The re-tracking is done once forcing the track to pass through the most likely target vertex position, and the χ^2/N_{dof} of this fit is stored as ‘chisq.target’. Then the re-tracking is attempted

Table 4.8: Track Cuts

Requirements
numHits > 14
$\chi^2/N_{dof} < 6$
numHits = 18 \vee pz > 18

Table 4.9: Dimuon Cuts

Variable	Requirement
$ dx $	< 2
$ dy $	< 2
$ dpx $	< 3
$ dpy $	< 3
dpz	[30–120]
mass	[0–10]
x_T	[0–1]
x_F	[-1–1]
x_B	[0–1]
trackSeparation	< 175
χ^2	< 10
p_x^+	> 0
p_x^-	< 0
$\text{sgn}(\text{roadID1}) \times \text{sgn}(\text{roadID2})$	> 0

Table 4.10: Target/Dump Selection

	Variable	Requirement
Target	dz	(-300 – -60)
	chisq_dump - chisq_target	> 10
Dump	dz	(0 – 150)
	chisq_target - chisq_dump	> 10

again, this time forcing the track to pass through the most likely dump vertex position. The χ^2/N_{dof} is stored as ‘chisq_dump’. Then, each dimuon can be identified as ‘target’, ‘dump’, or ‘ambiguous’, based on the criteria in Table 4.10. Dimuons are ‘target/dump ambiguous’ if $|\chi_{\text{target}}^2 - \chi_{\text{dump}}^2| < 10$. These dimuons are not used in any current analyses.

4.3.2 Mass Spectrum Fitting

The four major components of the reconstructed dimuon sample, Drell-Yan, J/Ψ , Ψ' , and combinatorial background, each has a quite distinct mass distribution shape. For this reason, it is possible to determine the contributions from each component by fitting the total mass spectrum with a simple model of four assumed mass distribution shapes, letting the normalization for each component float.

For the three ‘physics’ components, the mass distribution shapes are most easily determined through Monte Carlo. Using large Monte Carlo samples, the geometric and trigger acceptances are applied, and the events are reconstructed with kTracker. Finally, track and dimuon cuts are applied to the reconstructed results, equivalently to real data. From this sample of reconstructed Monte Carlo dimuons, we can generate mass distributions for each of the physics components, cut into arbitrary kinematic bins, if desired.




Figure 4.20: Component-fitted mass spectrum for reconstructed dimuon events from Roadset 62, Deuterium target

The mass distribution of the combinatorial background component is more difficult to determine. Simulating the combinatorial background observed at SeaQuest has proved prohibitively difficult. Therefore, we must use data. The default solution is to use reconstructed tracks from events triggered by a “singles” trigger. This trigger requires only at least one trigger road to be satisfied in order to fire. However, the experiment is deadtime limited, and the prescaling factor on the singles triggers was set high enough that the reconstructed sample of singles is rather small. Fortunately, most of the events which satisfied the dimuon trigger do not actually contain a reconstructable dimuon, but many of them contain one reconstructed muon. Selecting only those events which fired the dimuon trigger, but only contained a single good track after reconstruction, we can form a representation of the combinatorial background by randomly pairing opposite-sign tracks together. Applying the vertexing algorithm and quality cuts as normal, this ‘Mix’ sample becomes a decent representation of the reconstructed combinatorial background.

A four-component fit to the reconstructed mass distribution, using the four mass shapes for Drell-Yan, J/Ψ , Ψ' , and Mix, yields qualitative agreement between the model and the data. From such fits, the

contribution from each component to the total dimuon sample can be immediately read off. To extract the yields for a particular component in kinematic bins, an independent mass fit for each bin must be performed. The mass shapes of the components within each kinematic bin are easy to obtain since the events from each component can be binned as well. It is important to determine the component shapes independently for each bin since they can vary significantly for different kinematic bins.

4.3.3 Mass Cut

For Drell-Yan analyses, a much simpler method can be used to isolate Drell-Yan events from the J/Ψ , Ψ' , and combinatoric backgrounds. As demonstrated by the mass fits, when the dimuon mass is greater than 4.2 GeV the Drell-Yan component is the largest contributor. Therefore, applying a cut at $\text{Mass} > 4.2 \text{ GeV}$ results in a dimuon sample that is 75% Drell-Yan signal. A rigorous subtraction of the combinatoric background will be performed for all analyses before the final results are published.

However, as discussed in detail in Section 4.5.1, it is clear that the Mix Background component becomes increasingly important nearer to $\theta = 0$ and $\theta = \pi$, even above 4.2 GeV. The data at the edges of the θ distribution have a much stronger influence over the angular parameter fit and extraction than the $\theta \sim \frac{\pi}{2}$ data do. Unfortunately, the experimental acceptance drops rapidly away from $\theta = \frac{\pi}{2}$. These three effects conspire to make the extraction of λ , μ , and ν especially difficult. Therefore, trading slightly increased statistical uncertainties for a smaller systematic effect of the combinatorial contribution, the angular analysis uses a mass cut of 5.0 GeV to isolate DY dimuon events. From the mass spectrum fit, this mass cut should yield a data sample that is $\sim 85\%$ Drell-Yan, and $\sim 15\%$ background.

4.4 Acceptance Correction

In principle, extracting the angular parameters, λ , μ , and ν from a sample of Drell-Yan dimuon data is fairly straightforward. One needs only to plot the data as a function of θ and ϕ , and then fit the histogram to the function:

$$\frac{d\sigma}{d\theta d\phi} \propto \sin \theta \left(1 + \lambda \cos^2 \theta + \mu \sin 2\theta \cos \phi + \frac{\nu}{2} \sin^2 \theta \cos 2\phi \right) \quad (4.4)$$

In practice, however, experiments never have a perfect 4π detector. Angular distributions can be quite sensitive to acceptance effects. Particularly, in fixed-target, forward, Drell-Yan experiments like SeaQuest, the overall acceptance is rather small. Figure 4.21 shows the single-track acceptance for Monte Carlo generated, high-mass, Drell-Yan dimuons as a function of single-track 3-momenta. The positive and negative muon acceptance regions are clearly separated, due to the focusing effect of the two magnets.




Figure 4.21: Single-track momentum acceptance for high-mass DY Monte Carlo dimuons

Furthermore, the θ acceptance is rather narrow, and centered around $\theta = \frac{\pi}{2}$. This is unfortunate, because the extraction of the angular parameters, especially λ , is much more sensitive to the low-acceptance regions of the θ distribution.

The experimental acceptance in ϕ is shown in Figure 4.23. The geometrical effect of the forward spectrometer alone introduces a significant $\cos 2\phi$ component. Coincidentally, the acceptance effect due to the Top/Bottom split in the Matrix1 trigger almost exactly cancels the phi-dependence introduced by the geometric acceptance. Thusly, the resulting ϕ acceptance of the triggered dimuon data is almost flat.

4.4.1 Geant-based Monte Carlo (GMC)

The angular distribution analysis relies entirely on Monte Carlo simulated data to correct for the acceptance effects.

The Geant4 based Monte Carlo (GMC), maintained by Bryan Kerns, is the primary simulation code for the SeaQuest experiment. The dimuon generator in the GMC is based largely on the event generator from the older, FORTRAN-based Monte Carlo (FastMC), used in the previous fixed-target dimuon experiments at FNAL. The event generators in both the GMC and FastMC have been updated with more recent PDF

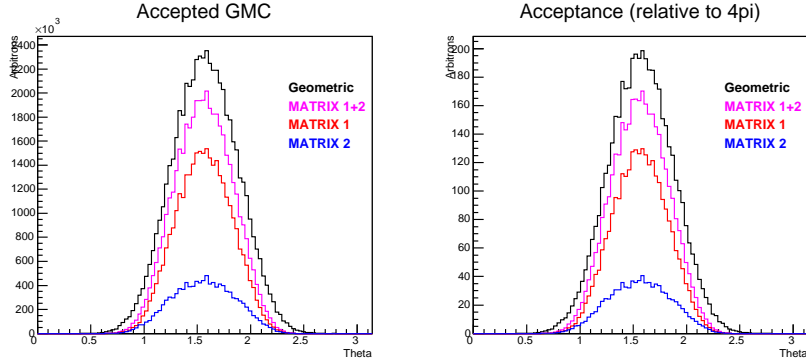


Figure 4.22: Geometric and trigger acceptances in θ . “Matrix 1” is the main physics trigger: opposite sign muons, on opposite sides (Top/Bottom) of the spectrometer. “Matrix 2” is the opposite sign, same side trigger. Therefore, the difference in acceptance between these triggers demonstrates the effect of the Top/Bottom separation in the trigger system.

fits and have been modified for the 120 GeV Main Injector beam energy used by SeaQuest. All previous FNAL fixed-target Drell-Yan experiments used the 800 GeV Tevatron proton beam. The GMC dimuon event generator can produce dimuon events for Drell-Yan, J/Ψ , and Ψ' production from beam proton interactions in the liquid targets, solid targets, and beam-dump.

The GMC uses Geant4 to simulate the propagation of muons through a detailed model of the SeaQuest spectrometer. Geant4 handles the effects of energy-loss in materials, multiple-scattering in materials, and deflection due to magnetic fields. The fields of FMAG and KMAG are interpolated from detailed field maps calculated by simulation.

Every time a detector component is moved, a new survey of the spectrometer is performed. The survey provides absolute space points for a few points on each detector frame. Using reconstructed magnet-off data, the track reconstruction programs are able to very precisely calculate the relative positions of all detector components in the experimental hall. For each period of data-taking, both the survey information and the ‘alignment’ information are used to generate a ‘geometry’, which specifies the positions and angles of all components of the spectrometer. The GMC uses these geometries to define the spectrometer model when simulating dimuon data. Therefore, the GMC is able to generate simulated data with a very detailed geometric model of the spectrometer separately for each SeaQuest run period. This detailed geometry simulation is required for good acceptance correction in any analysis measuring absolute cross-sections or shapes of kinematic distributions, such as the angular distribution analysis.

The GMC output consists of MySQL tables which correspond exactly to the format of the real data, in addition to other tables which hold the GMC ‘truth’ information. This means that the reconstruction programs can run on simulated data and real data semi-agnostically. The reconstructed Monte Carlo is

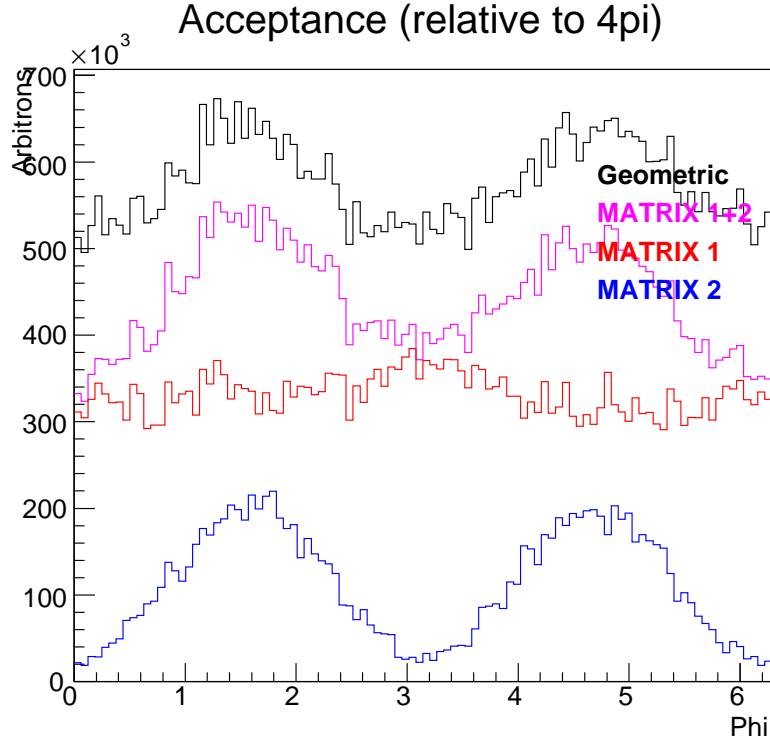


Figure 4.23: Geometric and Trigger Acceptances in ϕ . “Matrix 1” is the main physics trigger: opposite sign muons, on opposite sides (Top/Bottom) of the spectrometer. “Matrix 2” is the opposite sign, same side trigger. Therefore, the difference in acceptance between these triggers demonstrates the effect of the Top/Bottom separation in the trigger system.

critically important for the calculation of acceptance corrections for non-ratio analyses. It would also be quite useful for measuring reconstruction-induced biases which may also affect ratio analyses. Studies of such effects, which are assumed to be quite small, have not yet been performed.

Another mode of the GMC, which does not make use of the physics event generators mentioned above, is the ‘Gun’. In gun-mode, the GMC simple uses Geant4 to simulate the interactions of 120 GeV protons incident to the spectrometer from the beamline. The Geant4 simulation does not include such small cross-section interactions as Drell-Yan or heavy-quark production, but it does include pion-production. Therefore, it generates a sample of pion-decay single muons, which should represent the primary source of combinatoric background in SeaQuest data. In practice, due to the highly-irregular microstructure of the slow-spill extracted MI beam, simulation of the combinatoric background from coincident singles is quite difficult. A plausible simulation of the SeaQuest combinatoric backgrounds using the Gun GMC has not yet been achieved.

4.4.2 Bin-level, Dimuon-based Acceptance Correction

The simplest method for performing the acceptance correction is a bin-by-bin correction, using an acceptance calculated from cross-section weighted GMC Drell-Yan events. In this method, an acceptance factor is calculated for each bin in the analyzed variables. This requires two GMC productions: accepted and 4π . The first GMC sample contains only events which are in the geometric acceptance, passed the trigger requirements, and were successfully reconstructed by kTracker, passing the same quality cuts applied to the data. The second is a set of generated GMC events with no acceptance cuts applied representing the 4π sample of Drell-Yan dimuons generated by beam proton interactions in the target. The acceptance is then just the ratio of the reconstructed events to the 4π events, binned in the analysis bins. Dividing the intensity-corrected data by the GMC-calculated acceptance yields the acceptance corrected distribution.

Ideally, the acceptance correction would be done in a multi-dimensional binning. Binning the data and correction in as many kinematic variables as possible, especially those variables with a strong dependence on the acceptance, is preferable since correlations between the different kinematic variable are handled directly. If the acceptance correction is done by integrating over a kinematic variable with a strong acceptance effect, which is correlated to the analysis variable, then is it possible that small differences between the GMC-simulated distributions and the true distributions can cause a systematic bias in the acceptance correction.

Due to the limitations of time and the statistical precision of the data, the acceptance correction for the angular distributions analysis has only been binned in θ and ϕ .

4.5 Collins-Soper θ Distribution Results

As a simple first step into the angular analysis, I present the application of the analysis techniques described in the previous sections to the extraction of λ from the θ distributions of the D_2 and H_2 liquid target datasets.

4.5.1 Corrected θ Distribution

After the recorded data has been decoded and reconstructed, and the spill, event, dimuon, and track quality cuts have been applied, and the target selection cuts and Drell-Yan selection cuts have been applied, we are left with the ‘raw’ analysis data. Finally, the required corrections to the data can be applied, and physics result extractions can begin.

Figure 4.24 shows the uncorrected data on the D_2 and H_2 targets, for the Roadset 62 and Roadset 67 data-sets. The shape of the uncorrected θ distribution is due almost entirely to the geometric and




Figure 4.24: Uncorrected high-mass dimuon data distribution in θ

trigger acceptances, and reconstruction bias. These three effects are all corrected at once by the acceptance correction.

The first analysis correction is for the rate-dependent kTracker efficiency. This correction is applied to each dimuon individually, using the measured beam intensity associated with each event. The dimuons are weighted by the factor described in Section 4.2.1, which is calculated separately for each θ -bin. Therefore, the slight θ dependence in the kTracker reconstruction rate-dependent efficiency is corrected. The left panels of Figures 4.28 4.29 4.30 4.31 compare the data before and after this correction is applied. As expected, the magnitude of the corrected counts is much larger than the uncorrected counts. It is not yet obvious that the shape of the distribution has changed, but the acceptance correction makes the shape differences clear.

The analyzed Monte Carlo data provides the information necessary to correct for the geometric, trigger, and reconstruction acceptances in one step, as described in Section 4.4.2. The shapes of the kTracked, high-mass, Drell-Yan GMC events are shown in Figure 4.25, for the two liquid targets and the two analyzed roadsets.

The first observation is that the tails of these distributions are rather dramatically different than the




Figure 4.25: Total Acceptance in θ , from GMC

tails observed in the data. This difference is due to the contribution of combinatoric background ‘dimuons’ in the data sample, which are, of course, not present in the GMC Drell-Yan dimuon sample. Figure 4.26 shows the θ distributions for the ‘Mix’ background samples, requiring $\text{Mass} > 5 \text{ GeV}$. Unfortunately, the peaks of the ‘Mix’ exactly correspond to the tail of the data distributions, where they can easily dominate the relatively sparse signal events. Furthermore, the tails of the θ distribution provide the largest lever-arm in the angular fit, and are crucially important for an accurate extraction of the angular parameters.

It is not yet possible to apply the rate-dependent reconstruction efficiency correction to the ‘Mix’ background sample. Therefore, to determine signal/background ratios in the θ distributions, we must use the uncorrected data distributions, along with the kTracked GMC data and the ‘Mix’ background sample. This comparison is somewhat dubious since the kTracked GMC data does not suffer from the rate-dependent effects, but the data and the mix samples do. For now, however, it is the best we can do. Figure 4.27 shows the relative contributions of the signal and background to the data vs θ , using the 15% Mix and 85% GMC normalization obtained with the mass-spectrum fit to the $\text{Mass} > 5 \text{ GeV}$ data. This breakdown of contribution is not perfect because the data and mix do not have the rate-dependent reconstruction efficiency

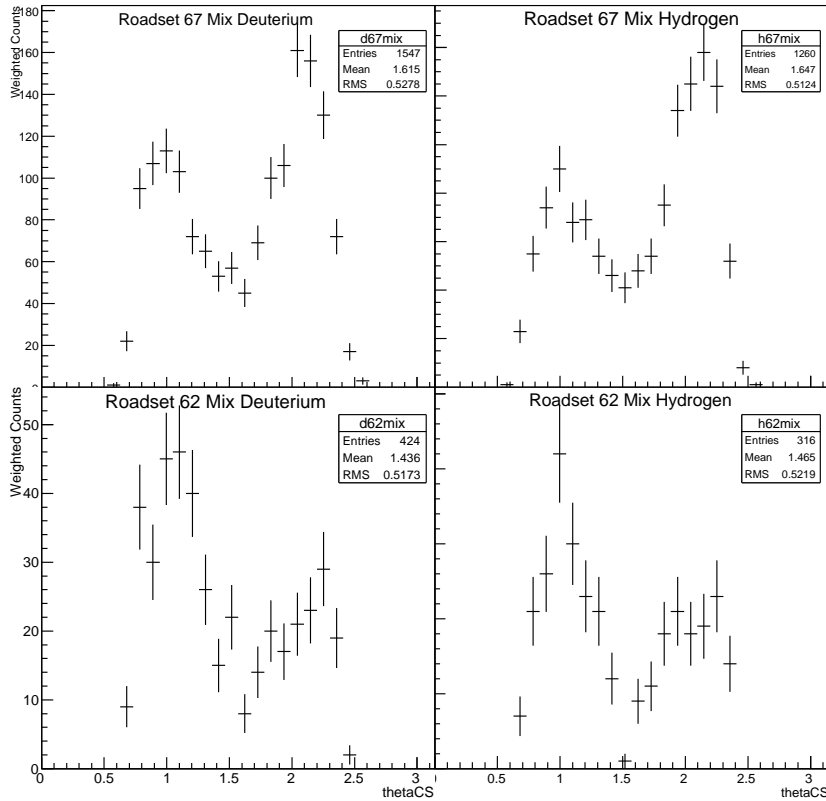


Figure 4.26: Mix Background distribution in θ

correction applied, and the GMC does not suffer from the reconstruction inefficiency. To ensure dominance of the signal in each bin, we must use the fiducial cuts: $1.1 \lesssim \theta \lesssim 2.1$.

Unfortunately, the inapplicability of the reconstruction rate dependence correction to the ‘Mix’ background means that it is currently impossible to correctly subtract the combinatorial background out of the corrected data. The subtraction cannot take place before the rate-dependence correction because the correction must be applied at the event level. Options for resolving this unfortunate limitation are discussed in Section 4.7.1.

Applying the fiducial cuts, the rate-dependence correction, and the acceptance correction, we finally arrive at the θ distributions, from which we can attempt to extract λ .

4.5.2 Extracted λ

The extraction of the polar-angle parameter, λ , from the final θ distributions is straightforward. The θ distributions are fit to the function:

$$\sin \theta (1 + \lambda \cos^2 \theta) \tag{4.5}$$

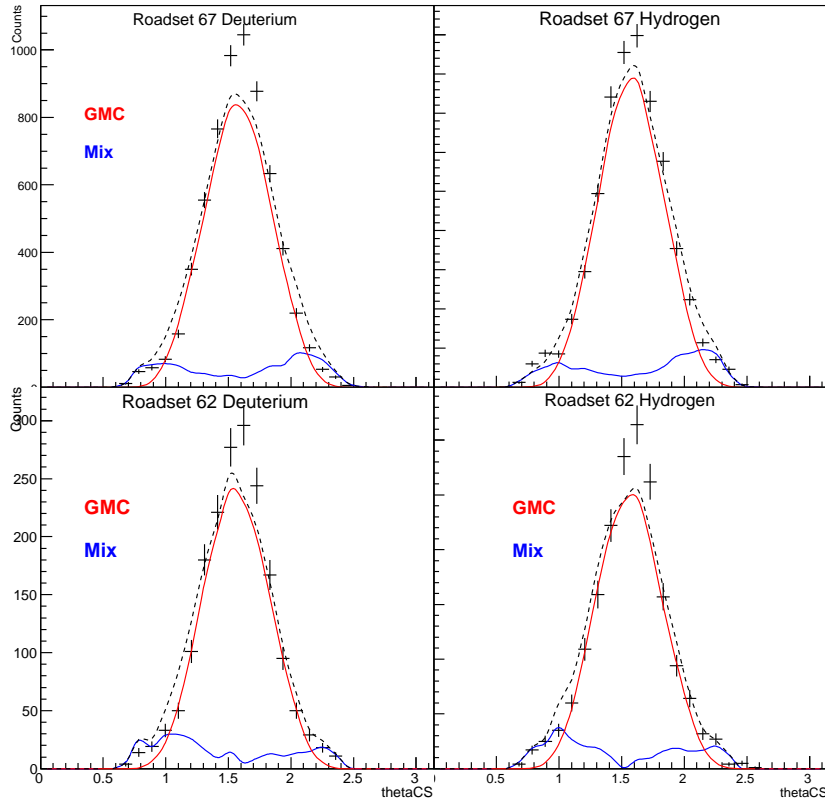


Figure 4.27: High-Mass DY θ distribution components

with λ as the only free parameter. The final θ distributions, and their associated fits and extracted λ values, are shown in Figures 4.28 4.29 4.30 4.30 for the four analyzed datasets. The extracted λ values and uncertainties are tabulated in Table 4.11.

Although angular distributions from proton-induced Drell-Yan have never before been measured at $\sqrt{s} \sim 15$ GeV, results from NA10 and E866 strongly favor the expectation that the λ -values measured at SeaQuest should be quite close to unity. The values extracted here are (mostly) consistent with that expectation in all four analyzed samples.

Roadset	Target	λ
67	Deuterium	0.804 ± 0.305
67	Hydrogen	2.050 ± 0.378
62	Deuterium	0.496 ± 0.539
62	Hydrogen	1.527 ± 0.630

Table 4.11: Extracted λ_{CS} results from the four analyzed datasets.




Figure 4.28: Application of Intensity and Acceptance Corrections to the θ distribution for Roadset 67, Deuterium target




Figure 4.29: Application of Intensity and Acceptance Corrections to the θ distribution for Roadset 67, Hydrogen target

4.6 ϕ Distribution

In preparation for attempting a full extraction of λ , μ , and ν from the two-dimensional angular distributions, we first check the one-dimensional ϕ distributions, and compare them to Monte Carlo. Keep in mind that the GMC is thrown with $\lambda = 1$, and $\mu = \nu = 0$, and so ϕ is thrown flat. Coincidentally, the acceptance in ϕ is also nearly flat. Figure 4.32 shows the ϕ acceptance for the four analyzed roadsets. Any significant azimuthal dependence in the final corrected ϕ distribution from data are either due to physics (Boer-Mulders!) or inadequate corrections.

The raw ϕ distributions are presented in Figure 4.33. In addition to a significant $\cos 2\phi$ component, the distributions are noticeably non-symmetric about π . Such an asymmetry cannot be generated by known physics, so it must reflect an effect of the acceptance.

Before correcting the data and interpreting any remaining $\cos 2\phi$ components as evidence of interesting physics, it is important to again consider the expected influence of the combinatoric background. Figure 4.34




Figure 4.30: Application of Intensity and Acceptance Corrections to the θ distribution for Roadset 62, Deuterium target




Figure 4.31: Application of Intensity and Acceptance Corrections to the θ distribution for Roadset 62, Hydrogen target

shows the ϕ distributions for the Mix Background sample from the four analyzed data sets. Unfortunately, the Mix sample has an *extreme* $\cos 2\phi$ component. Any observed $\cos 2\phi$ dependence in the corrected ϕ distribution cannot be attributed to physics until a satisfactory combinatoric background correction has been applied.

Figure 4.35 shows the relative contributions of the signal and background to the data vs ϕ , using the 15% Mix and 85% GMC normalization obtained with the mass-spectrum fit to the Mass > 5 GeV data. Despite the fact that the rate-dependent reconstruction efficiency has not been applied to either the Data or the Background in these plots, the data and the GMC+Background qualitatively agree. These plots are a clear demonstration that a very careful combinatoric background subtraction will need to be performed before a reliable ν -parameter extraction can be attempted.

Figures 4.36 4.37 4.38, and 4.38 demonstrate the application of the kinematically-binned, intensity dependent kTracker efficiency correction and the acceptance correction to the ϕ distributions for the four




Figure 4.32: Total Acceptance in ϕ , from GMC

analyzed data sets. The rate-dependent efficiency correction does alter the shape of the ϕ distributions, but not dramatically. In the center panels, the acceptance is plotted, normalized to the efficiency-corrected data distribution. As noted earlier, the acceptance corrections in ϕ are nearly flat. The split third panels show the acceptance corrected ϕ distributions, both with and without the efficiency correction. In each, the fit is simply to the function $1 + \alpha \cos 2\phi$ and does not imply a certain value for the ν parameter. The significantly non-zero value of α is assumed to be due to the azimuthal asymmetry of the combinatoric background. Implications on the ν parameter for $\sqrt{s} \sim 14$ GeV proton-induced Drell-Yan must wait for a rigorous combinatoric background subtraction.

Roadset	Target	α
67	Deuterium	0.212 ± 0.019
67	Hydrogen	0.201 ± 0.034
62	Deuterium	0.209 ± 0.034
62	Hydrogen	0.124 ± 0.035

Table 4.12: Extracted α results from the four analyzed datasets. α is qualitatively related to ν_{CS} .

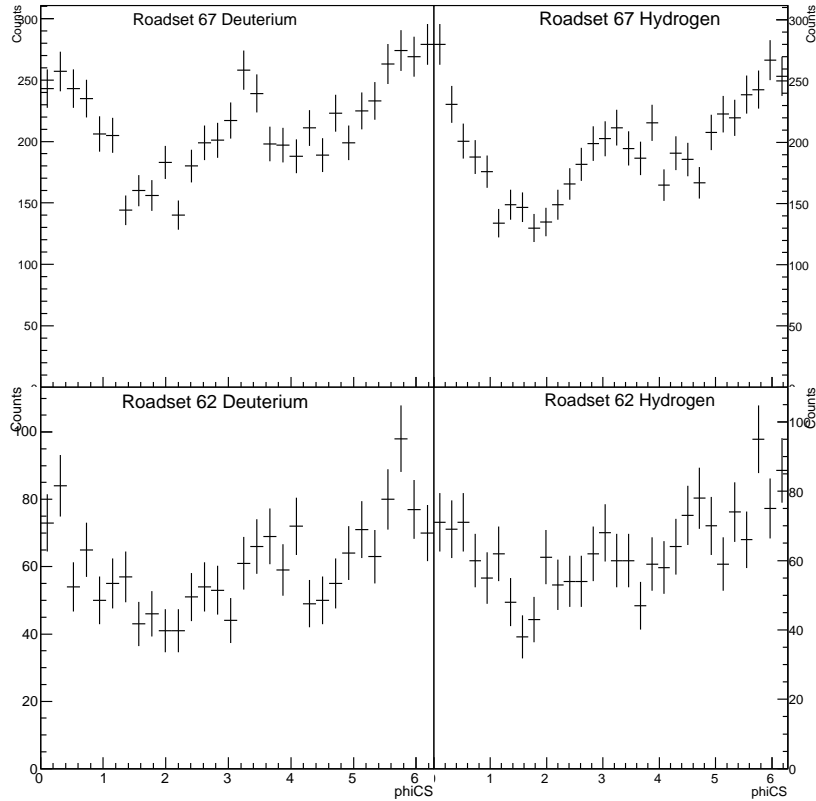


Figure 4.33: Uncorrected DY-region Data in ϕ

Table 4.12 tabulates the extracted values of α from the ϕ -distributions for the four analyzed datasets. The extracted α is qualitatively related to ν_{CS} but is currently dominated by effects from the combinatorial background.

4.7 Future Improvements

4.7.1 Rate-Dependent Reconstruction Efficiency for Mix Background

The influence of the combinatoric background can have a strong effect on the measured distribution of θ in dimuon events with $M > 5$ GeV. Therefore, an accurate extraction of λ requires subtracting out the contribution of combinatoric ‘dimuons’ from the Drell-Yan data sample. The more complex extraction of λ , μ , and ν from the two-dimensional θ and ϕ distribution will certainly also require a Drell-Yan dimuon sample with negligible combinatoric contamination.

The simplest procedure for removing the combinatoric ‘dimuons’ is simply to subtract them. The ‘Mix’ background sample generated by kTracker should be a good representation of the kinematic dependencies of




Figure 4.34: DY-region Mix Background in ϕ

the combinatoric background component of the data. The mass-spectrum fitting procedure, when applied to each dataset as a whole, produces a reasonably precise measurement of the relative contributions of signal and background to the mass distribution of the data. It is easy to apply a mass cut to the mass spectrum fitting output, to obtain the relative contributions of Drell-Yan signal and ‘Mix’ background above a certain mass (e.g. 5 GeV). With the fractional contribution in-hand, one can simply plot the mix background in the analysis bins (e.g. θ), scale the mix distribution such that integral matches the expected fraction of the integral of the data and subtract the result from the data.

The fatal flaw with this approach is that the rate-dependent reconstruction efficiency correction has been applied to all of the data dimuons. The true signal dimuons have been adjusted for reconstruction efficiency, as desired, *but so have the combinatoric-background-induced ‘dimuons’*. Therefore, in order to subtract the background out of the data, a rate-dependent, kinematics-dependent reconstruction efficiency correction must also be calculated for, and applied to, the ‘Mix’ background sample. It is not yet completely clear how to accurately measure this correction for the ‘Mix’ background sample. A simple first attempt might look like the following:

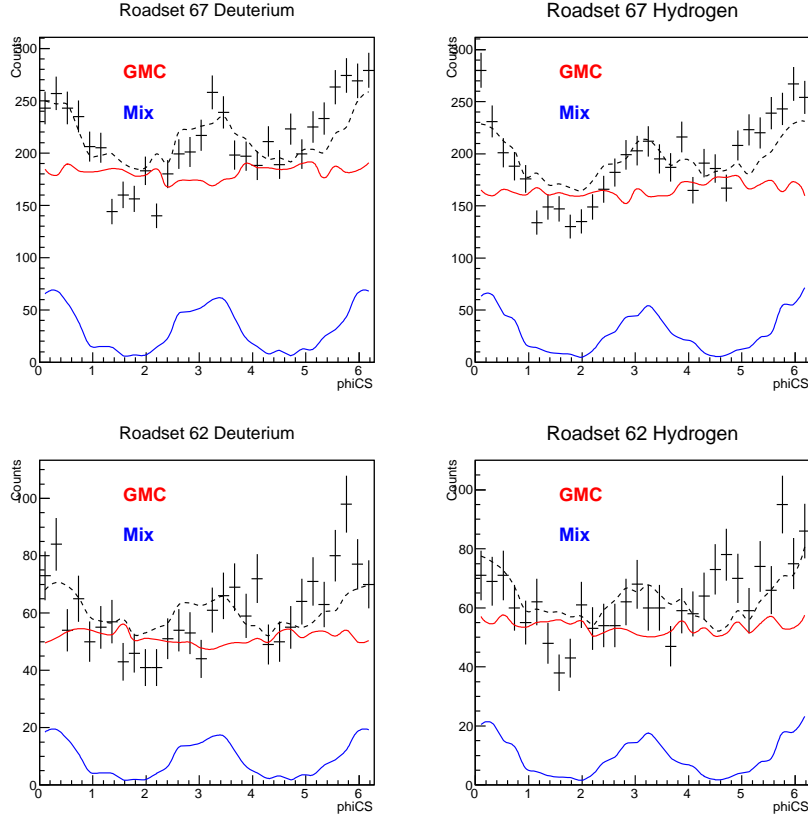


Figure 4.35: High-Mass DY ϕ distribution components

First, assume that the rate-dependence of the kTracker reconstruction is due entirely to the track-finding and track-fitting steps. Stated another way, assume that the vertexing step has no rate-dependent inefficiency. This assumption allows for the dimuon reconstruction efficiency to be factorized into a rate-dependent reconstruction inefficiency for each individual track.

$$\epsilon_{\text{dimuon}}(I) = \epsilon_{\text{track}}(I_1) \epsilon_{\text{track}}(I_2) \quad (4.6)$$

For the efficiency correction on the real data, $I_1 = I_2$ since both tracks necessarily come from the same event. Therefore, $\epsilon_{\text{track}}(I) = \sqrt{\epsilon_{\text{dimuon}}}$. Applying this to the ‘Mix’ sample:

$$\epsilon_{\text{dimuon}} = \sqrt{\epsilon_{\text{dimuon}}(I_1) \epsilon_{\text{dimuon}}(I_2)} \quad (4.7)$$

So, using the efficiencies calculated by the GMC+NIM3 embedding procedure, the efficiency of a mixed dimuon is simply the geometric mean of the efficiencies from the two events.

In the extraction of the rate-dependent reconstruction efficiencies from the GMC+NIM3 embedding, we




Figure 4.36: Application of intensity and acceptance corrections to the ϕ distribution for Roadset 67, Deuterium target




Figure 4.37: Application of intensity and acceptance corrections to the ϕ distribution for Roadset 67, Hydrogen target

assume a shape of $\epsilon_{\text{dimuon}} = e^{AI}$, where A is the only free parameter. We assume that A depends only on target position, spectrometer configuration, and dimuon kinematics. That is, the two tracks that contribute to a ‘Mix’ dimuon share the same value for A . Under this assumption:

$$\epsilon_{\text{dimuon}} = e^{AI_{\text{dimuon}}} = e^{\frac{1}{2}AI_1} e^{\frac{1}{2}AI_2} = e^{\frac{A}{2}(I_1+I_2)} \quad (4.8)$$

Therefore, $I_{\text{dimuon}} = \frac{1}{2}(I_1 + I_2)$; the effective intensity of a ‘Mix’ dimuon is the arithmetic mean of the intensities of the events of the contribution tracks.

With the above prescription, it would be possible to correct the ‘Mix’ dimuons for the rate-dependent reconstruction efficiency, using the intensities from the two events which each contributed one track to the ‘dimuon’. However, there are some issues with this approach. Firstly, it is not obvious that the efficiency dependence on dimuon kinematic variables can survive the track factorization. Secondly, in the current version of the reconstruction programs, the association between the tracks contributing to a ‘Mix’ dimuon




Figure 4.38: Application of intensity and acceptance corrections to the ϕ distribution for Roadset 62, Deuterium target




Figure 4.39: Application of intensity and acceptance corrections to the ϕ distribution for Roadset 62, Hydrogen target

and their original parent events is lost. In a future pass of the reconstruction, it would be quite easy to include original event information along with each track in the ‘Mix’ sample.

Chapter 5

Phenomenology of Drell-Yan Angular Distributions

This chapter details an intuitive methodology for understanding the non-trivial phenomenology of the QCD-improved Drell-Yan mechanism. Geometrically-derived, single-event definitions for the angular distribution parameters are obtained. First, this methodology is used to interpret the recently published CMS data on γ^*/Z dilepton production. The relative contributions from the $q\bar{q}$ -annihilation process and the qg -bremsstrahlung process are extracted from the p_T -dependence of λ . The influence of higher-order diagrams, through non-coplanarity of the quark axis with the hadron plane, is extracted through the p_T -dependence of ν . Next, under certain simplifying assumptions, the parameter definitions are used to derive the Lam-Tung relation. By varying the assumptions, many other parameter relationships are obtained, including some which are absent in the published literature. The parameter definitions also imply a triangular region in the λ - ν plane which should circumscribe experimentally extracted parameters. Excitingly, significant Boer-Mulders effects and significant non-transverse photon components are the only means by which measured angular parameters may trespass outside the triangle.

5.1 Interpretation of CMS Z -boson ‘Drell-Yan’ Results

The CMS experiment has published [50] high-precision angular distribution data on γ^*/Z dileptons around the Z boson peak, extending over the p_T range 0 – 300 GeV. They have performed a two-dimensional fit, and extracted the angular parameters in eight p_T -bins, for two rapidity ranges: $|y| \leq 1$ and $|y| \geq 1$.

5.1.1 Kinematic Derivation of Angular Distribution

The general expression for the angular distributions of the (parity conserving) Drell-Yan process in the Collins-Soper (CS) frame is:

$$\frac{d\sigma}{d\Omega} \propto 1 + \lambda \cos^2 \theta + \mu \sin 2\theta \cos \phi + \frac{\nu}{2} \sin^2 \theta \cos 2\phi + \rho \sin 2\theta \sin \phi + \sigma \sin^2 \theta \sin 2\phi \quad (5.1)$$




Figure 5.1: Relationship between the ‘natural axis’ (z') and the CS-frame

Due to the anti-symmetric dependence on ϕ , the $\sin \phi$ and $\sin 2\phi$ terms are experimentally inaccessible. Therefore, the experimentally observed angular distributions are completely described by the three parameters: λ , μ , and ν .

The CS frame is defined as the photon rest frame where the z-axis bisects the angle between the beam hadron momentum vector and the opposite of the target hadron momentum vector, and the x-axis lies in the plane defined by the hadron vectors.

Considering the $q\bar{q} \rightarrow \gamma^* \rightarrow l^+l^-$ subprocess for a single event, the angular distribution is:

$$\frac{d\sigma}{d\Omega} \propto 1 + \lambda_0 \cos^2 \theta_0 \quad (5.2)$$

where θ_0 is the angle between the quark and the positive lepton in the subprocess center-of-mass frame. The subprocess angular distribution is completely described by λ_0 . Furthermore, if the intermediate photon is completely transverse, then $\lambda_0 = 1$, and the angular distribution is simply $1 + \cos^2 \theta_0$ with flat azimuthal dependence.

Embedding the subprocess description within the CS frame, it is possible to derive an expression for the angular distributions in the Collins-Soper frame in terms of θ , ϕ , θ_1 , ϕ_1 , λ_0 , and θ_0 . θ_1 and ϕ_1 are the polar and azimuthal (respectively) angles of the quark momentum vector in the CS frame. This is simply a geometric manipulation since the intermediate photon is at rest in both the subprocess frame and the Collins-Soper frame. Figure 5.1 depicts the relationship between the angles. The quark momentum vectors are co-linear, and the ‘quark axis’ is labeled as \hat{z}' . The angle θ_1 characterizes the polar misalignment of the

quark axis and the CS z-axis. For non-zero θ_1 , the angle ϕ_1 quantifies the non-coplanarity of the quark axis and the hadron plane. In naive Drell-Yan, \hat{z}' coincides with \hat{z} , $\theta_1 = 0$, and ϕ_1 is undefined.

To relate the quark-axis angles to the CS frame angles, first notice that θ_0 obeys the following identity:

$$\cos \theta_0 = \cos \theta \cos \theta_1 + \sin \theta \sin \theta_1 \cos (\phi - \phi_1) \quad (5.3)$$

Plugging Eq. 5.3 into Eq. 5.2, we derive:

$$\begin{aligned} \frac{d\sigma}{d\Omega} \propto & 1 + \lambda_0 \left(\frac{2 - 3 \sin^2 \theta_1}{2 + \lambda_0 \sin^2 \theta_1} \right) \cos^2 \theta \\ & + \lambda_0 \left(\frac{\sin 2\theta_1 \cos \phi_1}{2 + \lambda_0 \sin^2 \theta_1} \right) \sin 2\theta \cos \phi \\ & + \lambda_0 \left(\frac{\sin^2 \theta_1 \cos 2\phi_1}{2 + \lambda_0 \sin^2 \theta_1} \right) \sin^2 \theta \cos 2\phi \\ & + \lambda_0 \left(\frac{\sin 2\theta_1 \sin \phi_1}{2 + \lambda_0 \sin^2 \theta_1} \right) \sin 2\theta \sin \phi \\ & + \lambda_0 \left(\frac{\sin^2 \theta_1 \sin 2\phi_1}{2 + \lambda_0 \sin^2 \theta_1} \right) \sin^2 \theta \sin 2\phi \end{aligned} \quad (5.4)$$

Comparing similar terms in Eq. 5.1 and Eq. 5.4 one can write down *single event* expressions for λ , μ , ν , ρ , and σ in terms of θ_1 , ϕ_1 , and λ_0 :

$$\begin{aligned} \lambda &= \lambda_0 \frac{2 - 3 \sin^2 \theta_1}{2 + \lambda_0 \sin^2 \theta_1} \\ \mu &= \lambda_0 \frac{\sin 2\theta_1 \cos \phi_1}{2 + \lambda_0 \sin^2 \theta_1} \\ \nu &= \lambda_0 \frac{2 \sin^2 \theta_1 \cos 2\phi_1}{2 + \lambda_0 \sin^2 \theta_1} \\ \rho &= \lambda_0 \frac{\sin 2\theta_1 \sin \phi_1}{2 + \lambda_0 \sin^2 \theta_1} \\ \sigma &= \lambda_0 \frac{\sin^2 \theta_1 \sin 2\phi_1}{2 + \lambda_0 \sin^2 \theta_1} \end{aligned} \quad (5.5)$$

Eqs. 5.5 are the most general forms of the kinematically-derived parameter definitions for photon-mediated Drell-Yan. (For Z -boson mediated ‘Drell-Yan’ there are additional terms due to parity violation[53]. This discussion will ignore those terms, as they do not alter the conclusions.)

5.1.2 Deriving Angular Parameter p_T -dependence

Before interpreting the CMS results using the kinematically-derived angular parameter definitions, we must derive expressions for the dependence of the angular parameters on p_T . Making use of derivations in the

published literature, in this section we will derive the $\mathcal{O}(\alpha_s)$ expressions for the p_T -dependence of λ and ν . We will derive expressions both for the $q\bar{q}$ and qG processes. (Note: For $p_T \gg k_T$, the $\mathcal{O}(\alpha_s^0)$ diagram does not contribute.)

Thews [51] derives the angular dependence of the λ and ν parameters in the Gottfried-Jackson frame for both the $q\bar{q}$ and qG processes.

For the $q\bar{q}$ diagrams:

$$\lambda_{q\bar{q}}^{GJ} = \frac{1 - w + w^2}{1 + 3w + w^2} \quad (5.6)$$

$$\nu_{q\bar{q}}^{GJ} = \frac{2w}{1 + 3w + w^2} \quad (5.7)$$

$$(5.8)$$

Where $w = \frac{p_T^2}{M^2}$, and p_T and M refer to the virtual photon transverse momentum and mass, respectively.

Lindfors [52] derives the conversion from the CS frame to the GJ frame for the A_0 angular parameter. (Recall that $A_0 = \frac{2-2\lambda}{\lambda+3}$)

$$A_0^{GJ} = A_0^{CS} + \sin^2 \beta^{CS} [1 - 2A_0^{CS}] + 2 \sin \beta^{CS} \cos \beta^{CS} A_1^{CS} \quad (5.9)$$

(NOTE! Lindfors [52] has a typo: $\sin \beta^{CS}$ should be $\sin^2 \beta^{CS}$)

Assuming $A_1^{CS} = 0$, $\sin^2 \beta^{CS} = \frac{w^2}{1+w^2}$ and rearranging...

$$A_0^{CS} = A_0^{GJ} \left(\frac{1 + w^2}{1 - w^2} \right) \quad (5.10)$$

Assuming Lam-Tung obedience ($A_0 = A_2$), we can convert the GJ-frame expressions for λ and ν into an expression for A_0 :

$${}^{q\bar{q}}A_0^{GJ} = \frac{2w}{1 + 2w + w^2} \quad (5.11)$$

Plugging in to Eq. 5.10, we derive:

$${}^{q\bar{q}}A_0^{CS} = \frac{w}{1 + w} \quad (5.12)$$

Finally, converting back to λ and ν :

$$\lambda_{q\bar{q}}^{CS} = \frac{2 - w}{2 + 3w} \quad (5.13)$$

$$\nu_{q\bar{q}}^{CS} = \frac{2w}{2 + 3w} \quad (5.14)$$

Equivalently, beginning with the Thews expressions for λ and ν for the quark-gluon process:

$$\lambda_{qG}^{GJ} = \frac{1 - 3w + 5w^2}{1 + 9w + 5w^2} \quad (5.15)$$

$$\nu_{qG}^{GJ} = \frac{6w}{1 + 9w + 5w^2} \quad (5.16)$$

Using Eq. 5.10, we find:

$${}^{qG}A_0^{CS} = \frac{5w}{1 + 5w} \quad (5.17)$$

Then we can derive the qG -process, CS-frame expressions for λ and ν :

$$\lambda_{qG}^{CS} = \frac{2 - 5w}{2 + 15w} \quad (5.18)$$

$$\nu_{qG}^{CS} = \frac{10w}{2 + 15w} \quad (5.19)$$

It is interesting to note that, for both the $q\bar{q}$ and qG processes, the derivation of the expression for A_0^{CS} involves canceling a common factor of $(1 - w)$ from the numerator and denominator. This means that *all* of the resulting angular parameter p_T -dependence expressions are undefined¹ at $w = 1$, or $p_T^2 = M^2$.

For both processes, $\lambda = 1$ and $\nu = 1$ at $p_T = 0$, and $\lambda \rightarrow \frac{-1}{3}$ and $\nu \rightarrow \frac{2}{3}$ for $p_T \rightarrow \text{inf}$. Also, both processes obey the Lam-Tung relation at all values of p_T . This is expected since the Lam-Tung relation is valid at $\mathcal{O}(\alpha_s)$. Despite these similarities, the two processes yield quite different shapes for the parameters in the intermediate p_T region, as seen in Fig. 5.2.

5.1.3 $q\bar{q}$ and qg Contributions and Quark Axis Non-Coplanarity from the CMS Data

The high-precision dilepton data from CMS allows for an enlightening comparison with the expected p_T -dependence of the angular parameters for the $q\bar{q}$ and qG processes.

The procedure proceeds as follows: First, use the λ data to extract the relative contribution of the $q\bar{q}$ and qG processes. Since λ only depends on θ_1 and not on ϕ_1 , the measured values of λ should be independent of any non-coplanarity. Using these relative process contributions, generate an expectation for the ν p_T -dependence, assuming $\cos 2\phi_1 = 1$ (coplanarity). Since any non-coplanarity can only reduce the measured value of ν (through the $\cos 2\phi_1$ factor), we can measure the average effect of the non-coplanarity by the deviation of the measured ν values below the expected p_T -dependence curve. In this interpretation, the non-coplanarity is *entirely* responsible for any observed Lam-Tung violation.

¹This should not affect the conclusion. I just thought it was interesting.




Figure 5.2: CS λ and ν vs p_T for qq and qg processes. Using $M = M_Z = 91.2$ GeV.

Fig. 5.3 shows the CMS data, plotted in terms of λ , ν , and $1 - \lambda - 2\nu$ as a function of p_T . Both the $|y| > 1$ and $|y| < 1$ rapidity data are shown, and they are quite consistent with each other. In the first panel, the $q\bar{q}$ and qG process expectations are shown in dashed curves. With the assumption that the fraction of the two processes is independent of p_T , a fit yields the solid curve, which corresponds to $58.8 \pm 1.6\%$ qG and $41.5 \pm 1.6\%$ $q\bar{q}$. The curve describes the data quite well throughout the p_T coverage of the data.

With the relative contribution of the two processes determined, we compare the expected curve for the weighted combination of the two processes against the CMS data on ν . Shown in the second panel of Fig. 5.3, the predicted curve (dotted) for ν is consistently, and significantly, slightly above the CMS data. Fitting with a simple scaling factor, we find a reduction of A_2 relative to the Lam-Tung prediction of $A_2/A_0 = 1$, caused by non-coplanarity. The best fit value is $A_2/A_0 = 0.77 \pm 0.02$. The solid curve shows the predicted shape for ν with $A_2/A_0 = 0.77$. The result is clearly in qualitative agreement with the data.

The last panel of Fig. 5.3 shows the quantity $1 - \lambda - 2\nu$, clearly demonstrating the Lam-Tung violation of the CMS data. The solid curve is the prediction, using $58.8 \pm 1.6\%$ qG and $41.5 \pm 1.6\%$ $q\bar{q}$, and $A_2/A_0 = 0.77$. The curve is in good agreement with the data, and reproduces the Lam-Tung violation of the data quite well.

In this framework, the angular distributions are generated by the misalignment of the quark axis relative to the Collins-Soper z -axis. In the naive subprocess, where the quark axis lies exactly along the Collins-Soper z -axis, the angular distribution of the leptons is simply given by $1 + \cos^2 \theta$ (assuming transversely polarized intermediate photons). At $\mathcal{O}(\alpha_s)$, the incoming quark or anti-quark can emit one gluon, and the “quark

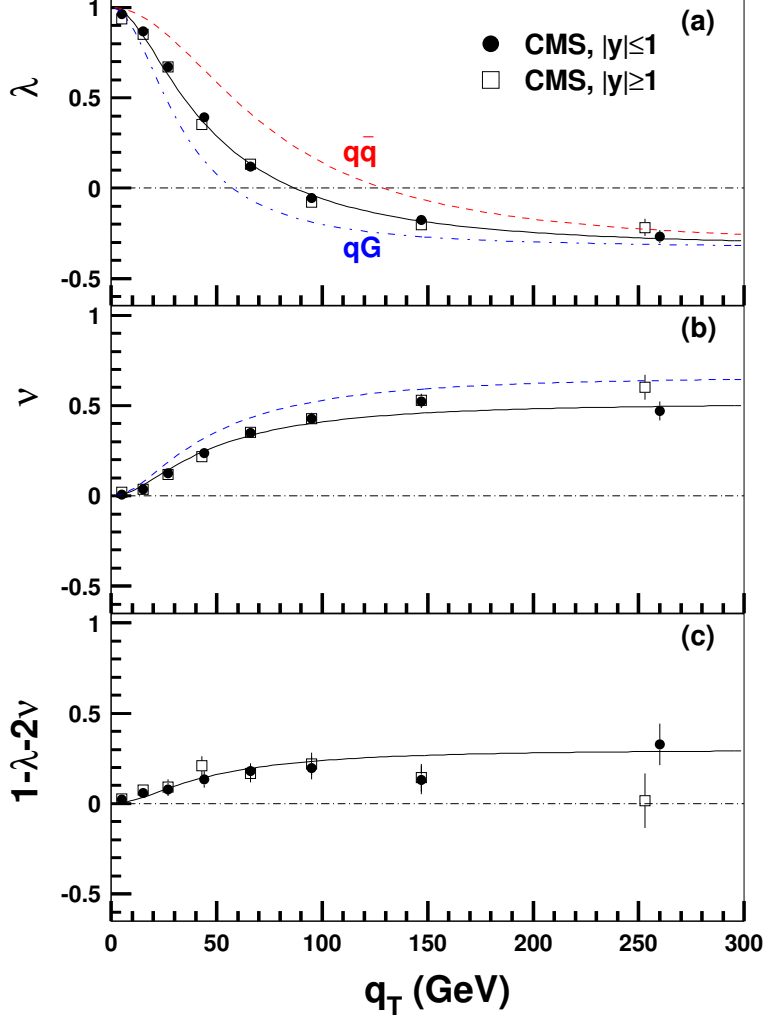


Figure 5.3: Comparison of the CMS data to the calculations described in Section 5.1.3. The figures demonstrate (a) extraction of the $q\bar{q}$ and qG contributions from the λ p_T -dependence, (b) extraction of the average degree of non-coplanarity from the ν p_T -dependence, and (c) comparison of the Lam-Tung violation between the data and calculations. From Ref. [53]

bremsstrahlung” process becomes relevant.

The $q\bar{q}$ -annihilation diagrams at $\mathcal{O}(\alpha_s)$ are shown in Fig. 5.4. For both of these diagrams, the quark axis must always remain in the hadron plane, assuming negligible intrinsic parton transverse momentum. Furthermore, the quark axis is always aligned with either the beam hadron momentum or the target hadron momentum. One of the partons emits a hard gluon, but the other must remain aligned with its parent hadron. Therefore, $\cos 2\phi_1 = 1$ and $\theta_1 = \beta$. Recall that, assuming transversely polarized intermediate photons, Lam-Tung violation can only occur due to non-coplanarity of the quark axis and the hadron plane. So, at $\mathcal{O}(\alpha_s)$, $q\bar{q}$ -annihilation will obey the Lam-Tung relation. Lam-Tung violation due to non-coplanarity can only occur in higher orders of α_s .

The qG diagrams at $\mathcal{O}(\alpha_s)$ are shown in Fig. 5.5. In the second diagram, the quark axis again must lie along either the beam hadron momentum or the target hadron momentum, since the incoming quark emits the photon before absorbing the gluon. In the first diagram, the quark axis is not restricted to the hadron momentum directions, but it still must remain coplanar with the hadron plane. The momentum of the intermediate, off-shell quark that emits the photon is the vector sum of the momenta of the two participating partons (one quark and one gluon). Under the assumption of negligible intrinsic parton transverse momentum, both of these partons are aligned with their parent hadrons. Therefore, the intermediate quark remains within the hadron plane and points somewhere between the two hadron momenta. Therefore, the Lam-Tung relation is still expected to hold for all contributing diagrams at $\mathcal{O}(\alpha_s)$ for the qG -process. Any Lam-Tung violation due to non-coplanarity must again come from higher-order diagrams. It is also interesting to note that, in the photon rest frame, the quark momenta before and after the photon emission are equal.

Therefore, the result of $A_2/A_0 = 0.77 \pm 0.02$ in the CMS data indicates a significant contribution of higher-order processes in pp dilepton production at the LHC.




Figure 5.4: Diagrams contributing to the annihilation process at $\mathcal{O}(\alpha_s)$




Figure 5.5: Diagrams contributing to the quark bremsstrahlung process at $\mathcal{O}(\alpha_s)$

5.2 Derivation of Parameter Relations

Within the kinematic picture presented above, it is interesting to vary the assumptions on quark co-planarity and photon polarization and derive relations between the angular parameters. In addition to re-deriving the Lam-Tung relation, more general equalities can be derived by allowing non-coplanar quarks and non-transverse photon polarization.

5.2.1 Deriving the Lam-Tung Relation

To derive the Lam-Tung relation from the kinematic description, one first must apply two limiting assumptions:

$$\begin{aligned}\lambda_0 &= 1 \\ \cos 2\phi_1 &= 0\end{aligned}$$

These two assumptions specify that the intermediate photon must be completely transversely polarized, and that the quark-axis must lie within the hadron plane, respectively. Both of these assumptions were implicitly made by Lam and Tung in their original derivation [23]. The *single event* parameter definitions then become:

$$\begin{aligned}\lambda &= \frac{2 - 3 \sin^2 \theta_1}{2 + \sin^2 \theta_1} \\ \mu &= \frac{\sin 2\theta_1}{2 + \sin^2 \theta_1} \\ \nu &= \frac{2 \sin^2 \theta_1}{2 + \sin^2 \theta_1}\end{aligned}\tag{5.20}$$

Combining the definitions for λ and ν , it is possible to derive the Lam-Tung relation:

$$\lambda = 1 - 2\nu\tag{5.21}$$

Therefore, the Lam-Tung relation is simply a result of the geometrical relationship between the natural axis and the experimentally observable angles in the photon rest frame, for Drell-Yan under the assumption of co-planar quarks. As discussed further in Section 5.1.3, diagrams up through $\mathcal{O}(\alpha_s)$ yield a co-planar quark axis (ignoring intrinsic k_T). Non-coplanarity can only result from intrinsic k_T and/or higher-order diagrams. So, in high-energy experiments where $p_T \gg k_T$ and the photon p_T is dominantly generated by $\mathcal{O}(\alpha_s)$ diagrams, the assumption of co-planar quarks should be approximately correct. Therefore, the Lam-Tung relation should be largely satisfied.

Note that the *single-event* parameters are not experimentally accessible. Experiments must extract $\langle\lambda\rangle$, $\langle\mu\rangle$, and $\langle\nu\rangle$ by fitting to a large collection of events. Fortunately, $\langle 1 \rangle = 1 = \langle \lambda + 2\nu \rangle = \langle \lambda \rangle + 2\langle \nu \rangle$, so the Lam-Tung relation is also valid for data-averaged parameter values.

Under the same assumptions, it is possible to derive an expression involving the μ parameter:

$$\mu^2 = \frac{1}{4}(1 - \lambda)(1 + \lambda - \nu) \quad (5.22)$$

For individual events, this equation is just as valid as the Lam-Tung relation. However, its non-linearity makes it impossible to derive the corresponding expression for the data-averaged parameters. Therefore, it is impossible to directly compare this relation to experimental data. The details of this limitation will be explored in Section 5.3.

5.2.2 Non-co-planar Quark Axis

Next, consider freeing the quark axis from the hadron plane. Phenomenologically, this could occur due to two gluon emission, or significant k_T of both annihilating quarks. Following the same kinematic prescription, and assuming $\lambda_0 = 1$, but keeping ϕ_1 free, we can derive more general definitions of the parameters:

$$\begin{aligned} \lambda &= \frac{2 - 3 \sin^2 \theta_1}{2 + \sin^2 \theta_1} \\ \mu &= \frac{\sin 2\theta_1 \cos \phi_1}{2 + \sin^2 \theta_1} \\ \nu &= \frac{2 \sin^2 \theta_1 \cos 2\phi_1}{2 + \sin^2 \theta_1} \end{aligned} \quad (5.23)$$

Clearly, since ν has ϕ_1 dependence but λ does not, it is impossible to derive an exact analogue of the Lam-Tung relation. However, μ also depends on ϕ_1 , and it is possible to derive an equation using all three parameters:

$$\mu^2 = \frac{1}{16}(1 + 3\lambda)(1 - \lambda + 2\nu) \quad (5.24)$$

For individual events, this equation should hold for all quark-axis orientations. Again, this equation is non-linear and cannot be directly compared to data.

5.2.3 Non-Transverse Photon

Next, assume that the quarks are co-planar with the hadrons, but that λ_0 can vary. The parameter definitions are:

$$\begin{aligned}\lambda &= \lambda_0 \frac{2 - 3 \sin^2 \theta_1}{2 + \lambda_0 \sin^2 \theta_1} \\ \mu &= \lambda_0 \frac{\sin 2\theta_1}{2 + \lambda_0 \sin^2 \theta_1} \\ \nu &= \lambda_0 \frac{2 \sin^2 \theta_1}{2 + \lambda_0 \sin^2 \theta_1}\end{aligned}\tag{5.25}$$

Combining the expression for λ and ν , we can derive a Lam-Tung-like relation for generalized photon polarization:

$$\lambda = \lambda_0 - \frac{\nu}{2}(\lambda_0 + 3)\tag{5.26}$$

Clearly this reduces to the Lam-Tung relation for $\lambda_0 = 1$.

Additionally, we may generalize Eq. 5.22 for unspecified λ_0 (written three different ways):

$$\mu^2 = \frac{2\nu\lambda_0^2(2\lambda + \nu)}{(2\lambda + 3\nu + \nu\lambda_0)^2}\tag{5.27}$$

$$\mu^2 = \frac{2(\lambda_0 - \lambda)(\lambda\lambda_0 + \lambda_0 + 2\lambda)}{(3 + \lambda_0)^2}\tag{5.28}$$

$$\lambda_0 = \frac{\mu^2 + \nu^2}{\nu - \frac{\nu^2}{2}}\tag{5.29}$$

5.2.4 Fully General

Finally, releasing both constraints required for the Lam-Tung derivation, we can derive an expression involving all three parameters, plus λ_0 , using Eq. 5.5:

$$\mu^2 = \left(\frac{1}{\lambda_0 + 3}\right)^2 \left(\frac{\nu}{2}(\lambda_0 + 3) + \lambda_0 - \lambda\right)(\lambda_0 + \lambda\lambda_0 + 2\lambda)\tag{5.30}$$

5.3 Comparison with Data

In this section, the derived *single-event* parameter definitions are used to derive constraints on the data-averaged parameters accessible by experiment. Then, these constraints are compared against results from four experiments: CMS [50], E866 [32, 33], E615 [30, 31], and NA10 [28, 29].

5.3.1 Constraints on Parameters

Setup

Begin with the usual kinematic derivation of the angular distribution. For the $q\bar{q} \rightarrow l^+l^-$ subprocess, assuming an intermediate transversely-polarized γ^* :

$$\frac{d\sigma}{d\Omega} \propto 1 + \cos^2 \theta \quad (5.31)$$

Transforming from the natural frame to the CS frame, placing no restrictions on the quark axis angles, θ_1 and ϕ_1 , we get:

$$\begin{aligned} \frac{d\sigma}{d\Omega} \propto & (1 + \cos^2 \theta) + \frac{\sin^2 \theta_1}{2}(1 - 3 \cos^2 \theta) \\ & + \left(\frac{1}{2} \sin 2\theta_1 \cos \phi_1\right) \sin 2\theta \cos \phi \\ & + \left(\frac{1}{2} \sin^2 \theta_1 \cos 2\phi_1\right) \sin^2 \theta \cos 2\phi \\ & + \left(\frac{1}{2} \sin^2 \theta_1 \sin 2\phi_1\right) \sin^2 \theta \sin 2\phi \\ & + \left(\frac{1}{2} \sin 2\theta_1 \sin \phi_1\right) \sin 2\theta \sin \phi \end{aligned} \quad (5.32)$$

Therefore, we can write down the *single event* angular parameter definitions. For comparison with experimental data, an average over many events must be applied. For ease of calculations, I will use the collider angular parameter convention.

$$\begin{aligned} A_0 &= \sin^2 \theta_1 \\ A_1 &= \frac{1}{2} \sin 2\theta_1 \cos \phi_1 \\ A_2 &= \sin^2 \theta_1 \cos 2\phi_1 \\ A_5 &= \frac{1}{2} \sin^2 \theta_1 \sin 2\phi_1 \\ A_6 &= \frac{1}{2} \sin 2\theta_1 \sin \phi_1 \end{aligned} \quad (5.33)$$

A_5 and A_6 cannot be measured experimentally due to their antisymmetric dependence on ϕ_1 .

Note that the fixed-target parameters are related to the collider parameters (assuming $\lambda_0 = 1$) by:

$$\lambda = \frac{2 - 3A_0}{2 + A_0}; \quad \mu = \frac{2A_1}{2 + A_0}; \quad \nu = \frac{2A_2}{2 + A_0} \quad (5.34)$$

Flat Bounds on A_0 and A_2

Using the *single event* definitions for A_0 and A_2 , and using $0 \leq \sin^2 \theta_1 \leq 1$ and $-1 \leq \cos 2\phi_1 \leq 1$, it is clear that:

$$0 \leq A_0 \leq 1 \quad \text{and} \quad -1 \leq A_2 \leq 1 \quad (5.35)$$

So, further discussion on bounds of A_0 and A_2 can be restricted to the rectangle defined by these bounds. These flat bounds are depicted in a 2-dimensional plot of A_2 -vs- A_0 in Figure 5.6.




Figure 5.6: Flat bounds on A_0 and A_2 .

Combined Bounds on A_0 and A_2

Combining the *single event* definitions for A_0 and A_2 , we can derive:

$$A_2 = A_0 \cos 2\phi_1 \quad (5.36)$$

Recalling that $-1 \leq \cos 2\phi_1 \leq 1$, we derive the following inequalities:

$$A_2 \leq A_0 \quad \text{and} \quad A_2 \geq -A_0 \quad (5.37)$$

Or, more concisely:

$$|A_2| \leq A_0 \quad (5.38)$$

Note that these bounds lie along the line defined by the Lam-Tung relation, and its reflection across $A_2 = 0$. Therefore, even if the Lam-Tung relation is violated due to non-coplanarity, the “Lam-Tung Inequality” $A_2 \leq A_0$ must still hold. The “allowed triangle”, defined by the bounds discussed so far, is depicted in Figure 5.7.

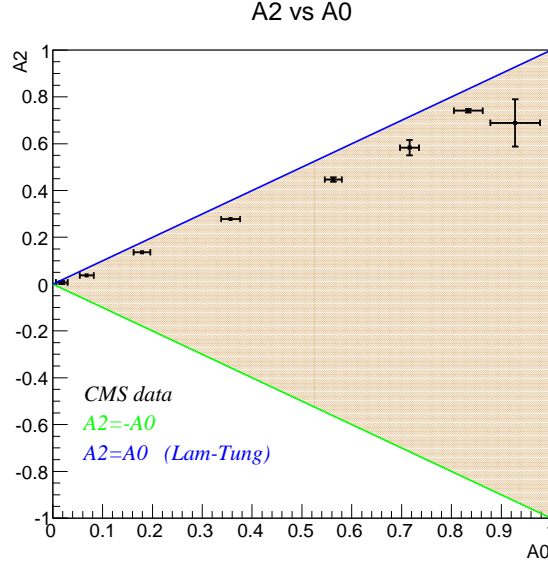


Figure 5.7: Bounds on A_0 and A_2 from the combined single-event parameter definitions.

Equality Relation

Using the *single event* definitions for A_0 , A_1 , and A_2 , we can derive an equality relation that should always be obeyed:

$$A_1^2 = \frac{1}{2}(1 - A_0)(A_0 + A_2) \quad (5.39)$$

This is equivalent to Equation 5.24. Notice that, in general, the right-hand side of the equality relation could be negative. However, this would imply that A_1 is imaginary, which is unphysical. Therefore, requiring $A_1 \in \mathbf{R}$, we obtain the following bounds:

$$A_0 \leq 1 \quad \wedge \quad A_2 \geq -A_0 \quad (5.40)$$

or

$$A_0 \geq 1 \quad \wedge \quad A_2 \leq -A_0 \quad (5.41)$$

However, since $A_0 \leq 1$, the region defined by the second set of bounds is excluded. These bounds correspond exactly to bounds already derived in the previous two sections.

5.3.2 Experimental Measurements

Since the angular distribution parameters can only be extracted from a large event sample, not for each event individually, we must consider the effect of averaging over a large number of events. Since A_0 and A_2 must




Figure 5.8: Kinematically derived bounds on A_0 and A_2 , defining the “allowed triangle”.

lie within the shaded triangle for each event, and since the triangle is a closed, convex area, the averaged values, $\langle A_0 \rangle$ and $\langle A_2 \rangle$, over *any* event sample must also lie within it. **Therefore, *any* experimental measurement of A_0 and A_2 must lie within the shaded triangle.** (With the assumptions of $\lambda_0 = 1$ and no effects from TMDs).

$$\begin{aligned}
 \langle A_0 \rangle &\leq 1 \\
 \langle A_2 \rangle &\leq \langle A_0 \rangle \\
 \langle A_2 \rangle &\geq -\langle A_0 \rangle
 \end{aligned}
 \tag{5.42}$$

Translating these constraints back into the λ, μ, ν convention:

$$\begin{aligned} \langle \lambda \rangle &\geq \frac{-1}{3} \\ \langle \nu \rangle &\leq \frac{1 - \langle \lambda \rangle}{2} \\ \langle \nu \rangle &\geq \frac{\langle \lambda \rangle - 1}{2} \end{aligned} \tag{5.43}$$

Now, we compare this allowed triangle with published data on λ and ν in Drell-Yan experiments.




Figure 5.9: CMS Data from Ref. [50]




Figure 5.10: FNAL E866 Data from Ref. [32, 33]




Figure 5.11: FNAL E615 Data from Ref. [30, 31]




Figure 5.12: NA10 Data from Ref. [28, 29]

The high-precision CMS data plotted on the allowed triangle in Figure 5.9 are consistent with the kinematically-derived constraints. All data points lie within the triangle. Recall that the data points represent different bins in p_T , from 0 GeV up to about 250 GeV, with the small- p_T points having larger λ values and the large- p_T points having smaller and negative λ values. The data points lie nearly along the line of the Lam-Tung relation. Therefore, the Lam-Tung violation in this data is relatively minor (compared with the available region within the allowed triangle). The distance from the Lam-Tung line increases with p_T . Most interestingly, the p_T -binned data points lie along nearly the entire length of the ‘‘Lam-Tung side’’ of the allowed triangle. Phrased another way, the low- p_T data appear to ‘saturate’ the high- λ point of the triangle, and the high- p_T data nearly saturate the high- ν point of the triangle. Since the precision of the data decrease for the high- p_T points, it is difficult to extrapolate the trend as $p_T \rightarrow \text{inf}$. With the more detailed interpretation of the CMS data in Section 5.1, the low- p_T extrapolation would predict to $\nu = 0.77\nu_{\text{max}}$ as $\lambda \rightarrow \frac{-1}{3}$.

The data from FNAL E866, plotted in Figure 5.10, are lower-precision and cover a much smaller p_T region (0 – 4 GeV). They are less consistent with the derived constraints. The E866 data appear to oversaturate the high- λ corner of the triangle. The uncertainties in λ are quite large, so it is difficult to quantify the apparent violation of the allowed triangle. The data from pp and pd interactions are mostly consistent with each other, considering the poor precision.

The angular parameter results reported by the E615 and NA10 collaborations, plotted in Figures 5.11 and 5.12, respectively, present a bit of a puzzle. When these data were reported, they were plotted only as λ , ν , and $2\nu - (1 - \lambda)$, separately. The results were interpreted as demonstrating significant a Lam-Tung violation, increasing with p_T . However, when these data are plotted on the ν vs λ plot, they quite clearly violate the derived constraint triangle. All Drell-Yan dileptons, generated by unpolarized quarks in unpolarized hadrons, and mediated by transverse photons *must* lie within the constraint triangle. Therefore, either the data are too poor, or some other process is dominating the dilepton production in these experiments.

One very interesting interpretation is that the observed violation is caused by non-zero valence quark TMDs. In particular, a non-zero Boer-Mulders function of the valence quarks in the pion and nucleon would cause a transverse momentum dependent polarization of the quarks within the unpolarized hadrons. If the proton and pion valence Boer-Mulders functions are the *same-sign*, they are expected [27] to result in a rise of ν with p_T , without altering the polar distribution away from $\lambda \sim 1$. Both the NA10 and E615 data are qualitatively consistent with this expectation.

Chapter 6

Conclusions

In this dissertation, I have presented my tripartite contribution to the study of hadron physics. On the experimental design and operation front, my efforts improving and operating the SeaQuest trigger system have been essential for the collection of analyzable data. On the data analysis front, I have blazed a trail for the challenging analysis of the Drell-Yan angular distributions with SeaQuest data. On the theoretical interpretation front, I have contributed to the development of a kinematical framework for deriving features of the Drell-Yan angular distributions. I have presented the methods and results of my contributions to these three fronts in the larger context of the coordinated efforts of my collaborators.

The FPGA-based SeaQuest trigger system is a significant technological improvement over the trigger systems used in previous fixed-target Drell-Yan experiments at FNAL. By implementing the trigger logic in FPGA firmware, the “trigger matrix” is easily adapted to changing experimental conditions. Furthermore, the timing alignment, “trigger matrix” generation, and firmware uploading procedures have been mostly automated, reducing the risk of user error in maintaining and operating the trigger system.

As should be expected when implementing a novel, complicated subsystem, many issues were discovered while operating the trigger system in real experimental conditions. The addition of the Level 0 v1495 boards accommodated the measurement of triggers lost due to DAQ deadtime. More importantly, Level 0 enabled rigorous pulser-testing of the Level 1 and Level 2 firmware, uncovering numerous previously unknown firmware bugs. Lastly, the Level 0 *in-situ* pulser test is used to check each newly-compiled firmware version (every time the roads are changed), to diagnose occasional firmware problems that generate no compile-time errors.

Due to the (late) realization that the internally-clocked TDC blocks could lose hits due to hits in adjacent buckets, the firmware for all Levels was modified to clock all firmware operations based on the MI-generated RF-Clock. This RF-Clock synchronization eliminated a possible source of rate-dependent trigger efficiency for which calculating a correction would pose a serious challenge. After this update, the FPGA-based trigger system was finally dead-time free.

The analysis of SeaQuest data has been quite challenging for analyzers of all physics topics. The principle

difficulty thus far has been the rate-dependence. By analyzing the Random RF-embedded Monte Carlo data, I found that the rate dependence is dominated by a rate-dependent inefficiency in the event reconstruction. Furthermore, I demonstrated that this rate-dependent inefficiency is also kinematically-dependent. I internally published the embedded data results in a simplified form, so that each analyzer can easily calculate and apply a kinematically-dependent, rate-dependent reconstruction efficiency to their data. This correction is absolutely essential for all analyses of SeaQuest data.

The large obstacle for SeaQuest analyses is a full characterization of the combinatoric background. Through much effort from many SeaQuest collaborators, significant progress has been made in the understanding of the combinatoric background present in the SeaQuest data. However, the full characterization will require calculating and applying a reconstruction rate-dependence correction to the data-derived estimation of the combinatoric background. Only then can the combinatoric background contribution be accurately subtracted from the rate-dependence corrected data.

I have made significant progress in the development of a procedure for the especially challenging analysis of Drell-Yan angular distributions with SeaQuest data. The extraction of the underlying angular distributions is highly sensitive to the spectrometer acceptance. I have developed a simple, purely Monte Carlo based correction for the spectrometer acceptance. Since the reconstruction rate-dependence is also dependent on the photon rest frame angles, the angular distributions analysis also absolutely requires the rate-dependence correction. The current results, while consistent with expectations, have large statistical uncertainties and show signs of significant uncorrected systematics. I have demonstrated that the combinatoric background conspires to present a special challenge to the angular distributions analysis.

Finally, I have extended the purely-kinematic derivation of the Drell-Yan angular parameter definitions developed by Oleg Teryaev and Jen-Chieh Peng. With the kinematic picture, they have developed derivations for the Lam-Tung relation and an equation for μ^2 in the case of coplanar quarks (as one expects for $\mathcal{O}(\alpha_s)$ Drell-Yan). They have also derived the associated inequalities for the case of non-coplanar quarks. While working through these derivations, I found an *equality* relating λ , μ , and ν , which holds even in the *non-coplanar* case. The discovery of Equation 5.24, reproduced below, is one of my proudest achievements.

$$\mu^2 = \frac{1}{16}(1 + 3\lambda)(1 - \lambda + 2\nu)$$

In addition, I contributed [53] to the application of the kinematic picture to the interpretation of high-precision Z-boson “Drell-Yan” data published by CMS [50]. The kinematic picture, in combination with the expected p_T dependence of the $q\bar{q}$ - and qg -processes already published in the literature, is one of the

first extractions of the relative contributions of the two $\mathcal{O}(\alpha_s)$ subprocesses to the Drell-Yan cross-section using the angular distributions. Furthermore, with this kinematic picture, we can naturally explain the observed Lam-Tung violation with non-coplanarity due to contribution from higher-order diagrams. This interpretation of the CMS (and CDF) data has been submitted for publication.

References

- [1] D. J. Gross and F. Wilczek, “Ultraviolet Behavior of Nonabelian Gauge Theories,” *Phys. Rev. Lett.*, vol. 30, pp. 1343–1346, 1973.
- [2] H. D. Politzer, “Reliable Perturbative Results for Strong Interactions?,” *Phys. Rev. Lett.*, vol. 30, pp. 1346–1349, 1973.
- [3] T.-M. Yan, “Naive Drell-Yan and its successor,” in *Sid Drell Symposium Stanford, California, July 31, 1998*, 1998.
- [4] J. D. Bjorken, “Asymptotic Sum Rules at Infinite Momentum,” *Phys. Rev.*, vol. 179, pp. 1547–1553, 1969.
- [5] R. P. Feynman, “Very high-energy collisions of hadrons,” *Phys. Rev. Lett.*, vol. 23, pp. 1415–1417, 1969.
- [6] M. Gell-Mann, “A Schematic Model of Baryons and Mesons,” *Phys. Lett.*, vol. 8, pp. 214–215, 1964.
- [7] G. Zweig, “An SU(3) model for strong interaction symmetry and its breaking. Version 1,” 1964.
- [8] J. Joyce, *Finnegans Wake*. Viking Press, 1939.
- [9] J. E. Augustin *et al.*, “Discovery of a Narrow Resonance in e^+e^- Annihilation,” *Phys. Rev. Lett.*, vol. 33, pp. 1406–1408, 1974. [Adv. Exp. Phys.5,141(1976)].
- [10] J. J. Aubert *et al.*, “Experimental Observation of a Heavy Particle J,” *Phys. Rev. Lett.*, vol. 33, pp. 1404–1406, 1974.
- [11] F. Abe *et al.*, “Observation of top quark production in $\bar{p}p$ collisions,” *Phys. Rev. Lett.*, vol. 74, pp. 2626–2631, 1995.
- [12] S. Abachi *et al.*, “Search for high mass top quark production in $p\bar{p}$ collisions at $\sqrt{s} = 1.8$ TeV,” *Phys. Rev. Lett.*, vol. 74, pp. 2422–2426, 1995.
- [13] S. W. Herb *et al.*, “Observation of a Dimuon Resonance at 9.5-GeV in 400-GeV Proton-Nucleus Collisions,” *Phys. Rev. Lett.*, vol. 39, pp. 252–255, 1977.
- [14] C. Quigg, “Elementary Particles and Forces,” *Sci. Am.*, vol. 252, 1985.
- [15] R. P. Feynman, “Space - time approach to quantum electrodynamics,” *Phys. Rev.*, vol. 76, pp. 769–789, 1949.
- [16] K. Gottfried, “Sum rule for high-energy electron - proton scattering,” *Phys. Rev. Lett.*, vol. 18, p. 1174, 1967.
- [17] S. D. Drell and T.-M. Yan, “Massive Lepton Pair Production in Hadron-Hadron Collisions at High-Energies,” *Phys. Rev. Lett.*, vol. 25, pp. 316–320, 1970. [Erratum: *Phys. Rev. Lett.*25,902(1970)].
- [18] P. Amaudruz *et al.*, “Proton and deuteron f_2 structure functions in deep inelastic muon scattering,” *Phys. Lett.*, vol. B295, pp. 159–168, 1992.

- [19] A. Baldit *et al.*, “Study of the isospin symmetry breaking in the light quark sea of the nucleon from the Drell-Yan process,” *Phys. Lett.*, vol. B332, pp. 244–250, 1994.
- [20] R. S. Towell *et al.*, “Improved measurement of the anti-d / anti-u asymmetry in the nucleon sea,” *Phys. Rev.*, vol. D64, p. 052002, 2001.
- [21] L. D. Isenhower *et al.*, “Proposal for Drell-Yan Measurements of Nucleon and Nuclear Structure with the FNAL Main Injector,” 2001.
- [22] J. Badier *et al.*, “Angular Distributions in the Dimuon Hadronic Production at 150-GeV/c,” *Z. Phys.*, vol. C11, p. 195, 1981.
- [23] C. S. Lam and W.-K. Tung, “A Systematic Approach to Inclusive Lepton Pair Production in Hadronic Collisions,” *Phys. Rev.*, vol. D18, p. 2447, 1978.
- [24] C. S. Lam and W.-K. Tung, “A Parton Model Relation Sans QCD Modifications in Lepton Pair Productions,” *Phys. Rev.*, vol. D21, p. 2712, 1980.
- [25] C. G. Callan, Jr. and D. J. Gross, “High-energy electroproduction and the constitution of the electric current,” *Phys. Rev. Lett.*, vol. 22, pp. 156–159, 1969.
- [26] M. G. Perdekamp and F. Yuan, “Transverse Spin Structure of the Nucleon,” *Ann. Rev. Nucl. Part. Sci.*, vol. 65, pp. 429–456, 2015.
- [27] D. Boer, “Investigating the origins of transverse spin asymmetries at RHIC,” *Phys. Rev.*, vol. D60, p. 014012, 1999.
- [28] S. Falciano *et al.*, “Angular Distributions of Muon Pairs Produced by 194-GeV/c Negative Pions,” *Z. Phys.*, vol. C31, p. 513, 1986.
- [29] M. Guanziroli *et al.*, “Angular Distributions of Muon Pairs Produced by Negative Pions on Deuterium and Tungsten,” *Z. Phys.*, vol. C37, p. 545, 1988.
- [30] J. S. Conway *et al.*, “Experimental Study of Muon Pairs Produced by 252-GeV Pions on Tungsten,” *Phys. Rev.*, vol. D39, pp. 92–122, 1989.
- [31] J. G. Heinrich *et al.*, “Higher twist effects in the reaction $\pi^- N \rightarrow \mu^+ \mu^- X$ at 253-GeV/c,” *Phys. Rev.*, vol. D44, pp. 1909–1932, 1991.
- [32] L. Y. Zhu *et al.*, “Measurement of Angular Distributions of Drell-Yan Dimuons in $p + d$ Interaction at 800-GeV/c,” *Phys. Rev. Lett.*, vol. 99, p. 082301, 2007.
- [33] L. Y. Zhu *et al.*, “Measurement of Angular Distributions of Drell-Yan Dimuons in $p + p$ Interactions at 800-GeV/c,” *Phys. Rev. Lett.*, vol. 102, p. 182001, 2009.
- [34] J.-C. Peng and J.-W. Qiu, “Novel phenomenology of parton distributions from the Drell-Yan process,” *Prog. Part. Nucl. Phys.*, vol. 76, pp. 43–75, 2014.
- [35] V. Barone, S. Melis, and A. Prokudin, “The Boer-Mulders effect in unpolarized SIDIS: An Analysis of the COMPASS and HERMES data on the $\cos 2\phi$ asymmetry,” *Phys. Rev.*, vol. D81, p. 114026, 2010.
- [36] J. Ashman *et al.*, “A Measurement of the Spin Asymmetry and Determination of the Structure Function $g(1)$ in Deep Inelastic Muon-Proton Scattering,” *Phys. Lett.*, vol. B206, p. 364, 1988.
- [37] A. Bacchetta, M. Diehl, K. Goeke, A. Metz, P. J. Mulders, and M. Schlegel, “Semi-inclusive deep inelastic scattering at small transverse momentum,” *JHEP*, vol. 02, p. 093, 2007.
- [38] J. C. Collins, “Leading twist single transverse-spin asymmetries: Drell-Yan and deep inelastic scattering,” *Phys. Lett.*, vol. B536, pp. 43–48, 2002.

- [39] M. Diefenthaler, “Azimuthal single-spin asymmetries in semi-inclusive deep-inelastic scattering on a transversely polarised hydrogen target,” *AIP Conf. Proc.*, vol. 915, pp. 509–512, 2007. [509(2006)].
- [40] L. L. Pappalardo, “Latest HERMES results on transverse spin effects in hadron structure and formation,” *Eur. Phys. J.*, vol. A38, pp. 145–148, 2008.
- [41] T. H. Chang *et al.*, “ J/ψ polarization in 800-GeV p Cu interactions,” *Phys. Rev. Lett.*, vol. 91, p. 211801, 2003.
- [42] C. N. Brown *et al.*, “Observation of polarization in bottomonium production at $\sqrt{s} = 38.8$ -GeV,” *Phys. Rev. Lett.*, vol. 86, pp. 2529–2532, 2001.
- [43] D. M. Alde *et al.*, “Nuclear dependence of dimuon production at 800-GeV. FNAL-772 experiment,” *Phys. Rev. Lett.*, vol. 64, pp. 2479–2482, 1990.
- [44] D. M. Alde *et al.*, “The A-dependence of J/ψ and ψ' production at 800-GeV/c,” *Phys. Rev. Lett.*, vol. 66, pp. 133–136, 1991.
- [45] J. C. Webb *et al.*, “Absolute Drell-Yan dimuon cross-sections in 800 GeV / c pp and pd collisions,” 2003.
- [46] A. Einstein, “Concerning an heuristic point of view toward the emission and transformation of light,” *Annalen Phys.*, vol. 17, pp. 132–148, 1905.
- [47] S.-H. Shiu, J. Wu, R. E. McClellan, T.-H. Chang, W.-C. Chang, Y.-C. Chen, R. Gilman, K. Nakano, J.-C. Peng, and S.-Y. Wang, “FPGA-based trigger system for the Fermilab SeaQuest experiment,” *Nucl. Instrum. Meth.*, vol. A802, pp. 82–88, 2015.
- [48] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Transactions of the ASME—Journal of Basic Engineering*, vol. 82, no. Series D, pp. 35–45, 1960.
- [49] S. Gorbunov *et al.*, “ALICE HLT high speed tracking and vertexing,” in *Proceedings, 17th Conference on Computing Applications in Nuclear and Plasma Sciences on Real Time (RT2010)*, 2010.
- [50] V. Khachatryan *et al.*, “Angular coefficients of Z bosons produced in pp collisions at $\sqrt{s} = 8$ TeV and decaying to $\mu^+\mu^-$ as a function of transverse momentum and rapidity,” *Phys. Lett.*, vol. B750, pp. 154–175, 2015.
- [51] R. L. Thews, “Tests for Two-body QCD Subprocesses in Dilepton Angular Correlations,” *Phys. Rev. Lett.*, vol. 43, p. 987, 1979. [Erratum: *Phys. Rev. Lett.*43,1968(1979)].
- [52] J. Lindfors, “Angular Distribution of Large q -transverse Muon Pairs in Different Reference Frames,” *Phys. Scripta*, vol. 20, pp. 19–22, 1979.
- [53] J.-C. Peng, W.-C. Chang, R. E. McClellan, and O. Teryaev, “Interpretation of Angular Distributions of Z-boson Production at Colliders,” 2015.