

© 2016 by Bryan P. Dannowitz. All rights reserved.

NUCLEAR DEPENDENCE OF PROTON-INDUCED DRELL-YAN DIMUON
PRODUCTION AT 120 GEV AT SEAQUEST

BY

BRYAN P. DANNOWITZ

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 Jen-Chieh Peng, Chair
Professor Naomi C. R. Makins, Director of Research
Professor Michael Stone
Professor Jim Eckstein

Abstract

A measurement of the atomic mass (A) dependence of $p + A \rightarrow \mu^+\mu^- + X$ Drell-Yan dimuons produced by 120 GeV protons is presented here. The data was taken by the SeaQuest experiment at Fermilab using a proton beam extracted from its Main Injector. Over 61,000 dimuon pairs were recorded with invariant mass $4.2 < M_{\gamma^*} < 10 \text{ GeV}$ and target parton momentum fraction $0.1 \leq x_2 \leq 0.5$ for nuclear targets 1H , 2H , C , Fe , and W . The ratio of dimuon yields per nucleon (Y) for heavy nuclei versus 2H , $R_{DY} = Y^{(A)}/Y^{(^2H)} \approx \bar{u}^{(A)}(x)/\bar{u}^{(^2H)}(x)$, is sensitive to modifications in the anti-quark sea distributions in nuclei for the case of proton-induced Drell-Yan. The data analyzed here and in the future of SeaQuest will provide tighter constraints on various models that attempt to define the anomalous behavior of nuclear modification as seen in deep inelastic lepton scattering, a phenomenon generally known as the EMC effect.

Dedicated to my grandfather, Ted, who took me to school every day.

Acknowledgments

This dissertation has manifested itself as the grand sum of the support that I've received from the colleagues, friends, and family in my life. First and foremost of them all, I give my thanks to my parents, Judy and Robert, who have never wavered in encouraging and supporting me in every way possible as I bounded through this long journey across many states, schools, jobs, and countries.

My sincere gratitude goes to my advisor and mentor, Naomi Makins, who first drew me to the University of Illinois and SeaQuest by her enthusiasm and fervor alone. Anything insightful that I may present or state here in this paper has probably come from what I have learned through her patient and exceptional, thoughtful teachings.

Thanks also go out to my fellow UIUC graduate students have been the bedrock of my academic life here. Without the constant camaraderie of my friend Evan McClellan, I fear I would not have made it through these past years with my sanity intact. Likewise, without the tenacious, problem-solving mind of Bryan Kerns, I'm nearly certain that we would not have solved as many problems at SeaQuest as we have.

Particular thanks must extended to one, Markus Diefenthaler, who by sheer force of will held SeaQuest together and kept the collaborative effort focused for so many years. Markus also played a pivotal position as a role model to Evan, Bryan, and myself, exemplifying what an experimental physicist should strive to be.

I would also like to state my appreciation to the SeaQuest collaboration for placing their trust (and data!) in the hands of a humble database administrator and second-year graduate student. I have learned very much in the company of the great scientific tapestry that is the SeaQuest collaboration, and I hope to carry all of the lessons that I've learned with me throughout my life. Of course, the data presented here would simply not be possible without the grand coordinated efforts by all of the collaboration members who came together to assemble and execute such an ambitious and challenging experiment.

My thanks also go out to the staff and faculty of the Nuclear and High Energy Physics Laboratory for their availability and willingness to help a graduate student in need. In particular, the weekly feedback and insight from Jen-Chieh Peng has been simply invaluable. It should be stated that the only reason that the

PMT studies and prototyping got off the ground is as a result of the assistance received from Steve Errede and Allison Sibert. Finally, life at the NPL would simply grind to a halt if it were not for the constant availability and helpfulness of Mike Suchor.

I would also like to take a moment to thank the Department of Physics, the Nuclear Physics Laboratory for providing my employment as a research assistant throughout the years, thereby giving me the opportunity to do something I love for a living. Of course, none of that would have been possible without the magnanimous support of the National Science Foundation¹, to which we all owe a great debt of gratitude.

I would like to extend my appreciation and thanks to my parents-in-law, Les and Stacey Yonamine, for their whole-hearted encouragement and generous support over these past years. It has been truly uplifting gaining such great new family members, and their frequent rations of *da ono kine Lion coffee* have fueled me through many a work night.

Thank you to my sister Amy, who helped to house my wife, myself, and far too many of our pets. We have been very fortunate to have such a caring family member so close to keep us all in good humor and in good company.

Finally, I offer my most heartfelt gratitude to my wife, Carly, for the incredible leap of faith she took moving from Hawai'i on over to landlocked corn belt. She has been my partner and my rock, always at my side as I put myself through the rigors of these graduate studies. Her patience, understanding, companionship, and encouragement have truly been what has kept me going all these years.

¹NSF Grant No. 1506416

Table of Contents

List of Tables	ix
List of Figures	xi
List of Abbreviations	xvi
Preface	1
Chapter 1 Theoretical Background	3
1.1 Di-lepton production in nucleon-nucleon collisions	3
1.1.1 The Drell-Yan Process	4
1.1.2 Drell-Yan Kinematics	5
1.1.3 Cross-Section	7
1.2 Nucleon Structure	9
1.2.1 PDFs, Structure Functions, and the Quark Parton Model	9
1.2.2 Drell-Yan Tests of the QPM	13
1.2.3 QCD-Improved Drell-Yan	15
1.3 The Nuclear EMC Effect	16
1.3.1 Historical Overview	17
1.3.2 Fermi Motion	18
1.3.3 Universal x -dependence	20
1.3.4 Nuclear Dependence of R	22
1.3.5 Q^2 Dependence of EMC Effect	22
1.3.6 Shadowing	23
1.3.7 A-dependence of EMC Effect Slope	24
1.3.8 $x > 1$ Plateaus and Short Range Correlations	25
1.4 The Role of Drell-Yan in Studying the Nuclear Quark Sea	27
1.4.1 Previous Measurements	28
1.4.2 Nuclear Dependence of Drell-Yan at SeaQuest	31
Chapter 2 Apparatus	33
2.1 Apparatus Overview	33
2.2 Main Injector Proton Beam	34
2.3 Beam Intensity Monitor	36
2.4 The SeaQuest Targets	39
2.5 Focusing and Analyzing Magnets	41
2.6 Beam Dump, Shields, and Absorbers	43
2.7 Tracking Detectors	43
2.7.1 Triggering Hodoscopes	44
2.7.2 Drift Chambers	45
2.7.3 Proportional Tubes	47
2.7.4 Mass Resolution from Chamber Resolution	49

2.8	Trigger	51
2.8.1	Design Requirements	51
2.8.2	Trigger Hardware	52
2.8.3	Triggering Modes	54
2.9	Data Acquisition Systems	55
2.9.1	Main DAQ	56
2.9.2	Scaler DAQ	56
2.9.3	Beam DAQ	57
2.9.4	Slow Control Readout	57
Chapter 3 Data Productions		61
3.1	Raw Data Processing	62
3.1.1	CODA Event Format	62
3.1.2	ROC Data Bank Formats	63
3.1.3	Decoding Raw Data	64
3.2	Online and Offline Production Processing	66
3.2.1	Online Monitoring and SeaScope	66
3.2.2	Offline Batch Processing	68
3.3	RDBMS Data Structure	68
3.3.1	Querying Language	69
3.3.2	Atomic Data Schema	70
3.3.3	Production Management	72
3.4	Data Quality	72
3.4.1	Bit Pattern Fields	73
3.4.2	Spill Data Quality	73
3.4.3	Other Data Quality Bits	73
3.5	Discussion and Retrospective	75
3.5.1	On Scalability	75
3.5.2	Key Practices for Success	77
3.5.3	Future Technologies	78
Chapter 4 Photomultiplier Tube Base Upgrade		82
4.1	PMT Basic Construction and Operation	82
4.2	PMT Base Design Iterations	84
4.2.1	Original Base	87
4.2.2	Prototype Base v1	87
4.2.3	Prototype Base v2	88
4.2.4	Prototype Base v3	88
4.2.5	Prototype Base v4	90
4.3	PMT Base Comparisons	90
4.3.1	Testing Apparatus, Measurements, and Procedure	91
4.3.2	Original vs. Prototype v1	94
4.3.3	Original vs. Prototype v1 vs. Prototype v2	95
4.3.4	Prototype v1 vs. Prototype v3	96
4.3.5	Prototype v1 vs. Prototype v4	96
4.3.6	Base Comparison Conclusions	97
4.4	Base Manufacturing and Installation	98
4.5	New Base Performance	100

Chapter 5	Tracking and Data Selection	102
5.1	Reconstruction of SeaQuest Dimuons	102
5.1.1	Data Reduction	102
5.1.2	Track Finding	106
5.1.3	Kalman Filter Track Fitting	107
5.1.4	Vertex Fitting	109
5.1.5	Monte Carlo Validation	109
5.2	Analysis Cuts	111
5.2.1	Data Selection	111
5.2.2	Run Level Cuts	112
5.2.3	Spill Level Cuts	113
5.2.4	Event Level Cuts	115
5.2.5	Dimuon Cuts	116
5.2.6	Track Cuts	117
5.3	Analysis Data Set	119
Chapter 6	Analysis	121
6.1	Normalization of Relative Yields	121
6.1.1	Live Proton Count and Beam Attenuation	122
6.1.2	Target Density Ratios	123
6.1.3	Detector Efficiency	124
6.1.4	Detector Acceptance	126
6.2	Rate Dependence	129
6.2.1	Weights as Corrections	130
6.2.2	Tracking Efficiency	130
6.2.3	Empty Target Background Correction	137
6.2.4	Residual Rate Dependent	140
6.3	LD2 Contamination	141
6.3.1	Contamination Measurements	141
6.3.2	Contamination Correction	143
6.4	Isoscalar Corrections for ^{184}W and ^{56}Fe	144
6.5	Summary of Analysis Steps	145
6.6	Corrected, Aggregated Yields	146
Chapter 7	Results and Interpretation	155
7.1	Drell-Yan Event Distributions	155
7.2	The A-Dependence of the Drell-Yan Process	159
7.2.1	Systematic Uncertainties	159
7.2.2	M_{γ^*} , x_F , and p_T Dependence	160
7.2.3	x_1 , x_2 , and $\cos\theta_\mu$ Dependence	161
7.2.4	Integrated Ratio versus A	167
7.3	Discussion of $R_{DY}(x_2)$	168
7.3.1	Natural Log of Ratios	168
7.3.2	Comparison to E-772	172
7.3.3	Models For Nuclear Modifications	172
Chapter 8	Conclusions	175
References		177

List of Tables

1.1	Kinematic variables relevant to the Drell-Yan process	7
1.2	Measured EMC slopes dR_{EMC}/dx for $0.35 \leq x \leq 0.7$ [42].	24
1.3	Measured SRC scale factors, as measured by CLAS [48] and JLAB E02-019 [46, 43].	26
2.1	Characteristics of the seven SeaQuest target positions. The “Spills/Cycle” is only a typical configuration and can vary according to needs and running configurations. The non-zero interaction length of the empty flask is due to the $51\mu\text{m}$ -thick stainless steel end-caps of the flask and the $140\mu\text{m}$ -thick titanium windows of the vacuum vessel that contains it.	40
2.2	Parameters of all hodoscope planes. z -positions of H2Y and H4 half-planes are offset slightly due to the half-planes themselves overlapping.	45
2.3	Parameters of all chambers. Those of primed planes are almost the same as of unprimed planes. The z -positions of U and V planes are relative to those of X planes.	46
2.4	Combination of D1 and D3m chambers per data taking period.	46
2.5	Parameters of the four proportional tube planes.	49
2.6	Trigger settings for Run 3 and beyond. Prescale figures shown are typical values and were adjusted as experimental needs were tuned. NIM-1 and NIM-2 triggers’ exact conditions could be changed and reconfigured from the counting room.	54
3.1	The Spill table dataQuality bitpattern definition.	74
3.2	The Event table dataQuality bitpattern definition.	75
3.3	The Hit table dataQuality bitpattern definition.	75
5.1	The multiplicity limits imposed on detector groups by the tracking.	105
5.2	The experiment-specific data sources and their details.	112
5.3	Some spill data quality criteria.	114
5.4	Specific spill ranges of excluded data for all roadsets.	115
5.5	A list of all quality and kinematic cuts for selecting good quality dimuons from the target.	118
5.6	Total raw dimuon yields after all spill, track, dimuon, target, and quality cuts. Not shown: empty flask target and “none” target (for background study).	120
6.1	Live proton sums for each target for each roadset.	122
6.2	Details pertaining to the measurables and physical properties individual nuclear targets.	123
6.3	LD2 kEfficiency fit parameters with the χ^2 evaluation of the fit to the data.	136
6.4	The calculated normalization constants for the background proportion function, $C(I)$	139
6.5	Results of the isotropic analyses of the liquid deuterium samples. Two tests are applied to each sample, and all numbers are in percentages.	141
6.6	Renormalized compositions (in percentages) when considering only the hydrogen isotopes.	142
6.7	Atomic percentages of the deuterium targets with asymmetric uncertainties in the d percentage.	143
6.8	Total yields for each target for each x_2 bin normalized to the live protons on target. The LD2 has also been corrected for contamination in the last column.	147
6.9	Total Roadset 62 yields for each target for each x_2 bin normalized to the live protons on target. The LD2 has also been corrected for contamination in the last column.	148

6.10	Total Roadset 67 yields for each target for each x_2 bin normalized to the live protons on target. The LD2 has also been corrected for contamination in the last column.	149
6.11	Adjustment made to deuterium yields (per live proton) to correct for deuterium contamination. Shown are the kEff+BG corrected yields before and after the contamination correction.	150
7.1	The mean dimuon kinematics in each bin of x_2 . These should be considered to be approximate, as they are calculated as a weighted average from the data set when corrected for rate dependence in x_2	158
7.2	Systematic error arising from the uncertainty in the kEfficiency exponential fit parameter.	160
7.3	The total estimated <i>fractional</i> systematic uncertainty for the ratio measurements versus x_2	161
7.4	The tabulated data of $R_{DY}(A/D)$ for C/D, Fe/D, and W/D, fully with statistical uncertainty and upper and lower asymmetric systematic uncertainties.	167
7.5	The ratio values with asymmetric <i>statistical</i> uncertainties after combining the results using the natural log transform method.	171

List of Figures

1.1	Dimuon mass spectrum from E-866 $p + h$ collisions at 800 GeV/c [1].	4
1.2	The Drell-Yan process, an s channel interaction consisting of the annihilation of a quark with an anti-quark.	4
1.3	A depiction of the Collins-Soper reference frame by Evan McClellan (UIUC). The CS variables are best understood by identifying the three defined planes and then referring to the specific angles between them.	6
1.4	Parton distribution functions for quarks and gluons and their x and Q dependence as calculated to NNLO by the CTEQ-TEA (CT) global analysis group[10].	12
1.5	Diagrams of (a) LO Drell-Yan, (b) vertex correction, (c-d) gluon emission, (e-f) “Compton” scattering [19].	16
1.6	The original EMC ratio measurement by the EMC collaboration [23]. Overlaid is the naive, expected ratio taking into account only the Fermi motion of the nucleons in the heavy nucleus.	17
1.7	SLAC E-87 was quick to report its own finding months later confirming EMC’s observation [24].	18
1.8	A compilation of data regarding the measurement of F_2^A/F_2^D for nine different heavy targets with an overlaid fit by Chen et al. [32].	19
1.9	The Fermi momenta K_F (in GeV) for various nuclei of atomic weight A [33].	20
1.10	The ratio of $\sigma^{C(N)}/\sigma^d$ covering nearly the entire $x \in [0, 1]$ interval [36]. The four qualitative regions are denoted by roughly chosen boundaries x_1, x_2, x_3 (not to be confused with the beam and target x_1 and x_2 used in the rest of this paper).	21
1.11	Measurements of the change in the σ^{Ca}/σ^d ratio with respect to $\ln Q^2$ over a broad range of x [25].	22
1.12	The slopes of the EMC ratio for $0.35 < x < 0.7$ vs the average nuclear density scaled by $(A - 1)/A$ [35].	24
1.13	Per nucleon cross section ratios vs x at $0.8 < x < 2$ as measured by JLAB-E02019 with a line fit to the “plateau” regions at $1.5 < x < 1.9$ [43].	25
1.14	Per nucleon cross section ratios vs x at $0.75 < x < 3$ as measured by CLAS with a line fit to the “plateau” regions at $2.25 < x < 2.8$ [44].	25
1.15	The EMC slopes plotted against the 2N-SRC scale factors [51]. The fit is constrained to go through deuterium at (1,0).	27
1.16	E-866 Drell-Yan σ^A/σ^{Be} for iron and tungsten as a function of x_2 [53].	29
1.17	E-772 Drell-Yan $\frac{A}{2H}$ ratios for several nuclear targets as a function of x_2 (x_t) [54].	30
1.18	A pairplot of the kinematic coverage of SeaQuest, based on weighted Geant4 Monte Carlo simulation, with $4.2 \text{ GeV} < M_{\gamma^*} < 10 \text{ GeV}$ requirement imposed. The diagonal shows a histogram of the single kinematic, while the off-diagonal plots show the contours (bottom left) and scatterplots (top right) of the correlations between variables.	32
2.1	Perspective view of the SeaQuest spectrometer apparatus.	34
2.2	Spill structure of the beam delivered to SeaQuest.	35
2.3	Beam profile detailed by SWIC detectors along the NM beam line.	36

2.4	A single high-intensity event with majority of all detector elements firing off. White space within the rectangles indicates inactive elements whereas red, blue, and green represent elements which have fired during that event. Track reconstruction in these cases is impossible.	37
2.5	The Beam Intensity Monitor (BIM) Čerenkov counter. Measurements are in inches.	38
2.6	The layout of the target table and its seven target positions, as seen from above.	40
2.7	Perspective drawing of FMAG’s aluminum coils embedded in an arrangement of iron slabs.	42
2.8	Spectrometer layout of FMAG, KMAG, and Detector Stations 1-4.	44
2.9	Proportional tube top ($x - z$ plane) view.	48
2.10	Proportional tube side ($y - z$ plane) view.	48
2.11	A simplistic depiction of a track passing through from Station 1, through KMAG where its path is bent by the magnetic field, and then straight on through Stations 2 and 3.	50
2.12	The trigger system at SeaQuest is composed of 9 CAEN v1495 FPGA modules. These modules output a trigger signal to the Trigger Supervisor, which tells the DAQ when to record data.	52
2.13	A block diagram of the major functions of the v1495 FPGA [62].	53
2.14	A standard Beam DAQ display showing beam characteristics and spill data. The bottom two plots show the output of the QIE on a ms scale.	58
3.1	(Left) The Prestart event, with event type 17; (Right) The Go event, with event type 18[64].	63
3.2	(Left) A standard physics event, with various event types, filled with a series of data banks; (Right) The ROC data bank format[64].	63
3.3	The SeaScape Online Monitor displaying the hit distributions for the Station 1 hodoscopes. Run selection is in the top red button selector, and the Spill selection is in the bar just below the top.	67
3.4	SeaScape displaying target position data for a sequence of spillIDs. Up to six different values can be shown at once against each other (scale-able for easier comparison) on each of the Slow Control, BeamDAQ, ScalerDAQ pages.	67
3.5	A visualization of the hierarchy and categorization of the data stored in the MySQL tables at SeaQuest. Data can be grouped into Run, Spill, and Event hierarchies, and categorized into raw data (green), calibrations (orange), BeamDAQ data (purple), slow control data (blue), reconstructed data (yellow), and aggregated metadata (red).	71
3.6	A traditional approach to HEP data handling at FNAL [68].	80
3.7	A proposal that suggests using a combination of horizontally-scalable NoSQL technologies to store all data [68]. Blobs are “binary large objects” that are application-dependent – they could be CODA or ROOT files, for instance.	80
4.1	(Left) Histogram of hodoscope ‘hits’ in a typical event; (Right) Histogram of high-intensity event, with marked sagging most noticeably in the middle of the y-measuring hodoscopes	83
4.2	A diagram of typical PMT operation. The circuit controlling the voltage-dropping resistors is the part that was upgraded in this chapter.	83
4.3	MOSFET showing gate (G), body (B), source (S) and drain (D) terminals. The gate is separated from the body by an insulating layer (white) [73]	85
4.4	Suggested voltage division schemes for gain vs. timing/linearity compromise [70].	86
4.5	The original PMT base inherited from the ARGUS and HERMES experiments.	87
4.6	The voltage division between subsequent stages for the original PMT base design when supplied with -1500 V.	88
4.7	The Prototype v1 board circuit diagram received from Fermilab Particle Physics Division [74]. The parts are denoted as R: resistor, C: capacitor, Q: MOSFET transistor, D: Zener diode.	89
4.8	The voltage division between subsequent stages for the Prototypes v1, v2, and v3 PMT base designs when supplied with -1500 V.	89
4.9	The Prototype v3 board: 3 more transistorized stages than the Prototype v1 design.	90
4.10	The (negative) voltage between subsequent stages for the Prototype v4 PMT base.	91
4.11	Inside of the light box, with the prototype board (left) wired up to a Philips XP-2008 PMT (middle), facing a fast-pulsing LED source (right).	92

4.12	Measurements of the (a) signal amplitudes and (b) bleeder and signal currents in the original and prototype v1 PMT bases. The HV was inadvertently set to -1600V V for this test instead of the intended -1500 V.	94
4.13	Measurements of the (a) signal amplitudes and (b) bleeder and signal currents in the original, prototype v1, and prototype v2 PMT bases.	96
4.14	Measurements of the (a) signal amplitudes and (b) bleeder and signal currents in the prototype v1 and prototype v3 PMT bases.	97
4.15	Measurements of the (a) signal amplitudes and (b) bleeder and signal currents in the prototype v1 and prototype v4 PMT bases.	97
4.16	(Left) The components of the disassembled original base and the manufactured new base PCB with attached daughter board. (Right) The original and new bases fully assembled, side-by-side.	98
4.17	The γ -decay energy spectrum of Cs-137 as seen through a scintillator-PMT-qVt setup [78]. Here, the Compton edge can be seen at channel 110.	100
4.18	The number of hits per event for the Station 1 hodoscopes for several intensity ranges. Shown is (left) the top x -measuring array and (right) the left y -measuring array. I_p is in terms of number of incident protons for the event.	101
5.1	(Left) A flow chart mapping out the decisions in how to treat clusters. (Right) A histogram of time differences of hits within a cluster. The peak below 10 ns is from electronic noise, and is the deciding factor in electronic noise classification.	104
5.2	(Left) The definition of the sagittas s_1 and s_2 , with the red asterisk denoting the actual hit positions in the stations. (Right) The ratio of the <i>true</i> ratios of s_1/s_2 across muons of all kinds of kinematics (within acceptance) shows a single mean value with a defined variance.	106
5.3	Shown is the application of the Kalman filter technique to track fitting [79]. The vertical lines correspond to the detector planes with indicated on it the measurement points and their errors. The cones represent the uncertainty in the reconstructed track parameters. As can be seen the knowledge of the track parameters is step wise updated with each measurement (Kalman filter technique applied from right to left).	108
5.4	The residuals of reconstructed track kinematics against MC truth values.	110
5.5	The residuals of fitted track vertex kinematics against MC truth values.	110
5.6	The p_z distribution of muons that form like-sign dimuon pairs, which is a source of background for DY signal events. In cases where $14 \leq \text{numHits} < 18$, imposing a cut on $p_z > 18 \text{ GeV}$ removes a majority of this background.	119
6.1	The Station 2 drift chamber occupancy for the four targets as a function of intensity. The error bars show a systematic deviation of occupancies from the mean for a given intensity bin. The inset plot shows a KDE plot of distributions for the intensity bin in which the values differ the most. The triangle markers represent the modes of the distributions.	125
6.2	The ratio of successful reconstruction to all thrown by MC for two chamber efficiency models. The linear decreases in efficiency are not significantly different from each other.	126
6.3	Ratios of accepted dimuons from heavy targets to deuterium for all kinematic regions of interest.	128
6.4	The rate dependence, depicting the decline in yields per triggering proton as a function of the overall environmental chamber intensity.	129
6.5	The rate dependence has a different behavior for different kinematics. Shown is the same plot as in Fig. 6.4, but in three statistically equal x_2 bins.	129
6.6	The linear tendency of detector multiplicity to increase with intensity is observed by looking to unbiased NIM3 events. Here, the occupancy per event of the drift chambers at station 2 are shown as a function of intensity.	132
6.7	The ratio of “messy” to “clean” dimuon samples from the Deuterium target in Roadset 67 data. This shows the linear relationship of the tracking efficiency to the intensity of the events. The shaded band indicates the 95% confidence band of the fit.	133
6.8	Tracking efficiency with the data broken into three statistically equivalent bins in the five primary kinematics, $(x_1, x_2, \theta_\mu, \phi_\mu, p_T)$. Significant kinematic dependence is seen in x_1 and x_2 .	134

6.9	The kinematics x_F and M_{γ^*} are investigated. A clear kinematic dependence exists in x_F while M_{γ^*} is largely consistent.	135
6.10	The efficiencies after re-weighting using kEfficiency values calculated when binned in x_F	135
6.11	Comparing the tracking efficiency curves across (left) the five targets and (right) two temporally different subsets (different roadsets) of data. Confidence bands removed from the target plot for the sake of clarity.	136
6.12	The intensity-dependence of the Empty target and "None" target, with fit to $p_0(1 + p_1 \cdot I_c^2)$ for each. 95% confidence band shown.	138
6.13	(Left) The remaining rate dependence behavior after corrections are made and (right) the rate dependence of a cross section ratio measurement with the linear fit and the 95% confidence band shown.	140
6.14	(Left) The ratio of free neutron to proton F_2 structure functions following a detailed parametrization [84] and the isoscalar correction factor [85], $R_{ISO}(x)$, for iron and tungsten.	145
6.15	A comparison of x_2 distributions between roadsets for each target, and all targets combined. The difference between the two normalized distributions shown by the red data points, with a red line at zero. The χ^2 shown is for the comparison between the roadset differences and the line at zero.	151
6.16	A comparison of x_1 distributions between roadsets for each target, and all targets combined. The difference between the two normalized distributions shown by the red data points, with a red line at zero. The χ^2 shown is for the comparison between the roadset differences and the line at zero.	152
6.17	A comparison of M_{γ^*} distributions between roadsets for each target, and all targets combined. The difference between the two normalized distributions shown by the red data points, with a red line at zero. The χ^2 shown is for the comparison between the roadset differences and the line at zero.	153
6.18	A comparison of x_F distributions between roadsets for each target, and all targets combined. The difference between the two normalized distributions shown by the red data points, with a red line at zero. The χ^2 shown is for the comparison between the roadset differences and the line at zero.	154
7.1	Event distributions for all combined data versus dimuon kinematic variables x_F, M_{γ^*}, p_T , and p_z	156
7.2	Event distributions for all combined data versus quark kinematic variables x_1, x_2 , and the cosine of the polar decay angle, θ_μ	157
7.3	The measured Drell-Yan per-nucleon cross section ratio against invariant dimuon mass. Only statistical uncertainty is shown.	162
7.4	The measured Drell-Yan per-nucleon cross section ratio against dimuon Feynman- x . Only statistical uncertainty is shown.	162
7.5	The measured Drell-Yan per-nucleon cross section ratio versus transverse momentum of the muon pair. Only statistical uncertainty is shown.	163
7.6	The measured Drell-Yan per-nucleon cross section ratio versus longitudinal momentum of the muon pair. Only statistical uncertainty is shown.	163
7.7	The measured Drell-Yan per-nucleon cross section ratio versus fractional momentum quantity, x_1 . Only statistical uncertainty is shown.	164
7.8	The measured Drell-Yan per-nucleon cross section ratio versus the cosine of the polar muon production angle, θ_μ . Only statistical uncertainty is shown.	164
7.9	The measured Drell-Yan per-nucleon cross section ratio versus fractional momentum quantity, x_2 . Only statistical uncertainty is shown. Overlaid is the data from the E-772 experiment.	165
7.10	The same plot, but zoomed into $R=[0.7, 1.2]$ for closer comparison with E-772 data.	165
7.11	The measurement of R_{DY} for the three roadsets. Their results are combined to render the values found in Figure 7.9.	166
7.12	The measured Drell-Yan per-nucleon cross section ratio versus fractional momentum quantity, x_2 , grouped into "high" p_T and "low" p_T . Only statistical uncertainty is shown.	166

7.13	Event distributions for all combined data versus quark kinematic variables x_1, x_2 , and the cosine of the polar decay angle, θ_μ	167
7.14	A depiction of the asymmetric phase space for A/B versus $\ln(A/B)$	168
7.15	The combined results of the log of R_{DY} from the three roadsets.	169
7.16	The combined results of the log of R_{DY} transformed back into a plot of R_{DY} by mapping the data points and error bounds with the exponential function.	170
7.17	The measured Drell-Yan per-nucleon cross section ratio versus fractional momentum quantity, x_2 . Overlaid is the data from the E-772 experiment. Only statistical uncertainty is shown. There exists an additional $\sim 3\%$ systematic uncertainty.	172
7.18	Virtual meson cloud model predictions for DY ratios for σ^{F^e}/σ^D [87]. Models shown are tuned for 800 GeV experiments, and may not apply exactly SeaQuest beam energies.	173
7.19	The SeaQuest results and a composite model prediction from Kulagin and Petti for DY ratios σ^C/σ^D , σ^{F^e}/σ^D , and σ^W/σ^D [89, 90].	174

List of Abbreviations

ACNET	Accelerator control system network
AD	Accelerator division (of Fermilab)
BIM	Beam intensity monitor
BNC	Bayonet Neill-Concelman signal cable connector standard
BOS	Beginning of spill
BPM	Beam profile monitor
CEBAF	Continuous Electron Beam Accelerator Facility
CODA	CEBAF On-line Data Acquisition
CS	Collins-Soper reference frame
CTEQ	Coordinated theoretical-experimental project on QCD
DAQ	Data acquisition
DC	Drift chamber
DIS	Deep-inelastic scattering
DY	Drell-Yan
EOS	End of spill
EOF	End of file
EPICS	Experimental physics and industrial control system
EVIO	CODA event input/output
FEE	Front-end electronics
FNAL	Fermi National Accelerator Laboratory (Fermilab)
FPGA	Field programmable gate arrays
IC	Ion chamber
LED	Light-emitting diode
LINAC	Fermilab Linear Accelerator
LO	Leading order

MI	(Fermilab) Main Injector
MOSFET	Metaloxidesemiconductor field-effect transistor
NDF	Neutral density filter
NLO	Next-to-leading order
NM	Neutrino-muon beam line
NNLO	Next-to-next-to-leading order
PCB	Printed circuit board
PDF	Parton distribution function
PID	Particle identification
PMT	Photomultiplier tube
RF	Radio frequency
RFQ	Radio frequency quadrupole
ROC	Readout controller
SEM	Secondary emission monitor
SHV	Secure high voltage connector standard
SLAC	Stanford Linear Accelerator Center
SRC	Short range correlations
TS	Trigger supervisor
QCD	Quantum chromodynamics
QED	Quantum electrodynamics
QIE	Charge (Q) integrator and encoder
RDBMS	Relational database management system
SQL	Structured querying language
SWIC	Segmented wire ion chamber
VME	Versa Module Europa

Preface

Matters are rarely as simple as we hope them to be, and atomic nuclei are no exception. When physicists began to get a handle on proton structure, it was hoped that nuclei could be simply understood as a simple collection of the same familiar protons and neutrons. Unfortunately, it turned out that *more is different*, and higher-order emergent behavior reared its head. It appeared that the distribution of partons within the proton became *altered* when the proton was confined within a nucleus – a phenomenon which is generally referred to as the EMC Effect. Since the discovery of this behavior, data and theories have proceeded to accumulate for over 30 years without a completely satisfying explanation, though recent breakthroughs give a reason for optimism.

The SeaQuest experiment was commissioned and began taking data in 2012 with the expressed purpose to, among other physics goals, shed some additional light on this phenomenon. I joined the collaboration in 2009 and have had the distinct pleasure of being a part of a large-scale, medium-energy experiment as it went from being an empty hall to a testing ground for nuclear physics. It has truly been a rewarding experience being able to contribute to a coordinated, monumental effort from the point of assembly all the way to its first publishable results (or at least very close). In this thesis, I will describe in detail these contributions to the experiment.

Chapter 1 will provide a broad outline regarding the background to the underlying physics and phenomenology. I describe the state of the field regarding the physics of unpolarized sea quark distributions along with what SeaQuest can contribute. Emphasis is placed on the phenomenon of modification of the parton distribution functions when in the presence of a nuclear medium.

In Chapter 2, the apparatus of the SeaQuest experiment is laid out, briefly covering every significant aspect, including the beam, targets, magnets, detectors, and data acquisition systems. For further information on my contributions to the SeaQuest experiment's apparatus and operations, I have laid out in the following chapters my work on the curation of the collaboration's data productions (Chapter 3) and my work upgrading the hodoscope photomultiplier tube bases for improved performance (Chapter 4). These account for a significant amount of my own time and effort on the experiment but are *not* required for understanding

the background and analysis regarding the primary measurement of this paper.

Chapter 5 summarizes how the data used in these studies are reconstructed and selected. This includes a brief overview of the tracking algorithm that reconstructs tracks and dimuons from the raw detector readouts, and then a thorough rundown of the data reduction, showing in great detail how good events are chosen from the whole set.

In Chapter 6, I step through the many stages of the analysis, going from the raw reconstructed data all the way through the target-to-target normalization and a litany of corrections in order to arrive at a final result. Particular emphasis is placed on the corrections pertaining to the rate dependence, which has been the key obstacle in preparing preliminary results for all SeaQuest analyses. This section goes perhaps into more detail than needed, which is intentional, as this document is, in part, meant to act as an explanatory road map for future students and researchers who wish to perform this same analysis on the full SeaQuest dataset in the future.

Chapter 7 proceeds to summarize the fully processed and corrected findings, providing a comprehensive review of R_{DY} , the per-nucleon Drell-Yan cross section ratio of heavy nuclei to deuterium, against various kinematics. Results are briefly discussed and compared against theoretical predictions, though more refined corrections and more statistics from further SeaQuest data taking are required to come to any firm conclusions.

Chapter 1

Theoretical Background

The topic of this paper is the exploration of the sub-structure of the *nucleon*, a particle that makes up an atom's nucleus which can be either a proton (p) or a neutron (n). By sub-structure, I refer to the nucleon's composition and the momentum distributions carried by the nucleon's constituent particles, quarks and gluons.

While a complete review of the history and physics behind nucleon structure and its investigative probes is beyond the scope of this paper, a brief overview of Drell-Yan, parton distribution functions, and nuclear structure phenomenology will help in understanding concepts and terminology relevant to this and later chapters.

1.1 Di-lepton production in nucleon-nucleon collisions

Studying the production of pairs of leptons resulting from hadron-hadron collisions has proven to be a powerful tool in probing nucleon structure and parton distribution functions (PDFs). Figure 1.1 shows the mass spectrum of the muon mode of di-lepton production in proton-nucleus collisions measured at the E-866/NuSea experiment at Fermilab [1]. The resonances of the J/Ψ , Ψ' , Υ , Υ' , and Υ'' can be seen atop a smooth, continuous distribution which decreases with mass. The process responsible for this distribution of $\mu^+\mu^-$ pairs is the Drell-Yan process [2], which is illustrated in the Feynman diagram in Fig. 1.2. This process is characterized by a quark annihilating with an anti-quark to form a virtual time-like photon which then decays into a pair of leptons. The study of this process can lend some insight into nucleon structure due to the fact that the mass and momentum of the di-lepton pair directly reflects the momentum distributions of the interacting quarks and anti-quarks. The most comprehensive and precise knowledge and models of these momentum distributions come primarily from deep-inelastic lepton scattering (DIS) and the Drell-Yan (DY) process.




Figure 1.1: Dimuon mass spectrum from E-866 $p + h$ collisions at 800 GeV/c [1].




Figure 1.2: The Drell-Yan process, an s channel interaction consisting of the annihilation of a quark with an anti-quark.

1.1.1 The Drell-Yan Process

The study of continuum di-lepton production in hadron collisions,

$$h_A + h_B \rightarrow l^+ l^- + X \quad (1.1)$$

can be used to study hadronic structure in a way that is complementary to the study of deep-inelastic scattering,

$$l + h \rightarrow l' + X. \quad (1.2)$$

In 1970, S. Drell and T.M. Yan were the first to suggest that, at high $Q^2(= M_{l^+l^-}^2) \geq 16\text{GeV}^2$, the quarks inside the hadrons h_A and h_B can be considered free fermions in the instantaneous moment that they interact. The Drell-Yan model addresses the dominant subprocess here,

$$q_A + \bar{q}_B \rightarrow \gamma^* \rightarrow l^+l^- \quad (1.3)$$

as an electromagnetic annihilation process. In this high- Q^2 kinematic space, the final state of the hadrons that contain these quarks becomes irrelevant. By the energy-time uncertainty principle [3] of $\Delta E \Delta t \sim \hbar$, the timescales under consideration are $< 10^{-25} s$, and the corresponding distances are $< 10^{-17} m$, where the size of a nucleon is $\sim 10^{-15} m$ [4]. At these short space-time scales, electromagnetic annihilation dominates this quark-quark interaction. This is reinforced by QCD and the running coupling constant, or *asymptotic freedom* [5], in which the coupling of the strong force gets weaker with higher Q^2 .

1.1.2 Drell-Yan Kinematics

In the center-of-mass frame, the Drell-Yan process can be broken down into three stages with three sets of kinematics (refer to Fig. 1.2). Beginning with the quarks, x is defined as the fraction of the hadron's momentum carried by the interacting quark or antiquark. Conventionally in fixed-target experiments, subscripts are assigned as x_1 and x_2 , which refer to the quark/antiquark from the beam and the antiquark/quark from the target, respectively. This x is called the Bjorken x , and is well-known in DIS processes to have a value of

$$x = -q^2/2p \cdot q \quad (1.4)$$

where p and q are the 4-momenta of the hadron and the photon. In DY, the same momentum fraction x is used to describe part of the fundamental basis of the interaction. In this case, q is the 4-momentum of the virtual photon and p is the 4-momentum of the hadron that contains the quark in question.

The next stage of the process is the virtual photon. This photon's properties are effectively equivalent to the 'dimuon' or 'dilepton', which are the terms more commonly used in referring to kinematics. The first of its relevant kinematics is its mass M_{γ^*} , which represents its energy and virtuality. The value x_F , or Feynman-x, is the fraction of the maximum possible longitudinal momentum carried by the virtual photon in the beam direction. The transverse momentum, p_T , and the azimuthal production angle, ϕ_{γ^*} are the remaining kinematics associated with the virtual photon.

The third stage regards the pair of leptons produced. In the frame of the virtual photon, there is a polar and azimuthal decay angle, θ_μ and ϕ_μ , respectively, for each of the decay muons. It becomes impossible,

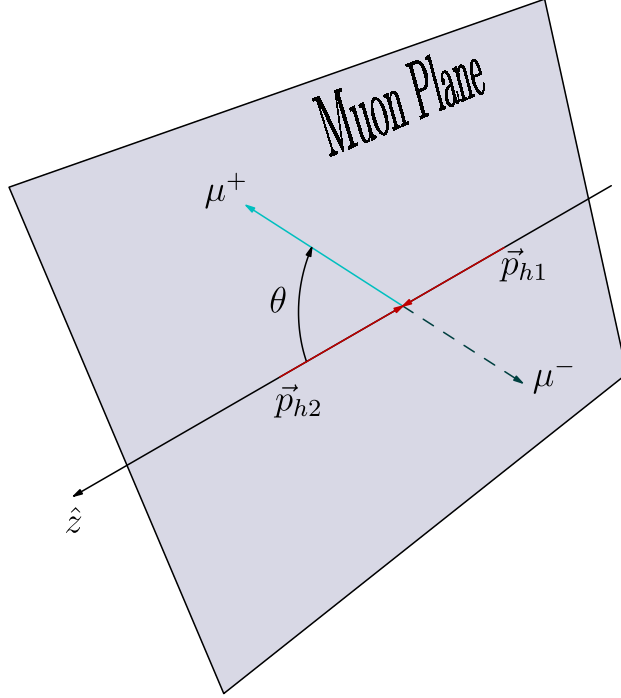


Figure 1.3: A depiction of the Collins-Soper reference frame by Evan McClellan (UIUC). The CS variables are best understood by identifying the three defined planes and then referring to the specific angles between them.

however, to reconstruct these variables, as the individual transverse momenta of the quarks are unknown, and thus the quark-antiquark annihilation axis is unknown. This is remedied by shifting the process into the Collins-Soper (CS) reference frame [6] which orients the reference axis to be parallel to the bisector of the angle between the interacting hadrons in the rest frame of the muon pair. A depiction of the Collins-Soper frame can be found in Fig. 1.3. In total, this brings a total of eight kinematic variables, summarized in Table 1.1.

Experimentally, six independent variables are measured, which form a basis by which all eight can be known. This is due to the fact that two pairs of variables, (M_{γ^*}, x_F) and (x_1, x_2) are correlated by the following, ignoring quark masses and considering $p_{\perp} \ll p_{\ell}$:

$$x_F \equiv \frac{p_{\ell}}{\sqrt{s}/2} \approx x_1 - x_2 \quad (1.5)$$

$$M_{\gamma^*}^2 \equiv E^2 - |p_{\ell}|^2 \approx sx_1x_2 \quad (1.6)$$

$$E = \frac{1}{2}(x_1 + x_2)\sqrt{s} \quad (1.7)$$

$$p_{\ell} = \frac{1}{2}(x_1 - x_2)\sqrt{s} \quad (1.8)$$

In this frame, the longitudinal momenta of the quarks are $x_1\sqrt{s}/2$ and $-x_2\sqrt{s}/2$, with \sqrt{s} being the center

Variable	Description
$x_{1/2}$	Momentum fraction of the beam/target quark
M_{γ^*}	Mass of the virtual photon (dimuon)
x_F	Fraction of the max. possible p_ℓ carried by virt. photon
p_T	Transverse momentum carried by the virt. photon
θ_μ, ϕ_μ	Polar and azimuthal decay angle of one of the muons, in the CS ref. frame
ϕ_{γ^*}	Azimuthal production angle
α	The fine structure constant
$K(x_1, x_2)$	High-order QCD correction term
\sqrt{s}	Center of mass energy of the hadronic collision
$\sqrt{\hat{s}}$	Center of mass energy of the $q\bar{q}$ collision
Q^2	Four-momentum of the intermediate time-like photon, squared
$q_i^{t/b}(x)$	The quark number density in the nucleon of the target/beam

Table 1.1: Kinematic variables relevant to the Drell-Yan process .

of mass energy of the hadronic collision. So, by measuring the 3-momenta of the μ^+ and μ^- , the quantities $(M_{\gamma^*}, x_F, p_T, \phi_{\gamma^*}, \theta_\mu, \phi_\mu)$, which is sufficient to calculate the remaining (x_1, x_2) in our approximation.

1.1.3 Cross-Section

The base measurement in particle physics is the cross section, or the effective area that quantifies the likelihood of a scattering event. Anything further such as quark probability distribution functions is merely an *interpretation* of this cross section. It is the case that this interpretation has been validated many times over, but it is important to keep in mind what is measured and what is interpreted / extracted in particle physics experiments.

As the participating quarks are asymptotically free within each hadron, there will be no correlations between the probability distributions of the annihilating particles, and the process is independent of the distributions. As a result, the cross section of the Drell-Yan process can be reduced to a function of the electromagnetic annihilation process and the quark probability distribution functions. With these components and some QCD considerations, we can construct it piece by piece.

The first step is to begin with the known hard scattering cross section of $\epsilon + \bar{\epsilon} \rightarrow l^+l^-$, where ϵ is an arbitrary particle. This cross section[7] is given by

$$\sigma(\epsilon\bar{\epsilon} \rightarrow l^+l^-) = \frac{4\pi\alpha^2}{3M_{\gamma^*}^2} e_f^2 \quad (1.9)$$

where α is the electromagnetic fine structure constant, e_f is the charge of the particle, and M_{γ^*} is the dilepton mass. With this, we add the QCD consideration that only $q\bar{q}$ of opposite color can annihilate with

each other into a colorless virtual photon. Possible combinations are $R\bar{R}$, $B\bar{B}$, and $G\bar{G}$ out of 3×3 possible cases. As such, an overall factor of $\frac{1}{3}$ is added to this cross section. Finally, factoring in the conservation of flavor (there can only be $u\bar{u}$, $d\bar{d}$, etc. combinations) and the quark structure of hadrons A and B, we use the product ($q_f^A(x_1)\bar{q}_f^B(x_2)$) of the quark probability distributions for finding quarks of the same flavor-antiflavor combination in the two hadrons. It must also be considered that the quark or antiquark may be found in either hadron A or hadron B. The product of these three factors leads us to the Drell-Yan cross section[8]

$$\frac{d^2\sigma}{dx_1 dx_2} = \frac{1}{3} \frac{4\pi\alpha^2}{3M_{\gamma^*}^2} \sum_f e_f^2 [q_f^A(x_1)\bar{q}_f^B(x_2) + \bar{q}_f^A(x_1)q_f^B(x_2)] \quad (1.10)$$

$$= \frac{4\pi\alpha^2}{9sx_1x_2} \sum_f e_f^2 [q_f^A(x_1)\bar{q}_f^B(x_2) + \bar{q}_f^A(x_1)q_f^B(x_2)] \quad (1.11)$$

where the sum is summing over flavors of quarks ($f \in \{u, d, s, \dots\}$). This can be evaluated in terms of the measurables M_{γ^*} and x_F via Equations 1.5 and 1.6

$$M_{\gamma^*}^2 \frac{d^2\sigma}{dM_{\gamma^*}^2 dx_F} = \frac{1}{3} \frac{4\pi\alpha^2}{3M_{\gamma^*}^2} \frac{x_1x_2}{x_1+x_2} \sum_f e_f^2 [q_f^A(x_1)\bar{q}_f^B(x_2) + \bar{q}_f^A(x_1)q_f^B(x_2)] \quad (1.12)$$

where x_1 and x_2 can be expressed as

$$x_1 = \frac{1}{2} \left[\sqrt{x_F^2 + 4\tau} + x_F \right], \quad x_2 = \frac{1}{2} \left[\sqrt{x_F^2 + 4\tau} - x_F \right], \quad \tau = \frac{M_{\gamma^*}^2}{s} \quad (1.13)$$

The cross section can also be represented by dimensionless variables in its scaling form,

$$s \frac{d^2\sigma}{d\sqrt{\tau} dy} = \frac{1}{3} \frac{4\pi\alpha^2}{3} \sum_f e_f^2 [q_f^A(x_1)\bar{q}_f^B(x_2) + \bar{q}_f^A(x_1)q_f^B(x_2)] \quad (1.14)$$

where we introduce the rapidity term y in describing x_1 and x_2 ,

$$y = \frac{1}{2} \ln \frac{E + p_\ell}{E - p_\ell} = \frac{1}{2} \ln \frac{x_1}{x_2}, \quad x_1 = \sqrt{\tau} e^y, \quad x_2 = \sqrt{\tau} e^{-y} \quad (1.15)$$

It should be noted that with collider experiments, it is conventional to refer to the hadrons A and B in terms of the beam and target hadrons, respectively, and as such, x_1 refers to the the quark in the beam hadron and x_2 refers to the quark in the target hadron.

In each of these equations 1.11, 1.12, and 1.14, the cross sections can be factored into two parts: one subprocess cross section and one part that has only a dependence on the parton distribution functions. They are independent of each other because one of the staples of the quark parton model is that the PDFs ($q(x)$)

and $\bar{q}(x)$) are independent of the process by which they are probed. This feature is commonly referred to as the “universality” of the PDFs.

1.2 Nucleon Structure

There are a few distributions, functions, terminologies, sum rules, and integrals that must be established and explained in order to speak the language of nucleon structure. Here, some of these are explained so that they may be used freely later. Also included are brief sections about how nucleon structure was tested by the Drell-Yan process, and how the QCD understanding of the nucleon was used to improve the understanding of the Drell-Yan process.

1.2.1 PDFs, Structure Functions, and the Quark Parton Model

In the context of hadron-hadron collisions, one can interpret the PDFs as follows: consider two hadrons A and B colliding; a parton of type a ($a \in \{u, d, s, g, \dots\}$) comes from A and carries with it a fraction of A’s momentum (x_A). The same goes for hadron B; a parton of type b comes from B and carries momentum fraction x_B . Now, the probability of finding the discussed parton from A at momentum fraction x_A is given by $q_{a/A}(x_A)$. Likewise, the probability of finding the discussed parton from B at x_B is $q_{b/B}(x_B)$.

In general, experimental data measuring F_2 is used in conjunction with some general constraints in order to arrive at the calculated PDFs. One constraint is to consider the known number of valence quarks in a nucleon. Since the function $q_f(x)$ can be interpreted to be the number density of quark flavor f as a function of momentum fraction x , the value $q_f(x)dx$ represents the number of quarks with flavor f with fractional momentum in the range of $[x, x + dx]$. Therefore, the known number of valence quarks of flavor f (N_f) provides the following condition.

$$\int_0^1 dx [q_f(x) - \bar{q}_f(x)] = N_f \quad (1.16)$$

Another condition considers the value $xq(x)dx$, which is the total momentum fraction of the hadron of the number of quarks $q(x)dx$. This provides a logical constraint is the summation of the momentum fractions of all partons must add up to the full nucleon momentum:

$$\sum_{q,\bar{q},g} \int_0^1 dx [xq(x)] = 1 \quad (1.17)$$

It should also be noted that the quark PDFs can be split with respect to their spin alignment (+) or

anti-alignment (-) with the nucleon's spin

$$q(x) = q^+(x) + q^-(x) \tag{1.18}$$

and the polarized (or helicity) PDF can be defined as

$$\Delta q(x) = q^+(x) - q^-(x) \tag{1.19}$$

though, due to the fact that SeaQuest uses neither a polarized $\Delta q(x)$. There are two standing proposals, however, to extend the SeaQuest spectrometer for an additional period of data taking, augmenting it with a polarized proton beam (FNAL E-1027) and/or a polarized target (FNAL E-1039). The polarized beam proposal adds two beam-polarizing apparatuses named *Siberian Snakes* to provide sufficiently polarized protons. The polarized target modification calls for the construction and installation of a new cryogenic liquid target system. Both proposals have been approved by the FNAL review committee and are likely to be acted upon near the end of the SeaQuest contract.

The structure functions F_1 , F_2 , and g_1 of the nucleon are interpreted in the QPM as the charge-weighted sums over the quark flavors of the corresponding PDFs:

$$F_1(x) = \frac{1}{2} \sum_f e_f^2 q_f(x) \tag{1.20}$$

$$F_2(x) = \sum_f e_f^2 x q_f(x) \tag{1.21}$$

$$g_1(x) = \frac{1}{2} \sum_f e_f^2 \Delta q_f(x) \tag{1.22}$$

where here f represents all flavors of quarks and anti-quarks. We see above the relation between $F_1(x)$ and $F_2(x)$ known as the Callan-Gross relation ($F_2 = 2xF_1$). This relation relies on the partons having spin- $\frac{1}{2}$ and no transverse momentum in the infinite momentum frame. The full expression is

$$F_2 = 2xF_1 \frac{1 + R}{1 + 2M_N x/\nu} \tag{1.23}$$

where M_N is the nucleon mass and $R = \sigma_L/\sigma_T$ is the ratio of cross sections for absorbing a longitudinal to that for a transverse photon. Experimentally, R is small[9] ($\lesssim 0.1$) for $x \gtrsim 0.1$ and for $Q^2 \gtrsim 5 \text{ GeV}^2$. Also, when assuming the infinite momentum frame, ν , the energy transferred by the photon is assumed to trend towards infinity. As a result, this relation reduces to the commonly-used Callan-Gross relation.

Due to the complex nature of lattice QCD¹ simulations, the parton distributions $q_f(x)$ and $\bar{q}_f(x)$ within the nucleus are determined empirically, with only a few rules based on theory. The primary probe used to measure them has been deep inelastic scattering (DIS), examining the inclusive jet production in order to get a clearer picture of nucleon structure. Data from many different experiments are combined to extract the unpolarized PDFs. The earlier parametrizations of the PDFs first relied on fits to the measured structure functions solely from DIS experiments which are primarily sensitive to the light quark distributions and are unable to distinguish between quarks and antiquarks of the same flavor. More modern parametrizations use many different processes to extract supplementary and complementary information that can be incorporated. Lepton-charge asymmetry observed in W^\pm production provides additional light quark distribution information. Jet production and photon measurements are also used to set constraints on gluon distributions. It should be noted that Drell-Yan dilepton production from such experiments as E605 and E866 that have contributed constraints on the light anti-quark distributions from the nucleon sea.

Current experimental determinations of these PDFs compiled by CTEQ (Coordinated Theoretical-Experimental Project on QCD) and TEA and are illustrated in Figures 1.4a and 1.4b. The CTEQ-TEA calculated these PDFs from data on inclusive, high-momentum transfer processes, for which perturbative QCD is expected to be reliable. In the case of deep inelastic lepton scattering, only data with $Q > 2$ GeV is used. Data in this region are expected to be relatively free of non-perturbative effects, such as higher twists or nuclear corrections. Thus, there is no need to introduce phenomenological models for nonperturbative corrections beyond the leading-twist perturbative contributions[10].

Looking to Figure 1.4 we see that at $x > 0.1$, u and d quarks dominate \bar{u} and \bar{d} quarks. The SeaQuest spectrometer (described in detail in the next chapter) is a “forward spectrometer”, which means that, for PDF-relevant purposes, the high- x_F ($x_F > 0$) kinematic phase space is explored, which translates to high x_1 and low x_2 . With this consideration, the Drell-Yan process measured at SeaQuest is probabilistically dominated by a quark from the beam annihilating with an anti-quark from the target. Further, all anti-quarks that exist in the target nucleons must come from what are called the *sea quarks*, or the virtual $q\bar{q}$ pairs that arise from gluons splitting.

Other interesting observations and measurements that have gone into these distributions are the observations of momentum contributions from all quarks and antiquarks and the momentum fraction from just antiquarks. The momentum sum

$$x_{tot} = \int_0^1 dx [F_2(x)] \simeq \int_0^1 dx [x(q(x) + \bar{q}(x))] \quad (1.24)$$

¹An approach to solving quantum chromodynamics situations by simulating quarks and gluons on a set of lattice points in spacetime and modeling their well-defined interactions.

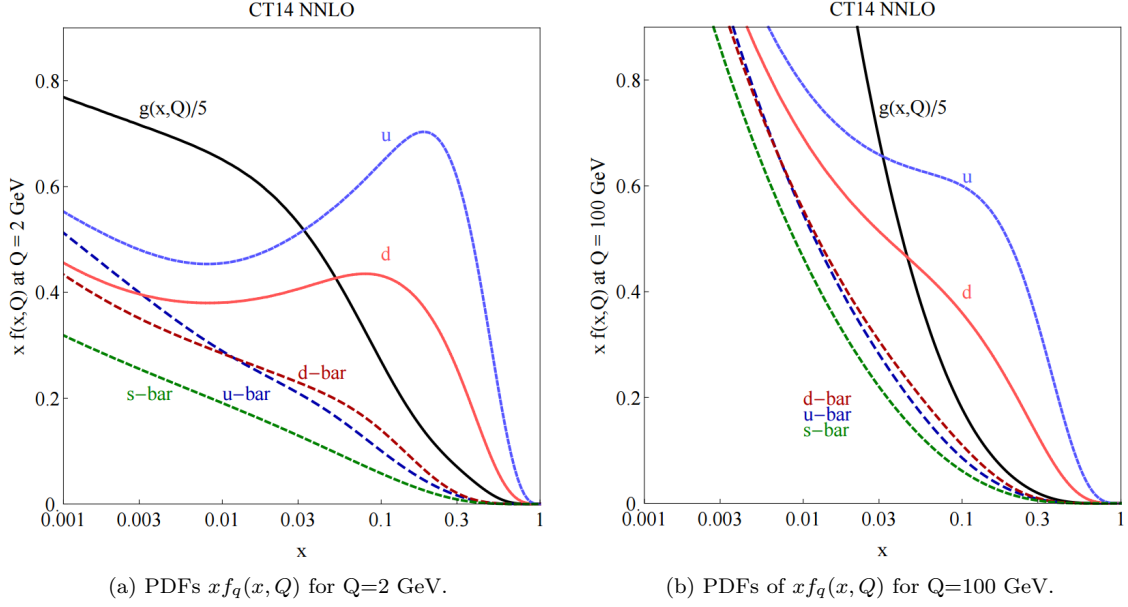


Figure 1.4: Parton distribution functions for quarks and gluons and their x and Q dependence as calculated to NNLO by the CTEQ-TEA (CT) global analysis group[10].

showed that x_{tot} for quarks is 0.512 ± 0.018 [11]. The rest was found to be attributed to the gluons, though they could not be detected directly by an electron beam in DIS since they carry no electric charge. The momentum fraction from just anti-quarks comes from clever use of another structure function, F_3 .

$$F_3^{\nu N}(x) = q(x) - \bar{q}(x) - 2s(x) + 2c(x) \quad (1.25)$$

$$F_3^{\bar{\nu} N}(x) = q(x) - \bar{q}(x) + 2s(x) - 2c(x) \quad (1.26)$$

$$F_3(x) = \frac{1}{2}[F_3^{\nu N}(x) + F_3^{\bar{\nu} N}(x)] = q(x) - \bar{q}(x) = u_v(x) + d_v(x) \quad (1.27)$$

Here, the v subscript denotes valence quark distributions, $q/\bar{q}(x)$ are

$$q(x) = u(x) + d(x) + s(x) + c(x) \quad (1.28)$$

$$\bar{q}(x) = \bar{u}(x) + \bar{d}(x) + \bar{s}(x) + \bar{c}(x) \quad (1.29)$$

and the individual $F_3^{\nu/\bar{\nu} N}$ structure functions are measured from νN and $\bar{\nu} N$ scattering and then combined. With F_3 , there is the Gross-Llewellyn Smith (GSL) sum rule[12]

$$\int_0^1 F_3(x) dx = 3 \quad (1.30)$$

which is to say that there are three valence quarks. The new information gained here with F_3 though is the measurement of $\int dx[xF_3(x)] = 0.341 \pm 0.036$, which gives the momentum fraction carried by the valence quarks. Put this together with the results of Eq. 1.24 and one can derive the conclusion that the anti-quarks in the nucleon possess between 13% and 17% of the nucleon's momentum[13].

The last interesting sum rule to discuss is the Gottfried sum rule which is focused on flavor asymmetries between the proton and its isospin partner, the neutron.

$$S_G = \int_0^1 dx \left[\frac{F_2^p - F_2^n}{x} \right] \quad (1.31)$$

$$= \int_0^1 dx \left[\sum_f e_f^2 [q_f^p(x) + \bar{q}_f^p(x) - q_f^n(x) - \bar{q}_f^n(x)] \right] \quad (1.32)$$

We assume charge symmetry insofar as $u^p(x) = d^n(x)$ and $\bar{u}^n(x) = \bar{d}^p(x)$, yielding

$$S_G = \int_0^1 dx \left[\frac{1}{3}(u(x) + \bar{u}(x) - d(x) - \bar{d}(x)) \right] \quad (1.33)$$

where these PDFs refer to the distributions in the proton. By some reorganization and grouping, this can be re-expressed in a more familiar way.

$$S_G = \int_0^1 dx \left[\frac{1}{3}[u(x) - \bar{u}(x)] \right] - \int_0^1 dx \left[\frac{1}{3}[d(x) - \bar{d}(x)] \right] - \int_0^1 dx \left[\frac{2}{3}[\bar{d}(x) - \bar{u}(x)] \right] \quad (1.34)$$

The first two integrals here are the definitions of the valence quark sum rule mentioned in Eq. 1.16, and this is known to be 2 and 1 for the u and d valence quarks of the proton, respectively. As a result, we arrive at the common representation of the Gottfried Sum Rule (GSR) [14].

$$S_G = \frac{1}{3} - \int_0^1 dx \left[\frac{2}{3}[\bar{d}(x) - \bar{u}(x)] \right] \quad (1.35)$$

If one were to assume that, within a proton, $\int dx[\bar{u}(x)] = \int dx[\bar{d}(x)]$, then this sum reduces to $\frac{1}{3}$.

1.2.2 Drell-Yan Tests of the QPM

The Drell-Yan process can be described within the frameworks of QPM and QCD, so it provides a good testing ground for these models. Even further, the very clean leptonic DY signature offers a good vantage point by which to confirm the parton picture experimentally. The testing process began with establishing the validity of the quark parton model via Drell-Yan. Once this is done, the following step would be to study

the deviation between measurements and what is expected. Here we step through some tests of the QPM that Drell-Yan was applied to in order to validate.

Point-like quarks

If the quarks in the QPM are point-like and without any substructure, then the DY cross section (expressed in Eq. 1.14) across Q^2 should exhibit scaling behavior, i.e. the parton distributions measured should be Q^2 -invariant. The experiments E-288 and E-605 at Fermilab compared the cross section yields across four different beam energies corresponding to \sqrt{s} values between 19.4 and 38.8 GeV. For each setting, the cross section was binned in x_1 and x_2 and compared across different beam settings (this is the same as binning in $Q^2 = M_{\gamma^*}^2 = sx_1x_2$).

If the what was being seen was a quark, and the quark was point-like, then the cross section would be Q^2 -independent. If there was substructure to the quark, then at higher Q^2 , sub-quarks would strike or annihilate with other sub-quarks, and one would observe the cross section increasing with Q^2 as new interaction phase space becomes available. The measurements of the experiments exhibited no Q^2 -dependence [15, 16] and lent support to the notion that what was being struck was indeed a point-like particle with no sub-structure.

Quark charge

One of the predictions of QPM was that the up (down) quarks possessed $e_u = +2/3$ ($e_d = -1/3$) times the elementary charge. The use of π^+ and π^- beams to induce Drell-Yan interactions was used in the Fermilab E-444 and the OMEGA experiments to measure the Drell-Yan cross section ratios of π^+ to π^- , which should approach a value of $\frac{e_d^2}{e_u^2}$. The results indicated a value of $\frac{1}{4}$ [17, 18], which is consistent with at least the magnitudes of the charges of the quarks.

Spin-1/2 quarks

If the quarks that annihilate are in fact spin- $\frac{1}{2}$ particles, then the virtual photon that the quarks annihilate to must be predominantly transversely polarized. This, in turn, should define the angular distribution of the leptons that are produced. As such, a transversely polarized photon should result in a dilepton distribution that goes as $dN/d\theta = 1 + \cos^2 \theta$ in the rest frame [19].

Several experiments sought to test this prediction, but the values of the measurement where $\cos \theta \approx \pm 1$ had a lack of data, allowing for a sizable amount of variation. Regardless, the data collected by CIP, NA3, ABCS, CHFMP, Omega, and Goliath all fit the angular distribution to the function $f(\theta) = 1 + \lambda \cos^2 \theta$ and arrived at results for λ consistent with $\lambda = 1$, though each had relatively sizable uncertainties [19]. This

is in agreement with what the simple Drell-Yan model predicts.

1.2.3 QCD-Improved Drell-Yan

According to the quark-parton picture, the simple Drell-Yan model predicts a certain magnitude of the cross-section and a certain measurement of p_T . Deviations from these predictions indicated a need to accommodate for higher-order effects that arise due to QCD effects.

It was determined over the course of many experiments that the shape of the cross section was predicted very well while the scale of the measurement was off by a non-negligible amount. The discrepancy was studied thoroughly in the form of an arbitrary correction constant, K , known as the K -factor:

$$\frac{d\sigma^{DY}}{dQ^2} \Big|_{exp} = K \cdot \frac{d\sigma^{DY}}{dQ^2} \Big|_{LO} \quad (1.36)$$

where LO stands for leading order (NLO, NNLO stand for next-to- and next-to-next-to-leading order). Relatively wide ranges of values of K were measured [17], from 1.6 to 3.1, with the combined value converging to around 2.

The expected measurement of p_T was also surprising, as the simple Drell-Yan predicted a simple, clean annihilation that carried no other p_T than the intrinsic p_T caused by nucleon confinement and Fermi motion, which should be on the order of 400 MeV (this value depends on the proton dimensions only). It turned out that not only did the measurement of p_T not match this constant value, but the Omega collaboration showed that the value had a dependence on the \sqrt{s} in the form of

$$\hat{p}_T^p = 0.45 + 0.025\sqrt{s} \quad ; \quad \hat{p}_T^\pi = 0.54 + 0.029\sqrt{s} \quad (1.37)$$

for proton-induced pion-induced DY, respectively [19].

These discrepancies can be understood when one considers gluon interactions that happen on very short time scales. When a quark participates in a DY interaction, it may emit a gluon before encountering the other quark. The effects of time dilation in the moving nucleon essentially result in this (*soft*) gluon happening “long before” the hard interaction occurs. This will not alter the factorized structure of the Drell-Yan cross section. It is, however, the case that gluon emission and absorption can occur at timescales on the order $\Delta t \sim 1/Q$ before the actual annihilation (i.e. *hard* gluons).

Several different hard gluon interactions that do affect the Drell-Yan cross section can be seen in Figure 1.5. In part (a), the LO DY vertex is shown, and the rest are the correction graphs that must be accounted for. Diagram (b) shows a vertex correction term, (c) and (d) show gluon emission in the final

state, and (e) and (f) show QCD “Compton” scattering [19]. Of these processes, diagrams (c-f) show a virtual photon recoiling off of an outgoing gluon, and these together easily account for the additional p_T that is produced.




Figure 1.5: Diagrams of (a) LO Drell-Yan, (b) vertex correction, (c-d) gluon emission, (e-f) “Compton” scattering [19].

The K factor has been found to be generally of the form $K = \exp[(\alpha_s/2\pi)(4\pi^2/3)]$, and has been calculated in considerable detail [20] to order $\mathcal{O}(\alpha_s^2)$ and are considered to be under control. Most importantly, it was found at the E-866 experiment that the differential cross section ratios in the next-to-leading order are almost identical to those in the leading order nuclear dependence (at 800 GeV) [21]. Further, similar results are given for the lower energy proton bombarding deuterium and tungsten at the Fermilab Main Injector (FMI, 120 GeV proton beam) [22] which is used at SeaQuest. Therefore, QCD corrections for nuclear targets (compared to free nucleons) can be considered negligible. As such, when taking ratios of cross sections of different nuclear materials, the K -factors are considered to cancel each other out.

1.3 The Nuclear EMC Effect

Approximately ten years before the NMC experiment’s discovery of the $\bar{d}(x) \neq \bar{u}(x)$ discovery, the European Muon Collaboration (EMC) at CERN, in 1983, measured the ratios of the structure functions of iron and deuterium over a large kinematic range using the DIS of muons [23]. The result has been repeatedly scrutinized, studied, and confirmed. It was discovered that the structure function of a nucleon bound in a nucleus is fundamentally modified from that of a free nucleon. The nucleus was naively understood as a simple convolution of many nucleons, and that the partonic structure functions of bound and free nucleons should be identical. This assumption was found to not hold against actual measurements, and the observed behavior holds true across all heavy nuclei. This effect can be interpreted as a *modification of quark and gluon distributions in bound nucleons by the nuclear environment*. In this section, the history and collective behavior known as “The EMC Effect” will be reviewed, a quantitative A-dependent measure of the effect will

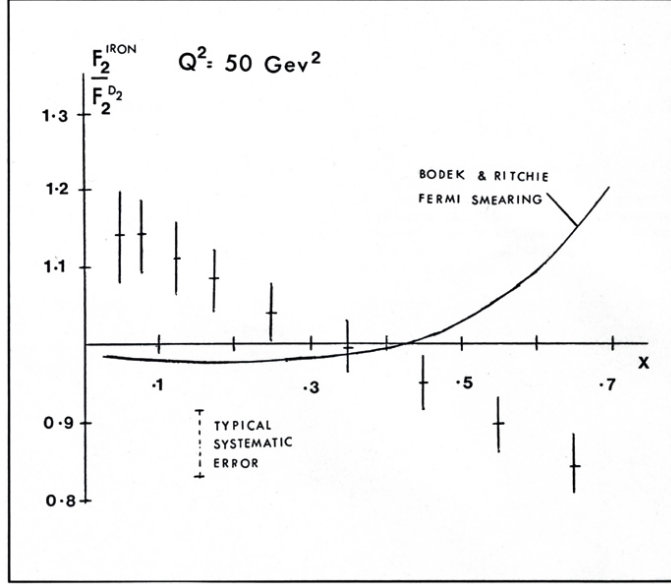


Figure 1.6: The original EMC ratio measurement by the EMC collaboration [23]. Overlaid is the naive, expected ratio taking into account only the Fermi motion of the nucleons in the heavy nucleus.

be defined, and models that attempt to explain the phenomenon will be discussed. Lastly, the contributions of SeaQuest and previous Drell-Yan experiments to the study of the EMC Effect will be explored.

1.3.1 Historical Overview

Unpolarized inclusive lepton scattering is a process that uses an energized lepton to investigate nucleon structure, and it specifically depends on two independent variables: Q^2 and Bjorken- x (sometimes referred to as x_B , but here left as x).

$$Q^2 = -q^2 \quad ; \quad x = \frac{Q^2}{2m_p\omega} \quad (1.38)$$

where q is the transferred four-momentum, x_B is the same Bjorken- x scaling variable as is in the PDFs in Section 1.2.1 above, m_p is the proton mass, and ω is the transferred energy in the proton rest frame. At energies of $Q^2 > 2(\text{GeV}/c)^2$, the parton substructure is revealed and the inelastic structure function, F_2 is able to be measured. The EMC group extracted the ratio $F_2^{Fe}(x, Q^2)/F_2^d(x, Q^2)$, where these are the average bound nucleon structure functions in iron and deuterium. This was done by measuring the DIS *per nucleon cross section ratio*, which can be expressed as [7]:

$$\frac{\sigma^{A_1}/A_1}{\sigma^{A_2}/A_2} = \frac{F_2^{A_1}(x, Q^2)}{F_2^{A_2}(x, Q^2)} \cdot \frac{\left[1 + 2 \frac{1+\omega^2/Q^2}{R_{A_1}-1} \tan^2 \frac{\theta}{2}\right]}{\left[1 + 2 \frac{1+\omega^2/Q^2}{R_{A_2}-1} \tan^2 \frac{\theta}{2}\right]} \approx \frac{F_2^{A_1}(x, Q^2)}{F_2^{A_2}(x, Q^2)} \quad (1.39)$$

where θ is the lepton scattering angle and $R_A = \sigma_L^A/\sigma_T^A$ is the ratio of the longitudinal to transverse cross section for nucleus A. The A-dependence of R_A will be discussed later, but for now, let us assume that R does not exhibit a nuclear dependence, and the per nucleon cross section ratio therefore approximates nicely to the ratio of the structure functions. For $0.4 \gtrsim x \lesssim 0.6$ where nucleon Fermi motion effects can be neglected, the expectation was to measure unity, which would indicate that the structure function of a deeply bound nucleon and loosely bound nucleons are identical. With this goal in mind, they chose ^{56}Fe and deuterium, as iron is the most tightly bound nucleus at $E_{\text{bind}}^{Fe} = 8.8 \text{ MeV/nucleon}$ and deuterium is as close to being a free nucleon as possible while still being a symmetric nuclei ($N = Z = A/2$) with a binding energy of only $E_{\text{bind}}^d = 2.2 \text{ MeV/nucleon}$. The reason to perform this experiment is that, if this measurement showed unity in the regions not affected by Fermi motion, then heavy nuclei could be used to gain a luminosity boost when studying the free nucleon structure functions. Instead, they observed the measured cross-section ratio being 1 at $x \approx 0.3$ and then steadily declining to a value of about 0.8 at $x \approx 0.7$ (Figure 1.6)

At the time of the measurement was presented, the result astounded the physics community. They could hardly believe that high momentum transfers that were many times larger than the nuclear binding energies that the quark distributions should be altered by the nuclear mean field. This result was confirmed only three months later by the SLAC-MIT-Rochester group that re-examined its data for DIS on steel [24] and found the same behavior, as can be seen in Figure 1.7. This opened the gates to an enormous amount of activity in both experiment and theory to fully characterize the phenomenon, resulting at present more than 1100 citations of the original EMC publication [23].

The compendium of experimental data and theory papers have been discussed in great detail in the excellent review papers by Geesaman [25] in 1995 and Norton [26] in 2003. As such, I will only summarize the main theoretical ideas and the most prominent interpretation of the effect. A visual summary of EMC data can be found in Figure 1.8 where data from EMC [23], SLAC [27, 28], BCDMS [29], NMC [30], and HERMES [31] are plotted for nine different targets.

1.3.2 Fermi Motion




Figure 1.7: SLAC E-87 was quick to report its own findings months later confirming EMC's observation [24].




Figure 1.8: A compilation of data regarding the measurement of F_2^A/F_2^D for nine different heavy targets with an overlaid fit by Chen et al. [32].




Figure 1.9: The Fermi momenta K_F (in GeV) for various nuclei of atomic weight A [33].

example, the nucleons in a tightly bound nucleus (like iron) must have a more precise spatial localization (smaller σ_x) than a nucleon in a more loosely bound nucleus (such as the deuteron). As a result, the magnitude of momentum variations will be greater in the case of iron than in the deuteron (an effect sometimes called “Fermi smearing”). The magnitude of the Fermi momenta of several nuclei can be found in Figure 1.9.

In the extreme case of a free nucleon (which deuterium is chosen to approximate as best as possible), there exists no Fermi motion, as the nucleon is unbound. As a result, the F_2 structure function must vanish as $x \rightarrow 1$ as it becomes very unlikely that a single quark will possess a large fraction of the momentum of the free nucleon without the help of quantum mechanical fluctuations and Fermi smearing. In the case of the deuteron, though, there is still some Fermi motion, as it still consists of two bound nucleons. Bodek and Ritchie were among the first to lay out a model for comparing this effect and how it would differ between different nuclei [33]. The model of how this would affect the structure function ratio is seen on the original EMC plot in Figure 1.6, given no other nuclear effects (such as the EMC Effect). It is important to note that this effect was known and predicted before the original EMC measurement, however it predicted an EMC ratio measurement of unity, which would then rise at higher- x .

The regime where this effect is dominant is far beyond the kinematic acceptance of SeaQuest to measure. There is still good reason to believe that, whatever the cause of the EMC effect, it likely overlays the effect of Fermi motion, which, depending on the model, can affect the ratio measurement above $x \approx 0.5$. Even with this estimate, there are very few SeaQuest statistics at this x -range.

1.3.3 Universal x -dependence

The SLAC E139 experiment provides very precise information for the high- x range ($x > 0.2$) for many nuclei. The high-energy electron DIS data from SLAC E61 [34] and HERMES [31] experiments along with the muon DIS data from NMC [30] supplements this with the very low- x range. Finally, the more recent

One of the more striking features of the $F_2^A(x)/F_2^d(x)$ is the sharp rise in the ratio value at high- x ($0.8 \lesssim x \lesssim 1$). This feature arises as a result of the quantum mechanical principle of Fermi motion, which is a result of the confinement of nucleons within the nucleus. If one considers the uncertainty principle inherent in all wave-like systems, $\sigma_x \sigma_p \geq \frac{\hbar}{2}$, there is a direct effect on the momentum distribution of nucleons when one considers the

spatial localization of bound nucleons in a nucleus. For

example, the nucleons in a tightly bound nucleus (like iron) must have a more precise spatial localization

(smaller σ_x) than a nucleon in a more loosely bound nucleus (such as the deuteron). As a result, the mag-

nitude of momentum variations will be greater in the case of iron than in the deuteron (an effect sometimes

called “Fermi smearing”). The magnitude of the Fermi momenta of several nuclei can be found in Figure 1.9.

In the extreme case of a free nucleon (which deuterium is chosen to approximate as best as possible),

there exists no Fermi motion, as the nucleon is unbound. As a result, the F_2 structure function must vanish

as $x \rightarrow 1$ as it becomes very unlikely that a single quark will possess a large fraction of the momentum of

the free nucleon without the help of quantum mechanical fluctuations and Fermi smearing. In the case of

the deuteron, though, there is still some Fermi motion, as it still consists of two bound nucleons. Bodek and

Ritchie were among the first to lay out a model for comparing this effect and how it would differ between

different nuclei [33]. The model of how this would affect the structure function ratio is seen on the original

EMC plot in Figure 1.6, given no other nuclear effects (such as the EMC Effect). It is important to note

that this effect was known and predicted before the original EMC measurement, however it predicted an

EMC ratio measurement of unity, which would then rise at higher- x .

The regime where this effect is dominant is far beyond the kinematic acceptance of SeaQuest to measure.

There is still good reason to believe that, whatever the cause of the EMC effect, it likely overlays the effect

of Fermi motion, which, depending on the model, can affect the ratio measurement above $x \approx 0.5$. Even

with this estimate, there are very few SeaQuest statistics at this x -range.




Figure 1.10: The ratio of $\sigma^{C(N)}/\sigma^d$ covering nearly the entire $x \in [0, 1]$ interval [36]. The four qualitative regions are denoted by roughly chosen boundaries x_1, x_2, x_3 (not to be confused with the beam and target x_1 and x_2 used in the rest of this paper).

data from JLAB-E03103 [35] that uses a relatively low energy beam (5.8 GeV) provides data for the high- x Fermi motion area ($0.85 < x < 0.95$). These together help to paint a full characteristic picture of the x -dependence of this σ^A/σ^d ratio, which has been found to be universal for all $A > 2$. A visualization of some of this data for carbon and nitrogen the measurement up into four qualitative regions of this ratio can be seen in Figure 1.10.

- $0 < x \lesssim 0.06$: The “shadowing” region where the ratio is less than unity and decreasing as $x \rightarrow 0$. The dominant contribution to the cross section in this kinematic regime is from sea quarks.
- $0.06 \lesssim x \lesssim 0.3$: The (perhaps poorly named) “anti-shadowing” region exhibits a very slight rise of the ratio above unity. In general, this range of values represents only a transition region between its two bookending effects.
- $0.3 \lesssim x \lesssim 0.8$: The ratio becomes significantly smaller than unity, decreasing with increasing x . This is of key focus for this paper, and this region is what is generally referred to as the “*EMC Effect*” [25] as it is the only part of the phenomenon that is not fully understood. As far as DIS goes, the sea quark contribution is essentially negligible, given the current understanding of quark PDF’s, and the behavior is dominated by valence quark distributions.
- $0.8 \lesssim x < 1$: The ratio increases rapidly with increasing x . This is dominated by a kinematic effect of the free nucleon cross section vanishing as $x \rightarrow 1$. It is also due to the effect of Fermi motion of the bound nucleons in the nucleus.

1.3.4 Nuclear Dependence of R

It was stated in section 1.3.1 that the per nucleon cross section ratio only reduces to the ratio of structure functions in the case that $R = \sigma_L/\sigma_T$ is A-independent. This particular measurement was the focus of experiments SLAC-E140 [9, 37], NMC [38], and HERMES [31]. A re-analysis of all SLAC data with improved corrections [37] showed that the difference in R between iron and deuterium is consistent with zero ($R_{Fe} - R_d = 0.001 \pm 0.018(\text{stat}) \pm 0.016(\text{sys.})$), and that there are no significantly different spin-0 constituents or higher twist effects in nuclei as compared to free nucleons. Later on, results from HERMES corroborated these findings. The ratio of R_A/R_d were derived from the NMC and SLAC results for $\Delta R = R_A - R_d$ and reported an averaged value of $R_A/R_d = 0.99 \pm 0.03$, providing further agreement on the A-independence of R .

A recent study by Guzey et al. [39] looked to data from SLAC E139, BCDMS, and EMC to examine the impact of R on the extraction of F_2^A/F_2^d and F_1^A/F_1^d from σ^A/σ^d data. They demonstrate the presence of a small but non-zero difference between R for nuclei and free nucleons. While this does appear to effect the extraction of the nuclear enhancement in the ratio of the F_1^A/F_1^d , the extraction of F_2^A/F_2^d appears to remain unaffected. As such, this thesis assumes that it is safe to assume that $\sigma^A/\sigma^d \approx F_2^A/F_2^d$.

1.3.5 Q^2 Dependence of EMC Effect

It is important when investigating structure functions and quark/gluon distributions to pay attention to the Q^2 sensitivity of the measurements. Nuclear structure functions can be directly interpreted as parton




Figure 1.11: Measurements of the change in the σ^{Ca}/σ^d ratio with respect to $\ln Q^2$ over a broad range of x [25].

distributions only if Q^2 is large enough to precisely define the light-cone direction and reduce higher-twist contributions. The first condition places a restriction on x^2/Q^2 and is usually well satisfied, even at relatively low values of x and Q^2 . However, the value of the higher-twist terms is under little theoretical control, and these terms can have specific nuclear dependences.

Experimentally, the verdict was very clear from the agreement between muon and electron DIS data that any Q^2 dependence must be very small since, at the same value of x , their average Q^2 differs by more than an order of magnitude. The lack of Q^2 dependence can be seen in Figure 1.11 which shows the derivative of DIS cross section with respect to $\ln Q^2$ over a broad range of x , combining NMC and E665 data for calcium [25]. *This absence of any observable dependence of the ratio on Q^2 justifies the usual practice of combining all the Q^2 data at a fixed x -value to display the x -dependence of the ratio.* The only measurement that seems to refute this conclusion is a measurement of σ^{Sn}/σ^C by NMC which observes a statistically significant, yet small Q^2 dependence at small values of x ($x \lesssim 0.06$) [38]. Due to a large amount of evidence to the contrary along with the fact that SeaQuest is only sensitive to $x \gtrsim 0.08$, this paper will move forward with the understanding that there is no Q^2 dependence in this ratio measurement. As such, this paper will frequently abbreviate such expressions:

$$\frac{F_2^A(x, Q^2)}{F_2^d(x, Q^2)} \rightarrow \frac{F_2^A(x)}{F_2^d(x)} \quad ; \quad \frac{2 \sigma^A(x, Q^2)}{A \sigma^d(x, Q^2)} \rightarrow \frac{2 \sigma^A(x)}{A \sigma^d(x)} \quad (1.40)$$

when in reality it is only the ratio that is Q^2 -independent.

1.3.6 Shadowing

Even though the “shadowing” region of the EMC effect ($x \lesssim 0.06$) is outside of the kinematic acceptance of the SeaQuest experiment ($x \gtrsim 0.08$), it is certainly worth briefly discussing this interpretation. This term “shadowing” has been introduced to explain the reduction of the nuclear cross sections in photoproduction, but in this particular case, has also been used in discussing very low- x modifications to the nuclear cross section in DIS.

The generally accepted explanation for this modification is that the parton distributions in bound nucleons actually *remain unchanged* compared to those in free nucleons. In the rest frame of the *nucleus*, the nuclear effect is attributed to the modification of the interaction of the virtual photon with the atomic nucleus by fluctuations of the virtual photon into quark-antiquark pairs. This pair then interacts with the nucleus via the strong interaction instead of electromagnetically. As a result of this increase in strength, the interaction no longer proceeds incoherently with all the rest of the nucleons, but *preferentially with those*

Nucleon	dR_{EMC}/dx [35]	dR_{EMC}/dx [28]	dR_{EMC}/dx (comb.)
2H			-0.10 ± 0.05 [41]
3He	-0.070 ± 0.029		-0.070 ± 0.029
4He	-0.199 ± 0.029	-0.191 ± 0.061	-0.191 ± 0.026
9Be	-0.271 ± 0.029	-0.207 ± 0.037	-0.243 ± 0.023
${}^{12}C$	-0.280 ± 0.029	-0.318 ± 0.040	-0.292 ± 0.023
${}^{27}Al$		-0.325 ± 0.034	-0.325 ± 0.034
${}^{40}Ca$		-0.350 ± 0.047	-0.350 ± 0.047
${}^{56}Fe$		-0.388 ± 0.032	-0.388 ± 0.032
${}^{108}Ag$		-0.496 ± 0.051	-0.496 ± 0.051
${}^{197}Au$		-0.409 ± 0.039	-0.409 ± 0.039

Table 1.2: Measured EMC slopes dR_{EMC}/dx for $0.35 \leq x \leq 0.7$ [42].

at the front surface. The other nucleons in the “shadow” of the nucleons in front are much less likely to interact, and thereby do not (or much less) contribute to the cross section. As a result, when calculating the per nucleon cross section ratio $\frac{\sigma^A/A}{\sigma^d/2}$ and the deuteron essentially has the same cross section as the heavy nucleus (in this specific regard), then the ratio value will be suppressed.

Quantitatively, this behavior will occur when the $q\bar{q}$ pair fluctuation exists for longer than the mean free path of the pair in a nucleus (distance between nucleons). The fluctuation length behaves as $\Delta d \approx 1/Mx$ where M , the effective mass of the $q\bar{q}$ pair is approximately 1.1 GeV [40], yielding a fluctuation distance of $\Delta d \approx 0.2 \text{ fm}/x$. This becomes larger than the mean free path $L \simeq 2.5 \text{ fm}$ when $x \lesssim 0.08$.

1.3.7 A-dependence of EMC Effect Slope




Figure 1.12: The slopes of the EMC ratio for $0.35 < x < 0.7$ vs the average nuclear density scaled by $(A - 1)/A$ [35].

Precise measurements of $R_{EMC} = (F_2^A/A)/(F_2^d/2)$ for various nuclei, from 3He to ${}^{197}Au$ have been analyzed, and the intermediate-to-high x region ($0.35 < x < 0.7$) shows a declining behavior that can be fit with a line. The slope of this line has become the focus of study when attempting to interpret the EMC effect, as this slope (dR_{EMC}/dx) exhibits a significant A-dependence. The magnitudes of the slopes appear to increase with increasing A . The table of values of dR_{EMC}/dx for various nuclei can be found in Table 1.2.

The initial attempt at understanding this was to interpret it as some effect of nuclear density. Since the effect of nuclear density on a given nucleus is due to the presence

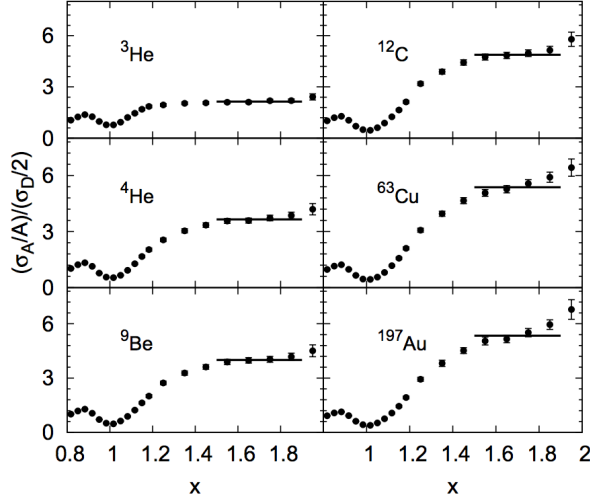


Figure 1.13: Per nucleon cross section ratios vs x at $0.8 < x < 2$ as measured by JLAB-E02019 with a line fit to the “plateau” regions at $1.5 < x < 1.9$ [43].




Figure 1.14: Per nucleon cross section ratios vs x at $0.75 < x < 3$ as measured by CLAS with a line fit to the “plateau” regions at $2.25 < x < 2.8$ [44].

of the other $(A - 1)$ nucleons, the slope $|dR_{EMC}/dx|$ was plotted against the average nuclear density scaled down by a factor of $(A - 1)/A$ in order to remove the struck nucleons contribution to the average nuclear density. The result for light nuclei was shown by Seely *et al.* to have an interesting result, as is seen in Figure 1.12. The most notable result was a definite overall trend with ${}^9\text{Be}$ representing a significant outlier. It has been theorized that ${}^9\text{Be}$ is an outlier due to its nuclear structure resembling two α -particles (${}^4\text{He}$ nuclei) with an extra neutron. This would mean that perhaps it is not the average nuclear density, but the *local* nuclear density or density about the struck nucleon. There is still a missing piece needed to represent this density.

1.3.8 $x > 1$ Plateaus and Short Range Correlations

What has recently spurred renewed interest in the EMC effect and its possible true origin is a series of measurements by JLAB of electron DIS cross section ratios where $x > 1$. In this kinematic space, the cross section from a free nucleon vanishes, but persists for the deuteron. It may not at first make sense that measurements can be made in this high x , but it is accessible. Such a behavior is interpreted as the manifestation of short-range nucleon-nucleon correlations (three-nucleon correlations for $x > 2$). These short-range correlations happen when two (or three) nucleons happen to get very close to each other and experience a large repulsive Coulomb force, resulting in the nucleons moving apart from each other with high *individual* momentum, but low *combined* momentum.

Nucleon	$a_2(A/{}^3He)$ [48]	$a_2(A/d)$ [46]	$a_2(A/d)$ [43]
2H	0.508 ± 0.025	1	1
3He	1	1.97 ± 0.10	2.13 ± 0.04
4He	1.93 ± 0.14	3.80 ± 0.34	3.60 ± 0.10
9Be			3.91 ± 0.12
${}^{12}C$	2.41 ± 0.17	4.75 ± 0.41	4.75 ± 0.16
${}^{56}Fe$	2.83 ± 0.18	5.58 ± 0.45	
${}^{63}Cu$			5.21 ± 0.20
${}^{197}Au$			5.16 ± 0.22

Table 1.3: Measured SRC scale factors, as measured by CLAS [48] and JLAB E02-019 [46, 43].

The JLAB-E02019 experiment [43] measured cross section ratios for 3He , 4He , Be , C , Cu and Au to deuterium. One can see in the results shown in Figure 1.13 that the cross section ratios of ${}^4He/D$, C/D and Cu/D rise with x until the value plateaus at $x \approx 1.4 - 1.5$. The key observation with these plateaus is that the ratio value at this plateau increases with A . Such behavior was already observed with less accuracy by an experiment at SLAC [45], and a similar pattern can be seen in the CLAS experiment which compared nuclear cross sections to tritium (3He) [44]. Here the measurements extend up to $x = 3$ and an indication of a second plateau is observed for $2.25 < x < 2.8$ (Fig. 1.14).

The variable $a_2(A/d)$ is introduced here as the ratio (in the plateau range) of the per-nucleon inclusive (e, e') cross sections for two nuclei. This value can be interpreted as the *ratio of probabilities to find high-momentum nucleons in those two nuclei* [46]. High-momentum nucleons were recently shown to arise primarily from nucleon-nucleon short range correlations, which is when two nucleons' wave functions overlap briefly, causing them to have high relative momentum but low combined momentum (relative to the ~ 400 MeV Fermi motion) [47]. Thus, this $a_s(A/d)$ ratio is called the ‘‘Short Range Correlation (SRC) Scale Factor’’² The current known values for the SRC scale factors are shown in Table 1.3. The variables $a_2(A/{}^3He)$ and $a_3(A/{}^3He)$ is also defined for first and second plateau ranges, respectively, but are not emphasized here.

It was first posited by Higinbotham *et al.* [49] that the magnitude of the EMC effect in a nucleus A is related to the measured SRC scale factor of the same nucleus if both are taken relative to deuterium. In Figure 1.15, the EMC slopes from Table 1.2 are plotted against the SRC scale factors from Table 1.3 to exhibit a strong correlation indicated by the straight line. It would seem unlikely that this is coincidental, and it can therefore be asserted with some confidence that the EMC effect seen in the valence quarks is largely due to these 2N-SRC [36].

²More specifically, it is called the 2N-SRC scale factor, as there exists 3N-SRC with respect to tritium.




Figure 1.15: The EMC slopes plotted against the 2N-SRC scale factors [51]. The fit is constrained to go through deuterium at (1,0).

There is a future study planned at JLAB to test this 2N-SRC - EMC effect connection, experiment E1211107, which has been approved to run as part of their 12 GeV program. The experiment will study the structure function of high-momentum nucleons in deuterium by performing electron DIS from nucleons in deuterium and identifying nucleons that are knocked out *backwards* towards the direction that the beam came from [50]. Its primary goal will be to simultaneously measure the EMC effect and SRC scaling factors in a wide variety of light and heavy nuclei.

1.4 The Role of Drell-Yan in Studying the Nuclear Quark Sea

There have been many models attempting to explain the A -dependence of such phenomena as the EMC effect. Each of these models provides a prediction regarding the individual sea/valence quark/antiquark distribution functions. In order to test these predictions, one can take stock of the many nuclear probes and assess the capabilities of what each can contribute.

In general, since deep inelastic scattering of electrons and muons is an electromagnetic interaction, it is not innately sensitive to the whether or not the struck quark was a valence quark or a sea quark and can only measure the charge-weighted (E&M+weak) sum of parton distributions. Through neutrino and antineutrino DIS it is possible to discern between valence and sea distributions. It is the case that most models make predictions for valence quark distributions at mid-to-high x that are not in tension with each

other. As such, measuring the valence distributions would not help much in discriminating between different models. Models do, however, differ in their predictions of the sea quark distributions at higher x ($x \gtrsim 0.2$) - but at high momentum fraction, the distributions of sea quarks diminish very quickly (Fig. 1.4). As such, antineutrino DIS has been unable to deliver precise measurements of sea quark distributions at high- x yet.

The Drell-Yan process, $pA \rightarrow l^+l^-$, uses the s -channel version of the same DIS process, but due to the nature of requiring interacting $q\bar{q}$ pairs, it provides a tool that allows one to cleanly and specifically investigate the sea quark distributions. With a forward spectrometer (such as is the case for E-772 and SeaQuest), where one accepts the $x_F > 0.3$ one is able to study the Drell-Yan interaction in the specific case of a valence quark from a free proton (the beam) annihilating with an antiquark from the quark sea in the target. As a result, a study of this process for various nuclei can provide a direct, precise measurement of the antiquark sea distributions that can assist in evaluating the plethora of EMC effect models out there, along with refining constraints on global fits of PDFs.

In this paper, the A-dependence of the Drell-Yan cross sections is studied, which directly measures the A-dependence of the antiquark distributions [52]. We define this ratio, R^{DY} as the ratio of Drell-Yan differential cross-sections, using Eq. 1.11 to get

$$R^{DY} \equiv \frac{d^2\sigma_A}{dx_1 dx_2} \bigg/ \frac{d^2\sigma_d}{dx_1 dx_2} \quad (1.41)$$

$$\approx \frac{\sum_f e_f^2 \cdot q_f^{\text{beam}}(x_1) \bar{q}_f^A(x_2)}{\sum_f e_f^2 \cdot q_f^{\text{beam}}(x_1) \bar{q}_f^d(x_2)}. \quad (1.42)$$

Assuming that the \bar{u} distribution is roughly similar to the \bar{d} distribution in the target, this term is dominated by up-quark contributions by a factor of eight: $e_u^2 = 4e_d^2$, and there are twice as many up-quarks from the beam proton as down quarks [25]. As a result, we can approximate

$$R^{DY} \approx \frac{e_u^2 \cdot q_u^{\text{beam}}(x_1) \bar{q}_u^A(x_2)}{e_u^2 \cdot q_u^{\text{beam}}(x_1) \bar{q}_u^d(x_2)} \quad (1.43)$$

$$= \frac{\bar{q}_u^A(x_2)}{\bar{q}_u^d(x_2)} \quad (1.44)$$

meaning that by measuring R^{DY} for pA to pd collisions with a forward spectrometer, a good measure of the nuclear modification of the anti up quark sea distributions can be achieved.

1.4.1 Previous Measurements

Even though Drell-Yan events for $x_F > 0.3$ do directly provide a measure the sea quark distribution, there is a limited amount of data available. FNAL-E-772 [54] was a 800 GeV fixed target experiment which provided a measure of σ^A/σ^d for C, Ca, Fe , and W as a function of kinematic variable x_2 . Soon to follow was FNAL-E-866 [53], also a 800 GeV proton beam fixed target experiment, which contributed a somewhat similar measurement of the DY cross section ratio for W and Fe with respect to Be . The results which can be seen in Figures 1.17 and 1.16 are able to show only a limited x -range of measurements of $0.05 \leq x_2 \leq 0.25$ for E-772 and $0.02 \leq x_2 \leq 0.1$ (again, the sea quark distribution drops off sharply with increasing x). This data does show that at low- x_2 ($x_2 \lesssim 0.07$), there is indication of shadowing occurring. This makes sense for DY




Figure 1.16: E-866 Drell-Yan σ^A/σ^{Be} for iron and tungsten as a function of x_2 [53].

as well as DIS, even though there is no need the $\gamma \rightarrow q\bar{q}$ fluctuations, as the number of sea antiquarks increases greatly at this low range of x and would thereby provide the same nucleon “shadow” as was described in Section 1.3.6.

For values above this shadowing region, the ratios appear to be consistent with unity. The range of x covered line up with the *transition region* (sometimes called “anti-shadowing”) between the shadowing and EMC effect region. In some DIS measurements, this region can be characterized by having a value in excess of unity. A measurement of unity in the Drell-Yan cross section ratio would seem to imply that whatever this effect might be, it is likely not caused by any nuclear modifications to the sea quark distribution.

Regarding the other models, there was a great expectation that the Drell-Yan cross section would sizeably increase due to virtual pions which mediate the internuclear force. It was surprising that E-866 and E-772 observed no modification of the sea quark distributions in large nuclei, which brings into question the long-standing model of mesons being exchanged between nucleons as the mechanism by which nuclei are held together [55].

To be able to truly draw any conclusions regarding various models, greater precision is needed, along with data at a higher- x range where models are more greatly differentiated. The SeaQuest experiment will attempt to do just this.

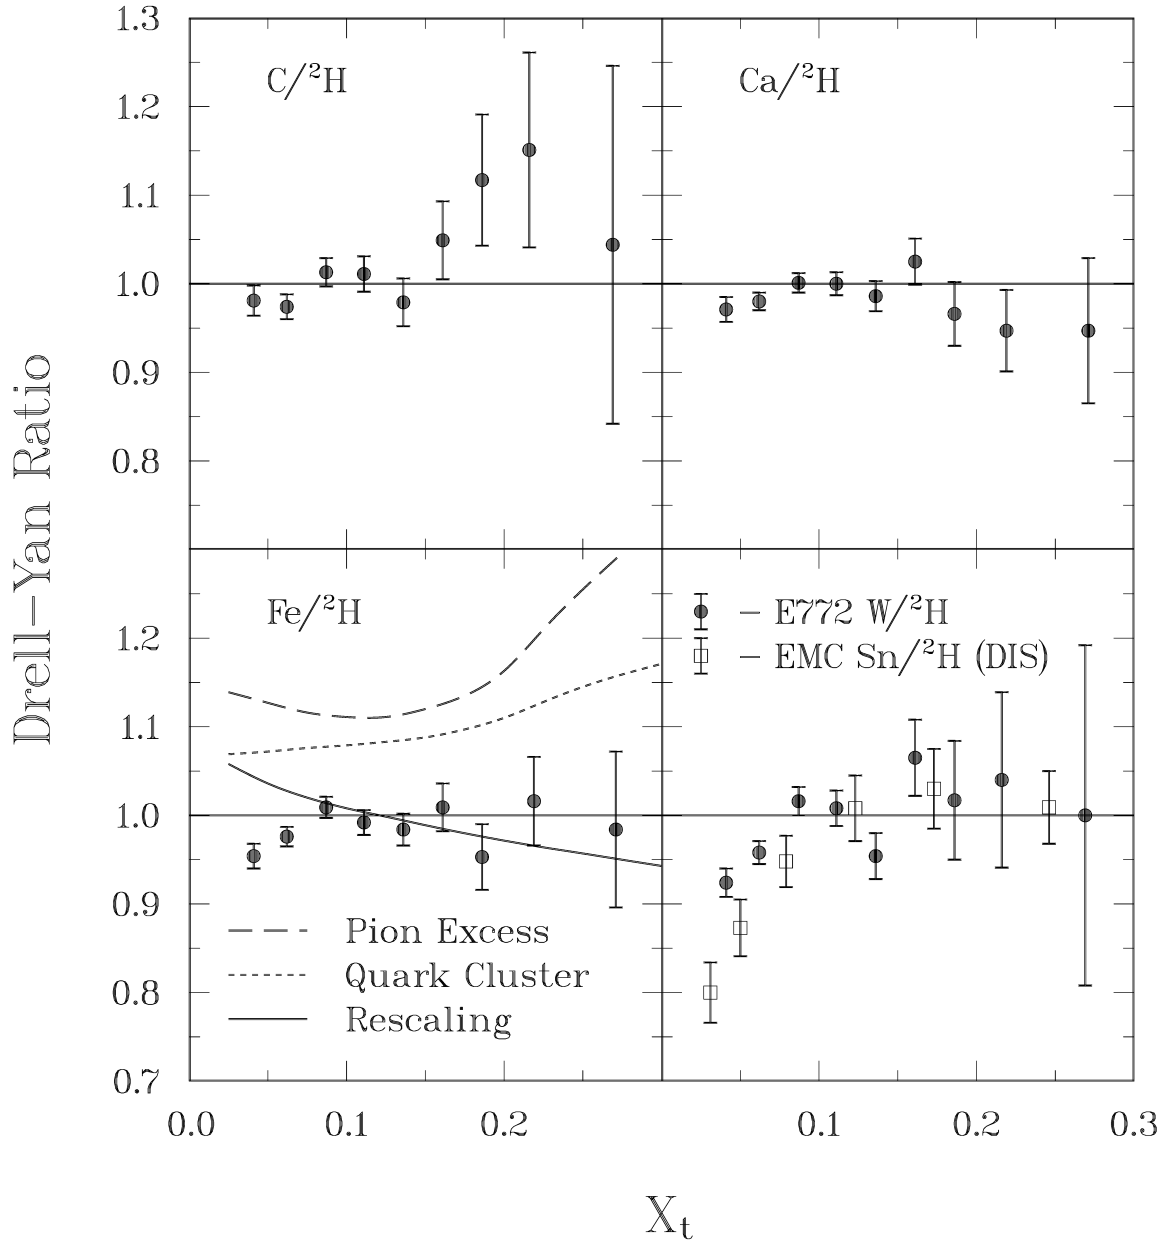


Figure 1.17: E-772 Drell-Yan $\frac{A}{2H}$ ratios for several nuclear targets as a function of x_2 (x_t) [54].

1.4.2 Nuclear Dependence of Drell-Yan at SeaQuest

The SeaQuest Experiment was proposed soon after the conclusion of the E-866 experiment. The main motivation was to explore the $\bar{d}(x)/\bar{u}(x)$ value observed by E-866 to more precisely ascertain its behavior as a function of x . With the ratio value being as large as ~ 1.7 at its peak and with it having a decreasing value at higher x (though with low statistical certainty), no models are able to accommodate such behavior, and so more data is needed.

Sea quark distributions at higher- x the SeaQuest collaboration at Fermilab is designed to detect proton-induced Drell-Yan cross sections on several targets: liquid hydrogen (1H), liquid deuterium (2H), carbon (^{12}C), iron (^{56}Fe), and tungsten (^{184}W). The beam source for the experiment is the Fermilab Main Injector which provides a 120 GeV proton source, yielding a \sqrt{s} of 15.1 GeV. This relatively lower beam energy as compared to 800 GeV ($\sqrt{s} = 38.8$ GeV) for E-772/E-886 provides an opportunity to study parton distributions at larger x . This is due to the fact that, as seen in Eq. 1.11, for a fixed value of x , the Drell-Yan cross section is inversely proportional to the square of the center-of-mass energy, s . Another benefit of utilizing a lower-energy beam is that the primary background, J/Ψ production, which limited the instantaneous luminosity in E-866, scales with s , and is thereby reduced. This allows SeaQuest the usage of a more intense proton beam, which is very useful when it comes to measuring a rare nuclear process like Drell-Yan.

Even though the emphasis of SeaQuest is on the \bar{d}/\bar{u} measurement, the periodic data taking on deuterium and the heavy nuclear targets will allow for a measurement of $R^{DY} \approx \bar{u}^A(x)/\bar{u}^D(x)$ for a large range of A and x , providing a more thorough glimpse into the modification of sea quark PDFs in the presence of a nuclear medium. The general kinematic coverages and correlations for SeaQuest are shown in Figure 1.18, as generated by Geant4 Monte Carlo simulation.

At the time of this writing, SeaQuest is the only fixed target proton-nucleon Drell-Yan experiment currently taking data in the world.

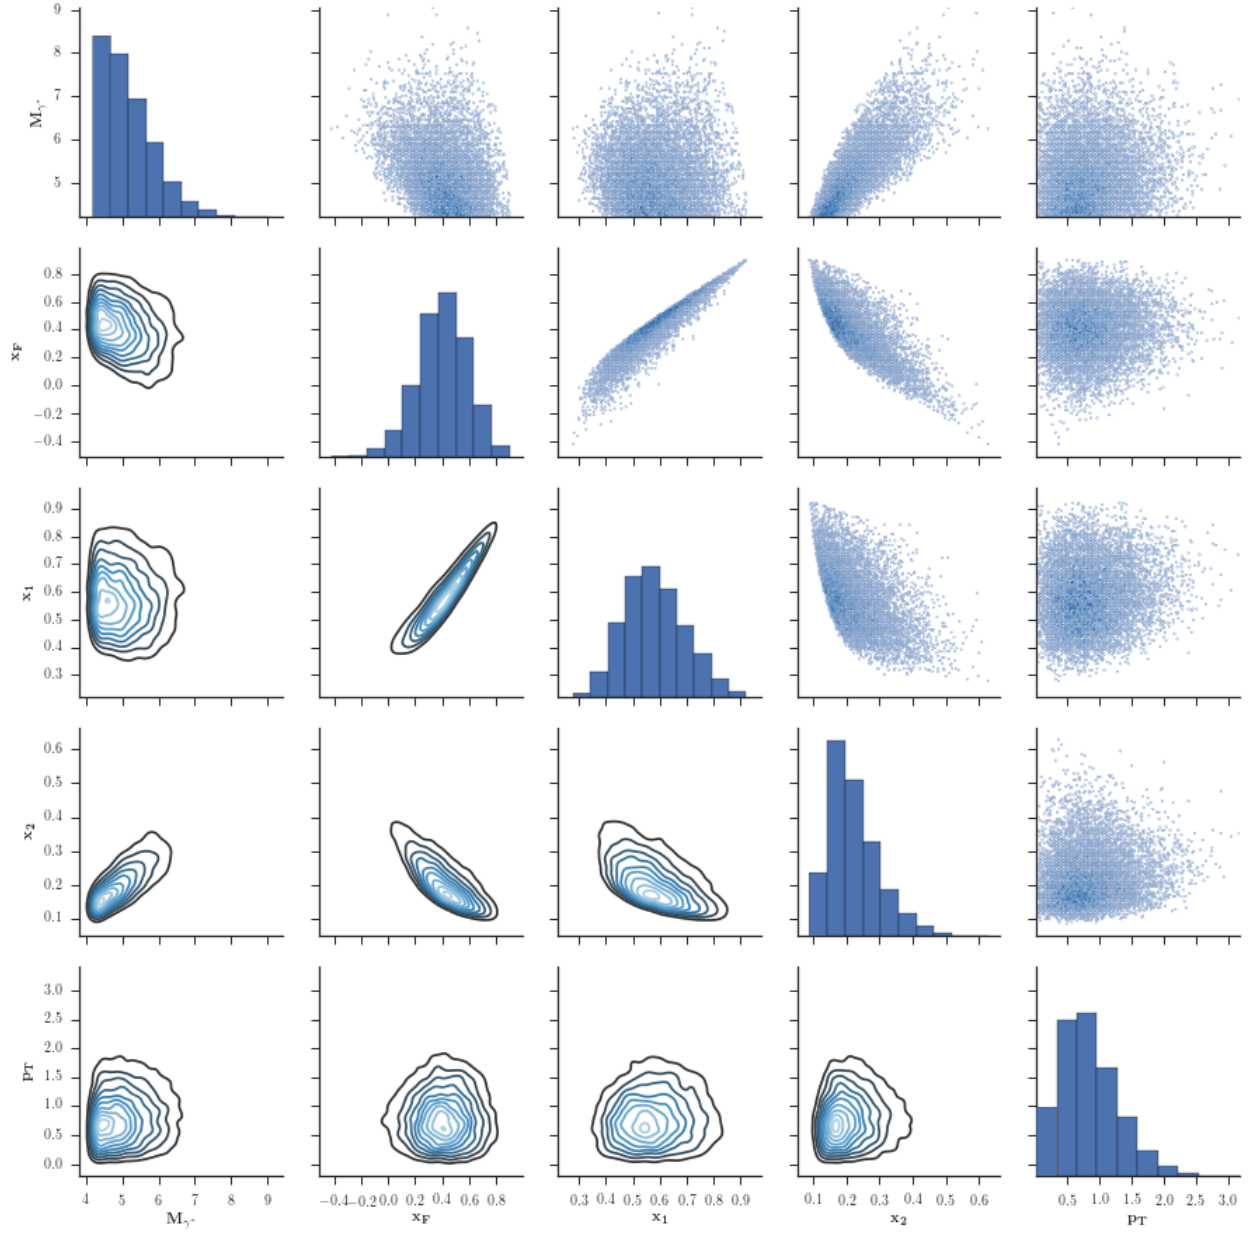


Figure 1.18: A pairplot of the kinematic coverage of SeaQuest, based on weighted Geant4 Monte Carlo simulation, with $4.2 \text{ GeV} < M_{\gamma^*} < 10 \text{ GeV}$ requirement imposed. The diagonal shows a histogram of the single kinematic, while the off-diagonal plots show the contours (bottom left) and scatterplots (top right) of the correlations between variables.

Chapter 2

Apparatus

SeaQuest is the operational name of Fermilab Experiment 906 (*E-906*) performed at the Fermilab Neutrino-Muon (*NM*) experimental area. The experiment was designed to take *high-intensity beam* at relatively *low center-of-mass energy*, provide *good mass resolution*, and allow for *accurate target-to-target systematic normalization*. The apparatus consists of a moving target table, two dipole magnets, 8 hodoscope planes, 24 drift chamber planes, and 4 proportional tube planes. Upstream of the target table (towards the beam source), there is also a Čerenkov counter for beam intensity monitoring and there are several segmented wire ionization chambers (SWIC's) for beam profiling.

2.1 Apparatus Overview

SeaQuest is a fixed-target experiment. In this style of an experiment, a stationary target is placed in the path of an accelerated beam of particles, as opposed to *collider* experiments where two accelerated beams are directed against each other, in opposite directions. The proton beam interacts with the target material and produces a variety of daughter particles. These daughter particles are tracked through a forward spectrometer and selectively filtered dependent on the purpose of the study.

The tuned and monitored 120 GeV proton beam is sent from the Fermilab Main Injector (MI) where the beam protons strike one of 7 targets. The high-momentum charged particles that are produced are focused into the acceptance of the experiment's detectors with the solid iron dipole magnet, FMAG, or NM3S. This solid focusing magnet also sweeps away low-momentum particles and acts as a beam dump / absorber.

The SeaQuest Spectrometer (Fig. 2.1) consists of the aforementioned focusing magnet, several tracking stations that record the positions of charged particles through the length of the spectrometer, and an analyzer magnet to bend the particles between tracking stations. The spectrometer measures particle momenta by recording the bend of each charged particle as it traverses the analyzer magnet, where the magnetic field is known. This is performed by reconstructing the trajectory of a particle in one-half of the spectrometer (before the analyzer magnet) and then similarly reconstructing the trajectory of particles in the other half. If




Figure 2.1: Perspective view of the SeaQuest spectrometer apparatus.

two trajectories can be matched up, then the particle momentum can be extracted by inspecting the change in the track’s direction as a result of passing through the magnet.

The spectrometer’s geometric design and event triggering selection is optimized to detect oppositely-charged pairs of muons while minimizing the sensitivity to various sources of unwanted backgrounds. Positive identification of muons is achieved by requiring signals in the hodoscopes for known muon “roads” along with requiring signals in the proportional tubes located at the farthest end of the experiment, past an iron wall. Electrons and any hadrons are stopped by the solid iron focusing magnet and iron wall farther down while muons will pass through them (relatively) unencumbered.

The coordinate system is defined as the following: the z -axis points along the beam direction, the y -axis points upwards vertically, and the x -axis lies along the horizontal direction in such a way that a right-handed coordinate system is formed. The terms *upstream* and *downstream* are often used when referring directions or regions in the experimental hall. *Upstream* refers to the direction towards the beam source ($-z$ direction) while *downstream* refers to the $+z$ direction of the beam trajectory. The origin of the coordinate system was chosen to be the point where the proton beam meets the *upstream*-facing surface of FMAG, the solid focusing magnet.

2.2 Main Injector Proton Beam

The Fermilab Main Injector (MI) receives protons that have been accelerated by the Radio Frequency Quadrupoles (RFQ), the Linear Accelerator (LINAC), and the Booster, and it continues to accelerate them




Figure 2.2: Spill structure of the beam delivered to SeaQuest.

from 8 GeV up to the nominal energy of 120 GeV. Along the way, radio-frequency cavity (RFC) accelerators in the LINAC and the MI “bunches” up the protons such that the beam has its characteristic 53.1 MHz structure. After the period of acceleration, the protons are then ‘scraped off’ slowly with each passing *turn* of the collected proton beam and sent down the Neutrino-Muon (NM) beam delivery line for approximately five seconds of every minute, called a “slow spill” or just “spill”. Beam is extracted using a resonant process, and the extracted beam retains the 53.1 MHz structure of the Main Injector RFC. Each “bucket” of protons is less than 2 ns long and the time between buckets is approximately 18.8 ns. The spill structure of the beam is depicted in greater detail in Figure 2.2.

The beam sent to SeaQuest is not uniform in time throughout the spill. There are beam buckets in the MI that are intentionally left empty so that the abort kickers can ramp to a full field during a gap in the beam. There are also buckets left empty to allow the injection kickers to inject 8 GeV protons from the Booster without disturbing buckets of protons already in the Main Injector. Typically, 498 of the 588 “RF buckets” in the Main Injector contain protons during the SeaQuest slow spill cycle. It is the case, however, for SeaQuest, that as the buckets are “scraped” off in a procedure called “slow spill extraction,” the intensity of the buckets delivered varies greatly throughout the slow spill. On *average*, each bucket will have $\mathcal{O}(10^4)$ protons and the spill has an intensity of approximately 2×10^{12} protons per second and therefore about

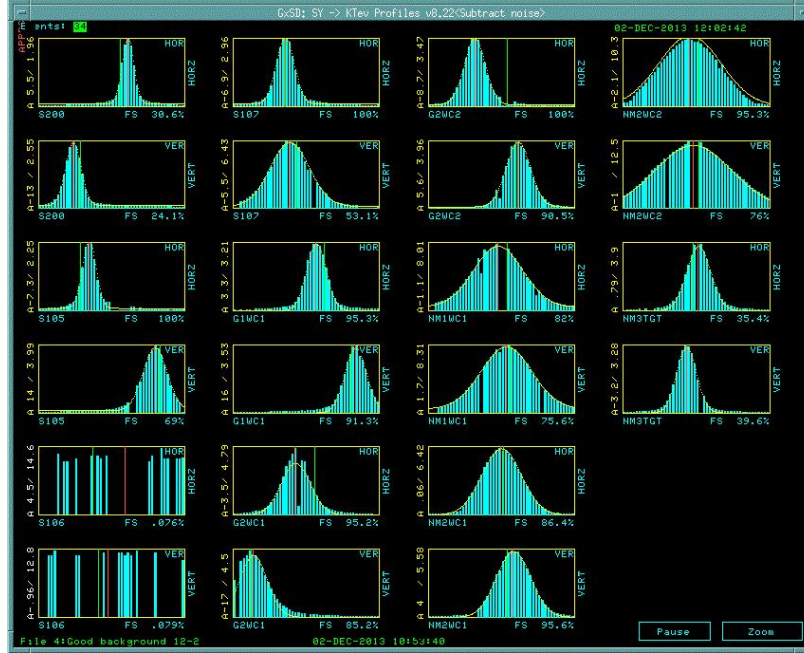


Figure 2.3: Beam profile detailed by SWIC detectors along the NM beam line.

1×10^{13} protons delivered per spill.

Several guiding and focusing magnets bend and deliver beam to the NM beamline which serves both the test beam facility and SeaQuest at NM4. The beam is focused to a width of $250 \mu\text{m}$. The profile, position, and intensity are measured along the NM beamline by several detectors. The intensity of the beam is monitored by an ion chamber (IC) and a secondary emission monitor (SEM) in the NM3 sector of the beamline. The beam profile and position are monitored by SWICs and beam position monitors (BPMs), respectively. The Accelerator Control Network (ACNET) display of the SWIC readout can be seen in Figure 2.3. The closest BPMs and SWICs to the spectrometer were located in NM2 enclosure. The beam profile does not maintain its $250 \mu\text{m}$ shape and spreads slightly as it moves towards the spectrometer. The final beam profile is measured by inspecting the upstream-facing side of the solid targets, and it was found to be approximately 6mm wide by 1mm high.

2.3 Beam Intensity Monitor

SeaQuest's trigger system (described in detail later) mostly fires on fake dimuons caused by two low p_T muons from unrelated pion decays. The hits in downstream hodoscopes from the pions combined with hits in the upstream hodoscope from two other unrelated particles frequently add up to a false dimuon signal. Since this type of fake trigger involves four unrelated particles, the probability that a trigger will occur




Figure 2.4: A single high-intensity event with majority of all detector elements firing off. White space within the rectangles indicates inactive elements whereas red, blue, and green represent elements which have fired during that event. Track reconstruction in these cases is impossible.

increases with I^4 , where I is the intensity of the beam bucket; the number of protons in the triggered beam bucket.

The SeaQuest data acquisition system (also described later in detail) can read out approximately 3000 events per second without significant dead time. During the commissioning run of SeaQuest, the trigger rate was very high and the trigger dead time was close to 100%. These triggers were taken at such high beam intensities that the occupancy of all SeaQuest detector elements was more than 50%, making pattern recognition essentially impossible (see Figure. 2.4). The Beam Intensity Monitor (BIM) was designed to solve this problem.

The SeaQuest Beam Intensity Monitor (BIM) senses when the beam intensity is above a programmable threshold. If an RF bucket with an intensity above this threshold is detected, the BIM sends a signal to inhibit certain triggers until the intensity once more falls below the threshold. The inhibit threshold is tuned frequently as trigger and beam conditions change, but it is typically set at $\sim 95,000$ protons per RF bucket. For reference, a full RF bucket at an intensity of 2×10^{12} protons per spill is $\approx 10,000$ protons.

The beam intensity is measured using an atmospheric pressure gas Čerenkov counter. A gas mixture of 80% Argon and 20% CO_2 is used as the Čerenkov radiator. The counter and readout electronics were designed to have $O(ns)$ time resolution, and a linear response over a large dynamic range. A diagram of the counter is shown in Figure 2.5. A 45 degree aluminized Kapton mirror directs light to a single photomultiplier tube. A *baffle* of black construction paper held parallel to the mirror ensures that the proton path length through the light-radiating gas with respect to the mirror is independent of beam position. A two-inch diameter 8-stage photomultiplier tube (PMT) is positioned close to the mirror so that all Čerenkov light




Figure 2.5: The Beam Intensity Monitor (BIM) Čerenkov counter. Measurements are in inches.

created between the baffle and the mirror falls directly on the aperture of the PMT. It was observed during the commissioning run that after exposure to $\approx 3 \times 10^{17}$ protons (~ 3 weeks of uninterrupted usage), the mirror reflectivity is significantly reduced in the beam spot, and the mirror then needs to be replaced.

The signal from the BIM is integrated and digitized using a custom charge (Q) integrator and encoder (QIE) integrated circuit board, which comes from a family of circuits used first by the KTeV experiment at Fermilab[56]. The chip is clocked with the Main Injector RF clock and provides an ADC (analog-digital conversion) every 18.8 ns clock cycle. The light incident on the photomultiplier tube is attenuated using neutral density filters (NDF's) so that the QIE least count corresponds to ~ 30 protons per beam bunch.

In addition to inhibiting triggers during high-intensity periods of beam, the BIM readout module also provides critical information used to calculate the number of protons incident on the SeaQuest targets while the experiment is ready and able to trigger. This value is needed to normalize SeaQuest cross section measurements. The BIM readout module provides the following:

- Sum of all ADC signals for the entire spill (QIESum).
- Sum while inhibit is asserted at trigger logic.
- Sum during trigger dead time (while not inhibited).

- A snapshot of beam intensity 16 buckets before and after the triggered RF bucket

These are used to calculate a ratio of protons that were ‘live’ (the experiment can trigger) via the following:

$$\%live = \frac{QIESum - (inhibit\ sum + dead\ time\ sum)}{QIESum} \quad (2.1)$$

$$liveProton = totalProton \cdot \%live \quad (2.2)$$

where “totalProton” is the intensity value recorded from the SEM detector located just upstream of the BIM Čerenkov counter. The SEM itself is calibrated by foil activation. The snapshot of the triggered RF bucket intensity along with the 32 surrounding RF bucket intensities is used for studies and corrections of the rate-dependent effects on detector efficiencies and reconstructed measurements.

2.4 The SeaQuest Targets

A wide range of atomic weights (from 2 to 184) is required to do an A-dependence study of the Drell-Yan process. At SeaQuest, the targets used are $^1H(\ell)$, $^2H(\ell)$, C , Fe , and W . In addition to the two liquid targets and the three solid targets, two positions on the target table were used for measuring background signal rates: an empty flask, identical to the flasks used for the 1H and 2H targets, and a single empty solid target holder. Colloquially speaking: the 1H target is interchangeably referred to as the liquid hydrogen target, ℓH_2 , $LH2$, or H_2 ; 2H is likewise referred to as the liquid deuterium target, ℓD_2 , $LD2$, or D_2 ; the empty flask is referred to as the “Empty” target; the empty solid target holder is referred to as the “None” target.

These are all mounted on a laterally-moving, remotely positionable table (in the $\pm x$ direction), able to move over a range of 91.4 cm. The table’s center is located at $(0, 0, -1.25)$ meters, directly in front of the upstream face of FMAG, the solid iron focusing magnet. Because of the ~ 5.0 cm diameter of the targets and the 6x1mm dimensions of the beam, the targeting efficiency was 100%. The details of the target materials are summarized in Table 2.1, and the layout of the target table can be seen in 2.6.

The H_2 gas used is “Ultra High Purity 5.0 Grade” or 99.999% pure. The deuterium has come from two different sources. The first of these is a Fermilab-provided supply of gas left over from previous bubble chamber experiments. This gas was known to have a small hydrogen contamination and was measured by mass spectroscopy to have a composition of 85.2% D_2 , 12.7% HD, 1.2% 4He , and 0.8% H_2 by mole. As the analysis of experimental data commenced, handling the ramifications of the D_2 impurity came under focus.

Position	Material	Density [g/cm ³]	Thickness [cm]	Interaction Length	Spills/Cycle (%spills)
1	H ₂	0.07065	50.8	0.06902	10 (43%)
2	Empty	NA	NA	0.0016	2 (9%)
3	D ₂	0.1617	50.8	0.1144	5 (22%)
4	None	NA	NA	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: Characteristics of the seven SeaQuest target positions. The “Spills/Cycle” is only a typical configuration and can vary according to needs and running configurations. The non-zero interaction length of the empty flask is due to the 51 μm -thick stainless steel end-caps of the flask and the 140 μm -thick titanium windows of the vacuum vessel that contains it.




Figure 2.6: The layout of the target table and its seven target positions, as seen from above.

Unexpected bottle-to-bottle variation in contamination became evident, and the sample-taking methodology itself for spectroscopy became suspect of introducing contamination. In order to simplify analysis and reduce the substantial complexity and cost of further gas analysis, SeaQuest switched to commercially available “Research Grade” D₂, which is better than 99.6% pure with virtually all HD to balance. The data analyzed in this paper deals with the impure D₂ target material before this switch. Further information on the D₂ composition and how it is handled in the context of analysis will be covered in Chapter 5.

Each of the three solid target positions is divided into three disks of 1/3 the total thickness provided in Table 2.1. These disks are spaced 25.4 cm apart to approximate the distribution of the liquid target, thereby minimizing target-dependent variation in spectrometer acceptance. The one exception to this is that during the Run II period the iron disks were more closely spaced (17.1 cm). The motivation to place these iron disks closer together than the rest during Run II is still unclear.

The target table is able to move between two different targets in about 30 seconds. This allows a change between targets in the 55 seconds between successive spills. With this frequent target interchange, the systematic uncertainties associated with drifts in beam characteristics, monitor gains, and detector efficiencies are reduced to a minimum when investigating A-dependent ratios. How much beam time each target received is determined by interaction lengths of the targets along with the amount of statistics desired for certain targets. As the flagship measurement of SeaQuest is the \bar{d}/\bar{u} asymmetry, more emphasis was placed on the hydrogen and deuterium than the nuclear targets. The spills per cycle and beam time allocation can be found in Table 2.1.

2.5 Focusing and Analyzing Magnets

Two large dipole magnets are used in the experiment to select forward going ($x_F > 0$) dimuons, reject low-momentum particles, and analyze their kinematic characteristics. The most upstream magnet, denoted “FMAG”, is a solid iron A-frame magnet with an aperture of 1.22m in the x -direction and 66 cm in the y -direction. It is assembled from 43.2 cm x 160 cm x 503 cm iron slabs, as shown in Fig. 2.7. The magnet has no air gap, and the iron has extremely high purity, allowing a 2000 A excitation current to generate a nearly constant, central magnetic field of 1.9 Tesla (yielding a 2.91 GeV/c total magnetic deflection). The field is generated by exciting the embedded aluminum “*bedstead*” coil to 2000 A at 25 V (50 kW). The current exciting FMAG is monitored by the Fermilab ACNET system and is broadcast to the SeaQuest slow data acquisition system every acceleration cycle. The excitation is also input to the beam-disabling safety system in order to prevent the beam from hitting the SeaQuest spectrometer when FMAG is not fully powered.




Figure 2.7: Perspective drawing of FMAG’s aluminum coils embedded in an arrangement of iron slabs.

FMAG also acts as the beam dump for the 120 GeV beam. There is a 5 cm diameter by 25 cm deep bore drilled into the upstream end of FMAG (recall, this is the origin of the experiment’s coordinate system). The 120 GeV protons that do not interact in the SeaQuest targets 125 cm upstream of FMAG interact in the central iron slab. Most of the 2.0 kW beam power is dissipated in this slab and is eventually conducted to the coils and external surfaces to be radiated away.

The downstream magnet, denoted “KMAG”, is a 300 cm long iron rectangular magnet with a 289 cm wide by 203 cm high central air gap. It was originally constructed by the KTeV collaboration[57] at Fermilab. It is excited to a central field of 0.4 T (0.402 GeV/c magnetic deflection) by 1600 A at 270 V (430 kW). The spatial distribution of the magnetic field in KMAG was measured by the KTeV group and re-verified by SeaQuest. In normal running conditions, both FMAG and KMAG bend muons horizontally in the same direction. This two-magnet configuration is often referred to as a focusing spectrometer.

The 2.91 GeV/c and 0.402 GeV/c magnetic deflections deliver a transverse-momentum (p_T) kick to the charged particles passing through the spectrometer. The magnets bend the paths of the muon in the $\pm x$ direction, with the sign depending on the orientation of the magnetic fields and the particles’ charges. Between Run II and Run III of data taking, the current direction was reversed, thereby reversing the direction of the magnetic fields. During Run II, the magnetic fields were pointing in the $-y$ direction, and in Run III, the magnetic fields were flipped to point in the $+y$ direction. This was done for two reasons: (1) to identify any left-right asymmetries in the experiment, and (2) to limit the amount of radiation on the electronics in the experimental hall, as a large number of positive particles were being swept directly towards the electronics racks during Run II.

2.6 Beam Dump, Shields, and Absorbers

In order to prevent damage to the downstream detectors from the beam and reduce signals from incidental radiation, the spectrometer is designed with a beam dump and two hadron absorber walls. Approximately 125 cm downstream of the target table is the water-cooled beam dump whose upstream face is located at (0.0, 0.0, 0.0) m. The beam dump is one of the many solid iron 5m blocks that fill and surround the FMAG coils. The whole length of the beam dump along the beam axis is equivalent to ≈ 35 nuclear interaction lengths of iron.

Between the downstream face of FMAG and Station 1, there is a 2 cm thick wall of borated polyethylene which is put in place as a fast neutron shield. This material is 5% boron by weight, with the rest being polyethylene. The polyethylene contains high hydrogen content, making it an effective fast neutron radiation shield, slowing down the fast neutrons down to thermal speeds. The boron in the material provides attenuation of thermal neutrons, thus reducing the levels of capture-gamma radiation elsewhere in the experiment. Borated polyethylene at this thickness is a common and optimal neutron shielding material for areas of low to intermediate neutron flux where the temperature is below 82°C [58]. These conditions make the downstream side of FMAG ideal for its placement.

Farther downstream, there is another hadron absorber wall located between Station 3 and Station 4. The absorber wall consists of a stack of 98 cm thick iron blocks. This is an equivalent of ≈ 6 nuclear interaction lengths. The purpose of this wall is to identify muons at the rear of the apparatus by effectively blocking all other types of particles. The only charged particles which can penetrate this absorber wall that weren't swept away by either magnet are muons.

2.7 Tracking Detectors

The tracking detectors are the instruments used for measuring the values of the kinematic variables of the dimuon pairs. Several different types of detectors are grouped together to form a detector *Station*. The types of detectors used are hodoscopes, wire chambers, and proportional tubes. There are four stations throughout the experimental hall that provide tracking information at different points along the spectrometer, numbered from 1 to 4 in order of increasing z . Station 1 is located between FMAG and KMAG. Station 2 is located at the downstream face of KMAG. Station 3 and 4 are just upstream and downstream, respectively, of the iron absorber wall. The Station layout can be seen in Fig. 2.8

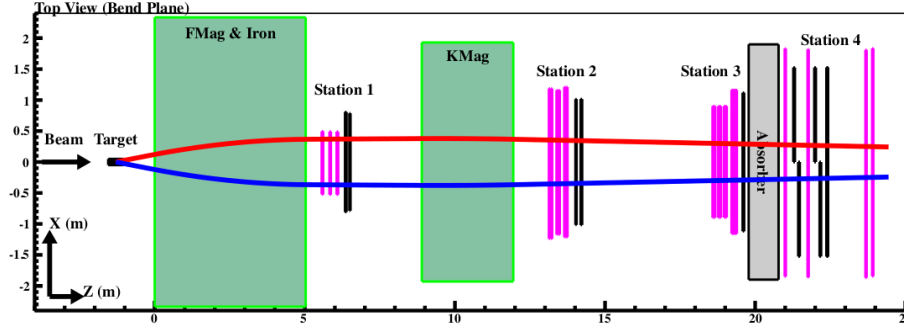


Figure 2.8: Spectrometer layout of FMAG, KMAG, and Detector Stations 1-4.

2.7.1 Triggering Hodoscopes

Hodoscope arrays are located at each of the four detector stations. These detectors' primary usage is to select events with two oppositely-charged muon tracks in them. Certain 'roads' through the spectrometer are defined in the fast trigger logic as $[H1, H2, H3, H4]$ paddle groups, and when two desired roads are observed in a given event, the trigger system tells the data acquisition systems to record that event's data. In addition to this, the hodoscopes provide analysts with the ability to discard or ignore certain hits in adjacent chambers for which there is no corresponding nearby hodoscope hit. This is useful in decreasing the hit multiplicities in the wire chambers, which in turn decreases the combinatoric complexity for reconstruction algorithms.

Each of the eight hodoscope planes is split into two halves: top and bottom in the case of planes with vertically-oriented paddles, or left and right for planes with horizontally-oriented paddles (denoted by 'T', 'B', 'L', and 'R', respectively). In each half-plane, the hodoscopes are a set of long rectangles arranged 'picket fence'-style with a small 0.3175 cm overlap, as to prevent any particles from possibly slipping between paddles. A single hodoscope detector element is composed of plastic scintillator material connected to Philips XP 2008 photomultiplier tubes (PMT) by plexiglass light guides. Stations 1, 2, and 4 each have two hodoscope planes, with planes of both vertically- and horizontally-oriented paddles (for measuring in x and y , respectively). Station 3 consists only of a vertically-oriented plane, and thus measures in the x -direction only, which is in the experiment's $x-z$ bend plane. The hodoscope planes are named according to detector station and the direction that they measure. For example, the y -measuring hodoscope plane in Station 2 is called "H2Y". The individual half-planes are named according to detector station and which half it is of the two. For example, the top half of the x -measuring hodoscope plane in Station 1 is referred to as "H1T". As such, the "H3X" detector is composed of "H3T" and "H3B". The detailed specifications of each hodoscope plane are given in Table 2.2. A precise alignment of the hodoscopes was achieved by examining the distributions of positions of tracked muons at each hodoscope plane when a given hodoscope element in that plane was

fired.

Detector	Paddle Width [cm]	Overlap [cm]	# of paddles (per half-plane)	Width \times Height (per half-plane) [cm] \times [cm]	z -position [cm]
H1X	7.32	0.3175	23	162 \times 70	666
H1Y	7.32	0.3175	20	79 \times 140	653
H2X	13.00	0.3175	16	203 \times 152	1421
H2Y	13.00	0.3175	19	132 \times 241	1400 (L), 1406 (R)
H3X	14.59	0.3175	16	228 \times 168	1959
H4X	19.65	0.3175	16	305 \times 183	2234 (T), 2251 (B)
H4Y1	23.48	0.3175	16	152 \times 366	2130 (L), 2146 (R)
H4Y2	23.48	0.3175	16	152 \times 366	2200 (L), 2217 (R)

Table 2.2: Parameters of all hodoscope planes. z -positions of H2Y and H4 half-planes are offset slightly due to the half-planes themselves overlapping.

2.7.2 Drift Chambers

Each of Stations 1, 2 and 3 is equipped with a drift chamber (DC) to measure the passing x and y positions of muons at its singular z -location. These measured positions are critical for reconstructing the trajectory of muons, and thereby their kinematics. Each DC contains six drift chamber planes, arranged in three pairs with parallel wire orientations (each pair referred to as a “view”). Wires oriented vertically (x -measuring) are referred to as being in the “X” view and at angles of $+14^\circ$ and -14° with respect to the y -axis are the “V” and “U” planes, respectively. The second plane in each view is offset from the first by one-half of a wire-to-wire distance (“*cell width*”) in order to resolve the left-right ambiguity of drift direction. This offset plane of each pair is referred to as the *primed* plane, and is denoted with a ‘ \prime ’. So, in each DC, there are X, X’, U, U’, V, and V’ planes, with primed and unprimed planes (like X and X’) constituting a view.

The individual drift chambers at Stations 1 and 2 are called “D1” and “D2”, respectively. Station 3 has two drift chambers since the desired acceptance area it has to cover is substantially larger than one DC can cover. These are split vertically to cover the top and bottom halves, and are called “D3p” and “D3m” where “p” and “m” stand for “plus” and “minus”. Table 2.3 summarizes the parameters of the DC’s. D1 and D3m have been upgraded during the data taking, as listed in Tab. 2.4. This original and upgraded versions are referred to as *DN.1* and *DN.2*, respectively.

The acceptance size of each chamber has been adjusted with a Drell-Yan event simulation in order to be as sensitive as possible to the x_2 range of interest. Particularly, the greater the acceptance width is, the higher the reach in x_2 is. This makes the hit-rate tolerance of the chambers a key feature because the spectrometer is exposed to a large number of background particles, particularly near the edges where the desired x_2 events

Chamber	Plane	Number of 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: Parameters of all chambers. Those of primed planes are almost the same as of unprimed planes. The z -positions of U and V planes are relative to those of X planes.

Run	Period	Chamber combination
1	2012 Mar.-2012 Apr.	D1.1, D3m.1
2	2013 Nov.-2014 Aug.	D1.1, D3m.2
3	2014 Nov.-2015 May	D1.1, D3m.2
3	2015 Jun.-2015 Jul.	D1.2, D3m.2
4	2015 Sep.-	D1.2, D3m.2

Table 2.4: Combination of D1 and D3m chambers per data taking period.

occur. It is particularly significant for the most-upstream station (i.e. Station 1) which receives the highest hit rates. Experimental data shows that the rate tolerances are 3.0 MHz/wire at D1, 1.6 MHz/wire at D2 and 0.7 MHz/wire at D3 with a beam intensity of 5×10^{12} protons/spill. The gas-amplification gain should not be degraded under these hit rates.

For Run 2 and beyond, the gas mixture used for almost all the chambers is Argon:Methane:CF4 (88%:8%:4%) with a drift velocity of about 20 $\mu\text{m}/\text{ns}$. A “fast gas” mixture used for D1.2 (upgraded D1) is Argon:CF4:Isobutane:Methylal (68%:16%:13%:3%) with a drift velocity of 50 $\mu\text{m}/\text{ns}$ and thereby a better hit-rate tolerance. This is “fast” in comparison to the $\approx 20\mu\text{m}/\text{ns}$ rest of the DCs. The spatial resolution of each plane is required to be 400 μm , which corresponds to a momentum resolution of $\Delta p/p = 0.03 \cdot p$ (GeV/ c). The resolution of dimuon invariant mass is dominated by the multiple scattering in FMAG; the chamber momentum resolution is about 10% of the total mass resolution at maximum.

The D1.1, D2, and D3m.1 chambers have been inherited from previous Drell-Yan experiments that have been conducted at Fermilab. Chambers D2 and D3m.1 have their origin in SeaQuest [59] and D1.1 is from SeaQuest’s direct predecessor, E-866/NuSea [1, 60]. Since these chambers have not been used for decades since E-866, they had to be refurbished by restringing $\approx 30\%$ of their sense wires due to them being loose or broken. Newly supplied electronic readout boards were also mounted on these chambers.

The D3p and D3m.2 chambers were designed and constructed specifically for this experiment in order to cover the large acceptance required at Station 3. D3p was newly constructed by the TokyoTech SeaQuest collaborators and was shipped from Japan to Fermilab. The first part of data taking, Run 1, was carried out using D3m.1 while preparing for the construction of D3m.2. The newer D3m.2 is wider than D3m.1 by 25 cm at each side, allowing the high- x_2 statistics on \bar{d}/\bar{u} and Y_A/Y_{2H} to increase by $\approx 20\%$ at $x_2 \sim 0.3$ and $\approx 10\%$ at $x_2 \sim 0.4$. The operational stability also improved, as D3m.1 suffered from frequent dead/noisy wires, HV trips, and leak currents. The D1.2 chamber was also designed and constructed for this experiment by the University of Colorado Boulder. As it is wider than D1.1 by 25 cm at each side and greater hit-rate tolerance, the anticipated statistics in the high- x_2 region is expected to increase still more. It was installed in the experimental hall near the end of Run 3.

2.7.3 Proportional Tubes

Downstream of Station 3 and the 1 m thick iron hadron absorber wall is Station 4 with its hodoscope planes and proportional tube detectors (prop tubes). At Station 4, the only beam-induced particles that remain that can leave tracks in an ionization detector are high energy muons. The prop tubes enable the task of muon particle identification (PID) for the experiment and consist of four planes. Each detector plane is made




Figure 2.9: Proportional tube top ($x-z$ plane) view. Figure 2.10: Proportional tube side ($y-z$ plane) view.

of 9 prop-tube modules, with each module assembled from 16 12-ft long 2" diameter prop-tubes staggered to form two sub-layers. The first and fourth planes are oriented along the horizontal direction (tubes parallel to the floor) to provide positional measurements in y , as shown in Fig. 2.9. The second and third planes are arranged vertically to measure x -position as shown in and Fig. 2.10.

Each plane of proportional tubes is composed of 9 *modules*, each having a set of 16 proportional tubes: 8 in line in a horizontal or vertical orientation and 8 more offset (“primed”) from the first 8 by 0.5 inches. This offset structure prevents any muons from going undetected between tubes and allows for left-right disambiguation for two hits in adjacent primed-unprimed sub-plane pairs in the same module. The modules are labeled 1–9 in increasing order from left to right for the x -measuring planes and from 1–9 in increasing order from top to bottom. So, the top module of the first vertical plane of prop tubes is P2H1. Additionally, in each module, the primed and unprimed sub-planes are referred to as f for “front” (facing upstream) and b for “back” (facing downstream). This substructure can be seen in detail on Fig. 2.9 and Fig. 2.10.

A single prop tube is made of 2-inch diameter aluminum tubes with a wall thickness of 1/16 inches. The central anode wire is a gold-plated $20\mu\text{m}$ diameter tungsten wire kept at approximately 1.95 kV. Considering the staggered nature of each plane’s substructure, one can in principle achieve a spatial resolution of 0.3 mm. During Runs 2 and 3 of data taking, a resolution of 0.5 mm for high energy muons was observed, which is more than sufficient for muon identification purpose. The gas mixture for prop-tubes is P-10 (Ar:Methane = 90:10) mixed with a 10% CF_4 gas (Ar: CO_2 : CF_4 = 70:20:10) which yields the maximum drift time about 400 ns. With this maximum, the prop tubes can handle a singles rate up to 2 MHz while normal operational hit rates are typically below 1 MHz.

A typical desired high-energy muon within spectrometer acceptance will traverse through two prop-tubes

in each plane and induces hit signals on two anode wires. The path of a track is reconstructed from the drift time measured on the two anode outputs, with a custom TDC board that provides 0.44 ns timing resolution. With the hit information reconstructed from readouts of the forward and backward planes in $x-z$ and $y-z$ direction, precise reconstruction of the track trajectories can be obtained. Ideally, 8 hits from the 4 planes are used to form a track pointing back to a hit upstream. If such is the case, a candidate muon track is successfully identified.

Detector Plane	Number of modules	# tubes per module	Width (x) × height (y) [cm] × [cm]	z -position of front (back) sub-plane [cm]
P1H	9	16	368×368	2,099 (+4)
P1V	9	16	368×368	2,175 (+4)
P2V	9	16	368×368	2,367 (+4)
P2H	9	16	368×368	2,389 (+4)

Table 2.5: Parameters of the four proportional tube planes.

2.7.4 Mass Resolution from Chamber Resolution

The mass resolution of the spectrometer is, in part, limited by the resolution of the tracking chambers and the distances between them. For an arbitrary set of hit positions in Stations 1, 2, and 3 (r_1, r_2, r_3) which are separated from each other in z by distances z_{12} and z_{23} , with KMAG in between Stations 1 and 2 supplying a P_{kick} , one can derive $\Delta P/P$. A line, or track segment, is first reconstructed between r_2 and r_3 . The slope (and its uncertainty) of this track segment in the $z-x$ plane is:

$$s_{23} = \frac{r_3 - r_2}{z_{23}} \quad , \quad \Delta s_{23} = \frac{1}{z_{23}} \sqrt{\Delta r_3^2 + \Delta r_2^2} \quad (2.3)$$

The values of Δr_n here are the position resolutions of the individual tracking chambers at each Station. One can use this slope to project a trajectory through KMAG and onto Station 1. The momentum is calculated by seeing where the actual hit position is in Station 1 and looking to the distance (d) between where the particle hit the station and where it would have hit had KMAG's magnetic field not existed. This is called a *sagitta analysis*, as the magnetic field moves the particle's trajectory as if along a circle, and we are looking to the particle's path along the sagitta of that circle. Ultimately, the momentum uncertainty is related to this distance as:

$$\frac{\Delta P}{P} = \frac{P}{P_{kick}} \Delta d \quad (2.4)$$

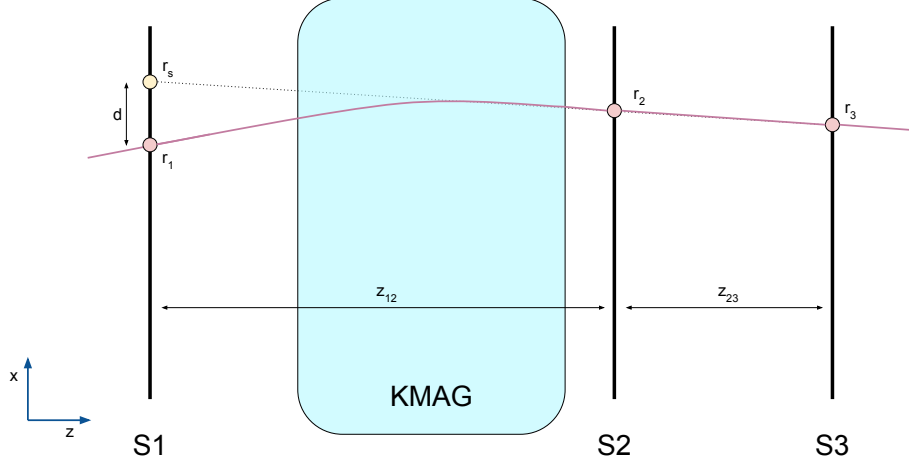


Figure 2.11: A simplistic depiction of a track passing through from Station 1, through K MAG where its path is bent by the magnetic field, and then straight on through Stations 2 and 3.

The expected sagitta point (r_s) is located at:

$$r_s = s_{23}z_{12} + r_2 \quad (2.5)$$

$$\Delta r_s = \sqrt{\Delta s_{23}^2 z_{12}^2 + \Delta r_2^2} = \sqrt{\Delta r_3^2 \left(\frac{z_{12}}{z_{23}}\right)^2 + \left(1 + \left(\frac{z_{12}}{z_{23}}\right)^2\right) \Delta r_2^2} \quad (2.6)$$

And the distance, d follows as:

$$d = r_s - r_1 \quad (2.7)$$

$$\Delta d = \sqrt{\Delta r_s^2 + \Delta r_1^2} = \sqrt{\Delta r_3^2 \left(\frac{z_{12}}{z_{23}}\right)^2 + \left(1 + \left(\frac{z_{12}}{z_{23}}\right)^2\right) \Delta r_2^2 + \Delta r_1^2} \quad (2.8)$$

The momentum resolution due to chamber resolution can therefore be described as:

$$\frac{\Delta P}{P} = \frac{P}{P_{kick}} \sqrt{\Delta r_3^2 \left(\frac{z_{12}}{z_{23}}\right)^2 + \left(1 + \left(\frac{z_{12}}{z_{23}}\right)^2\right) \Delta r_2^2 + \Delta r_1^2} \quad (2.9)$$

The dependence of the resolution on P/P_{kick} works such that the larger the P_{kick} is (w.r.t. particle momentum P), the more the track is bent, and the more precisely the original momentum is defined. Besides chamber resolution, other contributions to this momentum resolution are from multiple scattering through FMAG's iron beam dump and from the spectrometer's angular resolution.

The mass resolution is linked to momentum resolution via the following relation

$$\frac{\Delta M}{M} = \frac{\Delta P}{2P} \quad (2.10)$$

The factor of two arises from the fact that the Drell-Yan dilepton mass comes from two independent muons with the same momentum resolution. This mass resolution is important because a mass selection range is imposed upon analysis. Also, the mass resolution, by association, contributes to the $\Delta x_2/x_2$ resolution, which is the key dependent variable for SeaQuest measurements.

$$\frac{\Delta x_2}{x_2} \approx 0.57\Delta x_F + 0.012M^2 \frac{\Delta M}{M} \quad (2.11)$$

2.8 Trigger

The SeaQuest Trigger System is designed to quickly select candidate dimuon events from the high-rate, high-background environment using discriminated hodoscope signals.

2.8.1 Design Requirements

The trigger must select events of interest to the main physics goals with good enough efficiency (maximum signal-to-background ratio) to facilitate high-statistics analyses. In general, it is optimized to accept high-mass (4-10 GeV) dimuons originating from the targets and beam dump. The event selection is also designed to intentionally reduce acceptance of dimuons from other, higher-rate sources, such as J/Ψ decays and non-Drell-Yan dimuons originating from the beam-dump, though enough J/Ψ events are triggered on to allow high-statistics analysis of J/Ψ physics. Other goals of the trigger design are to keep the triggering rate low enough to maintain an acceptable DAQ livetime and to keep the trigger's internal throughput deadtime-free. By design, there should be no bottleneck at this stage, and the trigger should be capable of firing on any and all RF-buckets while the DAQ is live.

The hardware, firmware, and design should be sufficiently flexible to quickly accommodate changes in the spectrometer, beam conditions, and physics goals. Any changes to the geometric acceptance of the spectrometer, such as new/moved detectors or changes in the magnetic fields must be immediately reflected in the trigger selection in order to maintain high signal efficiency and good background rejection power. Similarly, a change in the beam duty factor or intensity should be accompanied by a change in the event selection, ensuring trigger rate optimization. Finally, the trigger should be capable of specific modifications to facilitate special runs for other physics goals. For these reasons, the design of the trigger system underscores

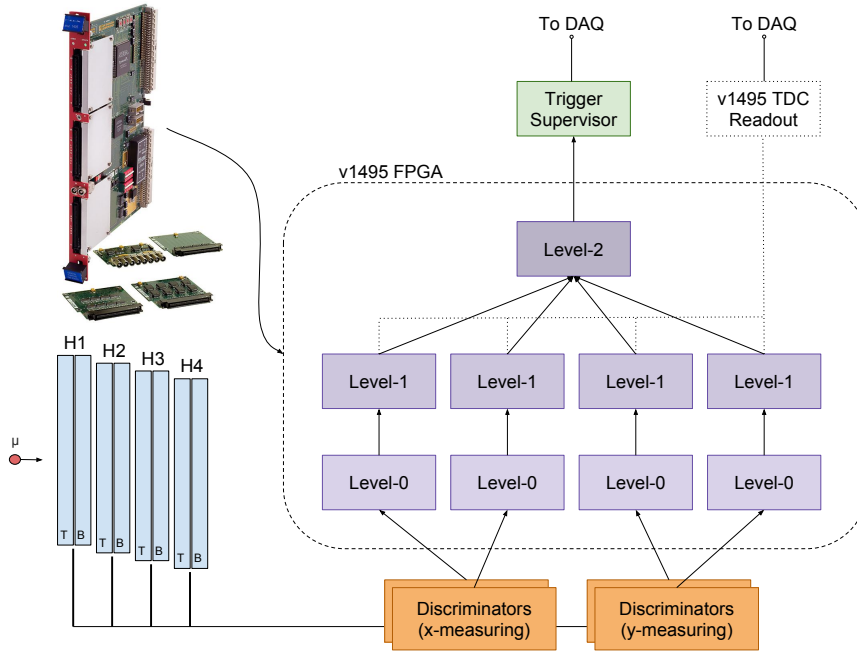


Figure 2.12: The trigger system at SeaQuest is composed of 9 CAEN v1495 FPGA modules. These modules output a trigger signal to the Trigger Supervisor, which tells the DAQ when to record data.

a significant need for flexibility in its hardware and design.

Lastly, the design of the trigger system should include self-diagnostic capabilities, allowing for constant monitoring of the trigger system’s performance. Internal pulser-testing is employed to test the function of each compiled firmware every time the trigger logic changes. For added transparency, data from the internal TDC’s is used by online and offline software to monitor the self-consistency of the trigger for each recorded physics event.

2.8.2 Trigger Hardware

The triggering system, which is depicted in Fig. 2.12, begins with the Stations 1-4 hodoscope arrays. Hodoscopes are used for this task as the time resolution and high speed is limited only by the speed of light through the scintillator paddles and secondary emission of electrons (through the stages of the PMTs). This occurs over the span of nanoseconds, which allows for fast triggering on RF buckets that occur every 19 ns during spills. The hodoscope planes, both x- and y-measuring planes, are separated into top and bottom sections as far as triggering is concerned. The signals from the hodoscope arrays are passed through a set of discriminator modules, which will only pass a signal on to the rest of the trigger if the raw PMT pulse crosses a present threshold.

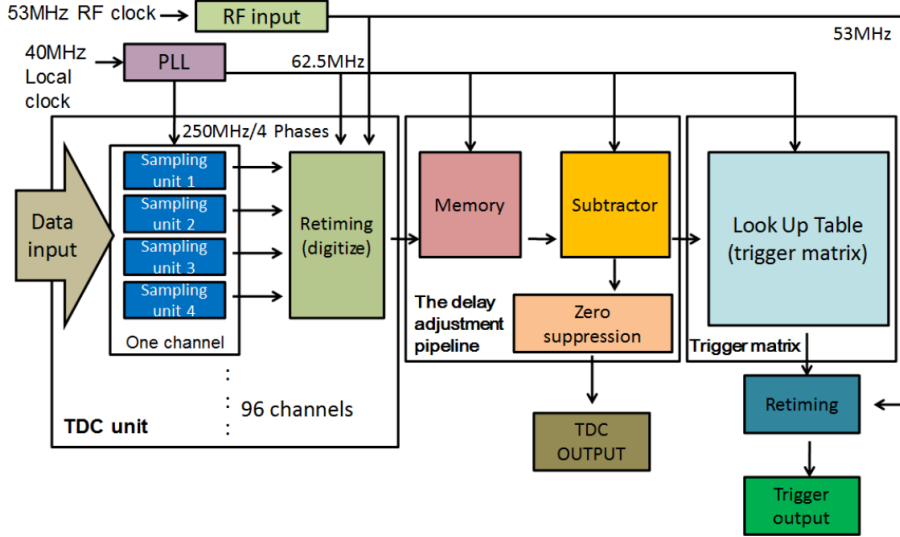


Figure 2.13: A block diagram of the major functions of the v1495 FPGA [62].

The output of the discriminators is passed along to one of the nine CAEN v1495 FPGA (Field Programmable Gate Arrays) VMEbus (Versa Module Europa) modules [61]. This hardware trigger consists of a single decision stage with a three-step parallel pipeline. These steps are referred to as the *level-0*, *level-1*, and *level-2* stage triggers. The detailed internal flowchart of the CAEN v1495 electronics can be seen in Fig. 2.13.

At the first step, level-0, there is nothing actively performed during nominal data taking. The signal is passed on through, unaltered, to level-1. As it stands, level-0 is not used for any of the triggering modes. This stage is, however, extensively used for the pulser testing and self-diagnostic purposes. The level-0 stage is capable of being programmed to pulse arbitrary patterns on to the level-1 stage, which is very useful for observing how level-1 (and subsequently, level-2) will behave in a controlled setting. Whenever a new firmware is installed on the v1495, a pulser test is performed to ensure that all design requirements are met before uploading the firmware for production data taking.

The level-1 stage records the hit signals from the x- and y-measuring hodoscope hits split into top and bottom partitions with respect to their location in the SeaQuest spectrometer. The hit patterns are tested via a lookup table against a preselected set of hodoscope hit patterns, or “*trigger roads*”. Each trigger road corresponds to a set of four hodoscopes: one from each station, with all being in the top or bottom half of the spectrometer. In this lookup table, there is a charge and approximate p_T value for each preselected road. These are known via studies of Monte Carlo studies of muon paths from di- or single-muon events originating from the targets or beam dump. The roads were selected such that as many good dimuons are selected while reducing roads that are dominated by background signal that might cause trigger and/or DAQ downtime.

Trigger	Condition	Sign	# μ	Prescale
FPGA-1	$T \wedge B$	Opposite	2	1
FPGA-2	$(T \wedge T) \vee (B \wedge B)$	Opposite	2	1000
FPGA-3	$T \wedge B$	Same	2	123
FPGA-4	$T \vee B$, all p_T	$+\vee-$	1	30000
FPGA-5	$T \vee B$, $p_T > 3 \text{ GeV}$	$+\vee-$	1	2000
NIM-1	Y-coincidence	NA	NA	30000
NIM-2	X-coincidence	NA	NA	1000
NIM-3	Random RF	NA	NA	1000

Table 2.6: Trigger settings for Run 3 and beyond. Prescale figures shown are typical values and were adjusted as experimental needs were tuned. NIM-1 and NIM-2 triggers’ exact conditions could be changed and reconfigured from the counting room.

During data taking, only the x-measuring hodoscopes were used for the v1495 trigger. As such, the two level-1 stages that look to the y -measuring hodoscope signals were unused. The level-1 trigger logic identifies four-out-of-four X1-X2-X3-X4 coincidences, which are characteristic of high p_T single muons produced from the targets or beam dump. Each time a candidate coincidence that satisfies a preselected road is found, the level-1 step outputs certain logical bits indicating the track’s charge, detector half (top or bottom), and p_T bin.

The third step of the triggering logic is the level-2 trigger, or the “*Track Correlator*” stage. This step takes the outputs of level-1 (charge, detector half, p_T bin) and finds if any combinations satisfy any of the preprogrammed triggering modes. If one of these modes was satisfied, a signal is sent to the Trigger Supervisor (TS) module that communicates with the DAQ system to record an event at a specific synchronized RF bucket.

2.8.3 Triggering Modes

The v1495 FPGA had five configurable sets of physics triggers were used in the trigger system for the majority of production-level data taking. These five are referred to interchangeably as “Matrix” and “FPGA” triggers, followed by the number 1-5 indicating which of the five. These modes are described in detail in Table 2.6.

Different trigger settings are used for capturing data for different purposes. The primary triggering mode is FPGA-1, which looks for a combination of two roads of opposite sign that reside in opposite halves of the spectrometer. FPGA-2 looks for good muon pairs that occurred in the same half. This, however, turned out to have a combinatoric problem that had two adverse effects: track reconstruction was difficult for these types of events, and the trigger fired off too often, causing a high trigger and DAQ downtime. As such, FPGA-2 was disabled for much of data taking. FPGA-3 trigger looks for the same kind of signal as FPGA-1, but with same-sign muon pairs. This type of event is useful for analyzing the experiment’s combinatoric background.

FPGA-4 and -5 are “singles” triggers that record events that are useful for estimating backgrounds and detector efficiencies. Each of these has a “*prescale factor*” that limits how many of each mode actually fires the trigger. For example, FPGA-3 has a prescale factor of 123, which means that, within a single spill, there must be 123 FPGA-3 events before it tells the DAQ to record one. This keeps certain high-frequency trigger modes from dominating the readout and causing undue deadtimes.

In addition to the v1495 FPGA triggering mechanism, the outputs of the hodoscope discriminators are also sent up to the control room to a set of NIM logic modules. These signals, however, do not describe each hodoscope paddle but instead are *ANDs* of the tops and bottoms of each plane. For example, if one or more paddles in H1T fires, then a positive signal is sent out indicating that H1T fired. These signals can be used to form a rudimentary trigger on situations such as 4- or 3-out-of-4 coincidences in the X- or Y-direction hodoscope planes. In general, NIM-1 is used for Y-coincidences, and NIM-2 is used for X-coincidences.

The third NIM trigger mode is of particular importance, as it is used for estimating a great deal regarding backgrounds and rate dependence issues. The NIM-3 trigger is called the “Random RF” trigger, as it is controlled by two different clocks beating against each other. When the two clocks are in coincidence, they will send a signal to the Trigger Supervisor to record the event corresponding to the next immediate RF bucket. This trigger allows the experiment to get a random sample of the events that are occurring at the spectrometer, notably without a bias that a selective trigger can incur.

2.9 Data Acquisition Systems

The whole of the SeaQuest experiment spans two sectors of the NM beamline (NM3 and NM4) and records not only detector data but also accelerator and atmospheric conditions. Designing a data acquisition (DAQ) for this experiment, therefore, requires a bandwidth and timing which cannot be provided by a single central computer system. SeaQuest has therefore employed four separate DAQ subsystems called “Main DAQ”, “Scaler DAQ”, “Beam DAQ” and “Slow Control”. The Main DAQ records the main detector information and the trigger timing. Scaler DAQ records various scaler readouts once per spill to reduce bias from the dead time of Main DAQ. Beam DAQ records information from a beamline Čerenkov detector read out by its QIE board. The Slow Control is a catch-all for everything else, from magnet current to radiation monitors that are read out once per spill.

2.9.1 Main DAQ

The MainDAQ is powered by a Jefferson Lab developed software package named CODA (CEBAF¹ On-line Data Acquisition). The Trigger Supervisor could receive up to 12 different kinds of triggers. The first four triggers (FPGA 1-4) can be prescaled up to 24 bits. The second four triggers (FPGA 5, NIM 1-3) can be prescaled up to 16 bits. The rest of triggers (NIM4, flush trigger, Beginning of Spill (BOS), End of Spill (EOS)) are not prescalable. The Main DAQ can configure which triggers are enabled at the Trigger Supervisor level and what the prescale factors are for each triggering mode. Once the TS receives a trigger, it will count the prescale factor and, based on how many triggers of that type have fired so far, it will decide if the Main DAQ will accept the trigger or not.

Once the TS accepts the trigger, it sends an “accepted trigger” signal to the other front end electronics, such as the detectors’ TDC readouts. The Main DAQ system is connected to the TS and the CPU’s in the VME crates via a local network. These CPU’s are known colloquially at SeaQuest as readout controllers, or “ROC’s”. The TS connects with each VME crate, and when the TS sends an “accepted trigger”, an interrupt signal is sent to each of the ROC’s. The ROC’s will then start reading the various front-end electronics (FEE).

The TS also sends a signal to the QIE of the BIM, and the QIE retains information regarding the ± 16 RF buckets around that triggered RF bucket to assist in investigating the beam quality around each of our event triggers. The output of the QIE is encoded into a scaler latch card format. If the beam intensity in RF buckets neighboring the trigger is higher than the user-select threshold, then the board will issue a veto signal to the TS to ignore this trigger.

The Main DAQ’s deadtime is considerable, as it has to communicate with the TDC’s, which have the longest copy-in-progress time of $32 \mu\text{s}$. The average TDC readout time is approximately 300 ns per 32bits (one hit), and as such, the slowest ROC which contains 7 TDCs has, on average, $150 \mu\text{s}$ deadtime. This deadtime is accounted for under the umbrella term of “*busy*” time and factored into the calculation of the “live protons” received by the experiment.

2.9.2 Scaler DAQ

The Scaler DAQ is used to ensure that the detector and trigger systems are working properly. The readout controller CPU is an MVME5500 [63] with four scalers. One of these scalers is triggered by the coincidence of a 7.5 kHz gate generator and the beam spill signal. This records the 7.5 kHz response of the hodoscope arrays. The other three scalers are triggered by the End and Beginning of Spill (BOS and EOS) signals

¹The old acronym inside of an acronym bit. CEBAF: Continuous Electron Beam Accelerator Facility

and record spill-level rates. Data collected by these spill-level scalers are the number of times each Main DAQ trigger is satisfied (both raw and accepted), the intensity of the beam, and the rates of the hodoscope arrays. The readout of this VMEbus DAQ is done using CODA very similarly to the Main DAQ but on a completely different system. An independent program analyzes the data in real-time in order to monitor the performance of the detectors and triggers, as well as the quality of the beam.

2.9.3 Beam DAQ

The Beam DAQ is composed of a Čerenkov detector in the proton beam (the Beam Intensity Monitor discussed earlier in this chapter), a QIE board, and a custom C++ program to control their operation and read out. The Beam DAQ is responsible for recording the 53 MHz structure of the beam, i.e. the intensity of each RF bucket. Its calculation of the 53 MHz duty factor $DF = \frac{\langle I \rangle^2}{\langle I^2 \rangle}$ is the primary measure of beam quality that accelerator operators use for tuning. In short, the duty factor is a measure of the *uniformity* of the intensity of the beam. If we were to ignore the buckets intentionally left empty, and if it were to be the case that every RF bucket had the same number of protons, the 53 MHz duty factor would be 100%.

There four types of data recorded by the QIE have already been covered in the Beam Intensity Monitor section: QIESum, beam inhibit, busy time, and the ± 16 RF bucket intensities. The Beam DAQ commences read out of each of these blocks of data when the EOS signal is seen. The block of QIE data for all buckets is about 300 MB. To read this much data in time to analyze it and be ready for the next spill, the DAQ program utilizes multithreaded processes. Three threads are used to read the data from the QIE board's three ethernet chips, and up to eight threads are used to analyze the data. Analyzed data is displayed on a public web page so that shifters and accelerator operators can monitor the quality of the beam (Fig. 2.14). The fully processed data is also written out to tab-delimited ASCII files, which are archived and also uploaded to our online MySQL server.

2.9.4 Slow Control Readout

The Slow Control readout system is an aggregation of many different readouts from various parts of the experiment. This is also the umbrella term for other critical monitoring and bookkeeping tasks that must be performed. The one common thread to them all is that they only *need* to be read out / performed once per spill in order to give perspective on the conditions of the experiment for each given spill. The term “*Slow*” in Slow Control is adapted from the term “Slow Spill”, which is used to characterize the method of beam delivery to SeaQuest. The Slow Control is run by an array of Python, Perl, and Bash scripts that, for the most part, run independently. These scripts primarily interact with and read out from EPICS, ACNET,

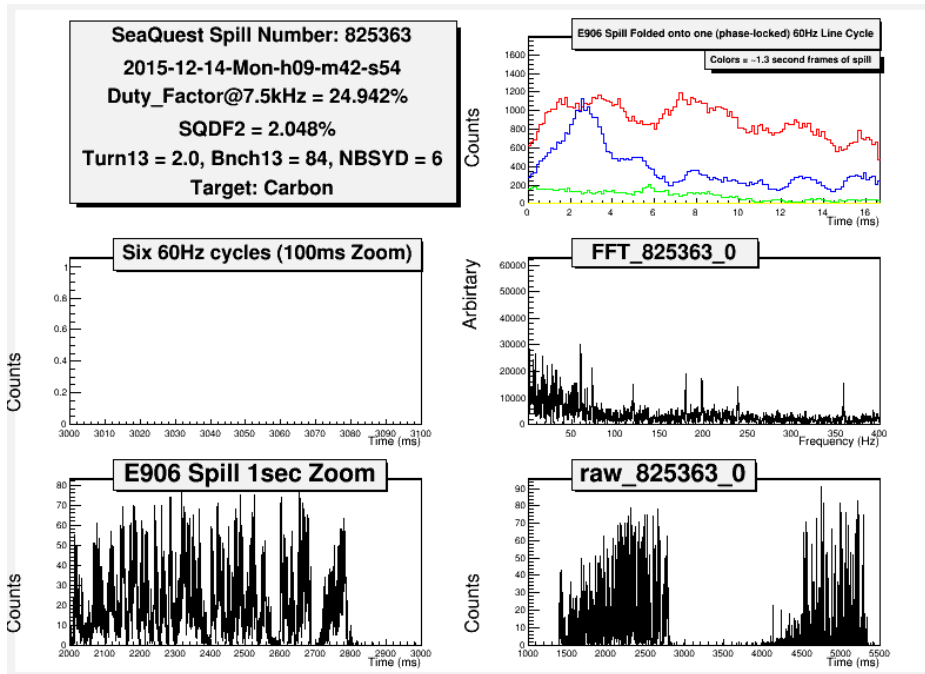


Figure 2.14: A standard Beam DAQ display showing beam characteristics and spill data. The bottom two plots show the output of the QIE on a ms scale.

and the local Spill Counter.

EPICS, ACNET, and Kiethley

The EPICS software package (Experimental Physics and Industrial Control System) is used for communicating any number of variables across the experiment’s local network. EPICS is used for feeding readouts from many different sources and formats, and making them available to any other system that may have need of certain information, but does not natively communicate with a different system’s format. Bridging this gap, EPICS is able to gather information regarding the target, beam, and environmental data into one place.

The EPICS server is installed on SeaQuest’s target control computer, where the target position, pressures, temperatures, and proximity sensors are monitored and controlled. These values are critical to track and monitor, as they affect what target is in position at any given spill, along with target temperatures and pressures, from which the cryogenic liquid targets’ densities can be calculated. Additionally, there is information regarding the operational mode of the target table, such as whether or not the target table is being manually controlled or if it is automatically repositioning itself. In Slow Control parlance, information from the target computer that is sent to the EPICS server is labeled as “Target” type data.

Several values are read into the EPICS server from the Accelerator Network (ACNET), which monitors

many critical systems. Beam intensities (S:G2SEM) and FMAG / KMAG currents are two of the major factors that must be included for any analysis. There are additional readouts with esoteric names that can record anything from beam pipe vacuum to upstream radiation monitors. It is important to note that the FMAG and KMAG variables from ACNET are read out to the EPICS server every 3s, particularly because it is unsafe to deliver beam to the sensitive detectors without FMAG being fully powered. As such, the FMAG current is read into the beam interlock system. In the Slow Control feed, this type of information is labeled as “Beam” data.

Finally, there is a Keithley multimeter in SeaQuest Hall near the electronics racks by Station 3. This device has the responsibility of relaying information regarding ambient temperature, air pressure, and humidity up to the control room. This information on its own is not typically used for any analyses, but in the case that, for example, a rise in humidity results in any electrical device failures, the monitoring data exists to diagnose it as such. This data is gathered by a ‘kscan’ program that is run on the target machine, and its operation typically takes 11s to return a result. Data from the Keithley multimeter that is put into the Slow Control feed is labeled as “Environment” data.

Spill Counter

Perhaps one of the most critical components of the data acquisition is the coordination between the many sources of data in the DAQ. The Main DAQ, Scaler DAQ, Beam DAQ, and Slow Control do not share a single clock or network, so some measures need to be taken to ensure that data from different sources regarding a single spill’s worth of data actually all correspond to the same spill. The concept of a locally maintained Spill Counter arose as a solution to this problem, and it is handled in the form of a simple Python script and a flat text file containing a single integer.

The enacted solution begins with placing an initialized spill counter (*spillID*) value (currently $O(10^5)$) into an otherwise empty text file at a specified location. Then, when the ACNET readout receives a BOS/EOS signal from AD, a simple python script performs the following:

- Read the current spillID from the file
- Broadcast this value over the EPICS server
- Insert a single ASCII-encoded event into the Main DAQ and Scaler DAQ CODA streams containing only that spillID number
- Waits for 20s while other scripts may want to use the current spillID value to connect its value to the spill that just ended

- Updates the file to contain spillID + 1

This simple approach has been very effective in its usage so far, and the spillID incremented via this method has been shown to be accurate upon scrutiny of data contents across DAQs. This can be easily done by looking at a series of full and empty spills, as it is obvious which spills had a nominal beam spill and which did not by investigating the output of each DAQ. The weakness lies in the fact that this system relies on a single file, which is susceptible to any number of file access and timing issues. At some point in the future, it would be advisable to have all spillID-using systems to read the current value from the broadcasted EPICS server value instead of having several applications attempting to access and read the single file (even if just read-only).

Python Slow Control Script

With all of this going on to aggregate data from several sources, there is one final script that takes this aggregated data and makes it readily available for the rest of the DAQ systems. A `read_slowcontrol.py` script, once every 60s supercycle, puts the following together into a single ASCII text file:

- Clock time (“vxticks”) of the Trigger Supervisor
- Spill count data
- Beam data
- Environmental data
- Target data
- Timestamp

Once all of this is written to a file, this is sent to the Main DAQ and Scaler DAQ to be integrated into the full CODA file as a plain text event with a simple, tab-delimited format.

With data in all these several formats, some coordinated effort has to be put into processing and storing it. Further, with these disparate sources, there must be some coordinating factor that unifies it so that users can associate any one piece of information with the rest. The following chapter discusses the solutions implemented by SeaQuest in the methods and choices of processing and storing all of this data.

Chapter 3

Data Productions

Data analysis is a highly dynamic activity – particularly for a nascent, large-scale experiment. As a result, having a flexible and scalable data management and analysis framework is of particular importance. Different data components may be needed at different points of study, brand new readouts may be added in the middle of data taking, and analysis styles themselves may differ from user to user – all of which are either facilitated or impeded, depending on the flexibility of the framework implemented.

In an attempt to provide such a framework to its users, SeaQuest makes use of the capabilities of an array of MySQL servers, leveraging relational database technologies. While MySQL and database systems are nothing new, even to physics experiments, they are not widely adopted among the physics community. Often, in the cases that they are, the resulting product tends to be ineffective or inefficient.

During my tenure at SeaQuest, I was primarily in charge of the effort to decode the raw, serialized data from the data acquisition systems and store them in a manner in which it could be readily accessible by collaborators worldwide. Beyond that, an emphasis was placed on curating the tracked physics data and preparing it in such a way that facilitates quick and simple analysis from many different analytical software packages.

In this chapter, the details regarding the processes and formats that define the online and offline data productions will be described in some technical detail. The technologies used to perform this were a combination of compiled C/C++ codes, MySQL, ROOT, and the Fermilab Computing Division’s primary computing grid. The end product provided to all collaborators was in the form of unified “schemas” (or, groups of tables) possessing all of the combined *run*, *spill*, *event*, *slow control*, and *tracked* information, along with specific data quality flags in place to assist in the selection of good, analyzable data. Finally, there is a retrospective discussion on the scalability and best practices, followed by speculation on how future data management technologies might be used at SeaQuest and other experiments.

3.1 Raw Data Processing

The three raw outputs of the data acquisition systems, as described above (and in Chapter 2) are Main DAQ CODA files, Scaler DAQ CODA files, and Beam DAQ ASCII files. Each raw data file corresponds to the data taken from certain subsystems over approximately one to two hours of running time. These three types require varying degrees of de-serialization, parsing, processing, and storage – a process as a whole defined as *decoding*.

All raw data files are backed up to long-term tape storage (managed by FNAL Computing Division), and the decoded and processed data gets stored on one of four MySQL servers to be used for analysis by the collaboration. Data is also output to ROOT files for ease of use by independent tracking programs and those who prefer navigating data in the raw ROOT TTree format.

3.1.1 CODA Event Format

A single CODA file can be described as being a chain of events. These can generally be divided into “CODA Events” for those related to the CODA file itself, and “Standard Physics Events” when containing experimental data read out by CODA and its subsystems. Each event can be represented by an array of unsigned 32-bit integers. When using the CODA EVIO (event input/output), one whole event is read out into such an array, which can vary in size from 6 to $\sim 50,000$ integers long, depending on the type of content.

CODA Events primarily correspond to the actions taken by a shift-taker using the *requi* application on the data acquisition computer. The types of events that fall under this category are Sync, Prestart, Go, Pause, and End. It should be noted that the Sync and Pause events do not generally get used at SeaQuest. These provide markers in the CODA file of what commands were given to CODA through the course of taking data.

Before data taking can begin the readout controllers and the trigger supervisor need to be loaded with a designated firmware in order to operate properly. Before a CODA file is started and created, this “Download” is performed, and then CODA is ready to transition from the downloaded state to the prestarted state. The CODA file is then initialized with a prestart event, whose format can be seen in Figure 3.1. The most important information stored in the prestart event is the run number, which is used universally through the experiment.

Once in a prestarted state, a “Go” action is initiated to transition it into data taking mode. This action generates a Go event. If ever paused, a Go event would be created once data taking was resumed. Go events contain the current time and the number of events in the run so far, also seen in Figure 3.1. It should be noted that the “time” stored in these events is not trusted, as the data acquisition computers are not synced




Figure 3.1: (Left) The Prestart event, with event type 17; (Right) The Go event, with event type 18[64].




Figure 3.2: (Left) A standard physics event, with various event types, filled with a series of data banks; (Right) The ROC data bank format[64].

to an official external source.

The last CODA Event of importance is the End Event. This contains the number of events in the run, but primarily serves as a marker to any programs reading the CODA file that the EOF (end of file) has been reached, and to commence any post-run processes.

Most of the events in the CODA file, however, will be “Standard Physics Events”. The contents of these can be varied, but can be characterized generally as having a header containing the total length of the event (in number of 32-bit integer “data words”), followed by several “data banks”, which are typically filled with a series of ROC outputs (Fig. 3.2). The contents and formats of the ROC data banks are completely controlled by the ROC programming, and whatever set of 32-bit words the ROC has to read out, it will store in its “data words” block.

3.1.2 ROC Data Bank Formats

The data within the ROC data banks have their own definition of encoding, and it differs from event type to event type. Of the types most commonly used, there are scaler, QIE, slow control, and “JyTDC” formats. These formats were designed and implemented gradually and as-needed through the development of the SeaQuest data acquisition. Through these iterations, there was a consistent effort spent devising an output format for the TDCs that was as compressed as possible while performing very quickly (to reduce deadtime).

The first TDC format devised was compact, but required some computation and caused unnecessary dead time. After that, signals were encoded into a non-zero-suppressed format where a binary string such as 00011111111110000 would represent a signal, with each 1 representing 2.5 ns of contiguous signal. This then required writing too much data, which still caused some dead time, and the file size of the raw data would become bloated. Eventually, a quick, compressed, complex format called “JyTDC2” was adopted and became the standard for the MainDAQ encoding.

3.1.3 Decoding Raw Data

The CODA file decoding process is nearly identical in the cases of the Main DAQ and Scaler DAQ outputs, and only differ by content; the Scaler DAQ reads out specific scaler data, and the MainDAQ contains TDC readouts of the detectors. For each one to two hour *Run*, the CODA files can be well-described as the following sequence of events (and the data they contain):

1. Prestart Event (Run data)
 2. Begin Spill Event (Spill data, Scaler readout)
 3. Many Physics Events (Event data, TDC readout)
 4. End Spill Event (Spill data, Scaler readout)
 5. SlowControl Event (Slow control readout, Spill ID readout)
 6. Spill Counter Event (Spill ID readout)
- ...(Repeat 2-6 for each *Spill*)

The decoding programs use C and C++ in conjunction with Jefferson Lab’s CODA I/O library[64] to read these events and parse them according to their individual formats. Data from these CODA events are decoded and placed into hierarchical categories.

Run Level Data

Run-level data contains data and metadata pertaining to the entirety of the run that is recorded. At the time of the Prestart Event, the date and time of the run are stored, along with a readout of the specific settings of all non-trigger TDC boards. Settings regarding which triggers are enabled and what each of their prescale factors is also stored in the run-level data.

After the End Event is encountered, metadata is aggregated and stored regarding such items as the number of chamber hits, the types of triggers that were fired, the target positions used, average magnet

currents, and other useful metrics. All of these fields of settings and aggregated values can be used later on to quickly determine if a run is analyzable, or if it has characteristics that are desirable for specific niche analyses.

Spill Level Data

The *Beginning of Spill* (BOS) and *End of Spill* (EOS) events bookend the set of physics events for a given spill. At each BOS and EOS events, the 140 MHz VME scalers are read out. At the beginning of the spill, all scalers are zeroed out, and then read out again after the spill has ended.

Slow Control events are read out between spills, which contain data regarding the current spill identifier number, target systems, beam and radiation monitors, and environmental readings.

The spill identifier (*spillID*) is what is used to synchronize the data together across various data acquisition systems. As such, the *spillID* is read out redundantly in both Slow Control and Spill Counter events (which contain only the *spillID* value) to ensure that the data is appropriately labeled.

When the End Event (for the whole CODA file) is reached, the independently-recorded Beam DAQ data (recorded in an ASCII file) is read and stored with the rest of the Spill-level data.

Event Level Data

For each spill, $\sim 3k$ events are triggered to be recorded, though this number can vary greatly with the particular beam settings of a run. With each event, three types of information are stored: the trigger which fired the event, a measure of the beam intensity per RF bucket, and the full detector readouts. The detector readouts require the most processing of all the rest of the data. The CODA files contain the hardware addresses of each detector *hit*, along with a *TDC time*. The following steps briefly summarize the processing steps:

1. Mapping: Map the hardware address to a detector name and detector element number
2. Timing: Classify hits as in-time or not and calculate *drift time* from TDC time
3. R-T (time-to-space): Translate *drift time* to *drift distance*
4. After-Pulse Elimination: Remove hits that result from signal reflection and other electronic artefacts
5. Trigger Road Reconstruction: Use *v1495* TDC hits to reconstruct possible trigger roads that may have fired
6. Hodoscope Masking: Remove drift chamber hits that have no adjacent hodoscope hit

7. Trigger Road Masking: Same as hodoscope masking, but only using hodoscopes from reconstructed trigger roads

This fully processed data is then stored into one the experiment’s MySQL databases and/or a ROOT file.

3.2 Online and Offline Production Processing

There are two running modes of raw data processing: on-line and off-line. These are designed to suit two needs: (1) the need to monitor the experiment and the detector readouts, and (2) to perform a full analysis of the data taken.

3.2.1 Online Monitoring and SeaScape

For the on-line mode of productions, all Run- and Spill-level data is decoded, but only 1-in- n Physics Events are processed, where n is typically between 3 and 15, depending on the event data taking rates. This “*sampling mode*” is used in order for the decoding to reliably keep up with even high-intensity beam data. This data pipeline is important for on-line monitoring of the experiment by shift crews. With this up-to-the-minute stream of data, any abnormalities in features such as duty factor, wire maps can be observed directly and acted upon quickly.

In 2014, Dr. Sylvester Joosten, a then-UIUC PhD student working on the HERMES experiment, collaborated with SeaQuest to create a robust, responsive web page for online monitoring. This project, named *SeaScape*, was built using JavaScript, Python Flask, and used a Python MySQL API for database interfacing and data handling. The result of the combined efforts was a full-featured, globally-available, password-protected web app hosted on a Fermilab virtual server. Current capabilities of SeaScape consist of wire maps (histograms of hits) for all groups of detectors (Fig. 3.3), time series of ScalerDAQ, BeamDAQ and Slow Control (Fig. 3.4) values, and even a note-taking section for use by those observing remotely or those on shift.

Any run schema may be selected for displaying, and two distributions from two different schemas may be compared against each other. A spill selection slider bar is provided so that a user may step through the data for various ranges of spills. An export feature is provided for users to output the displayed figures to PDF files.

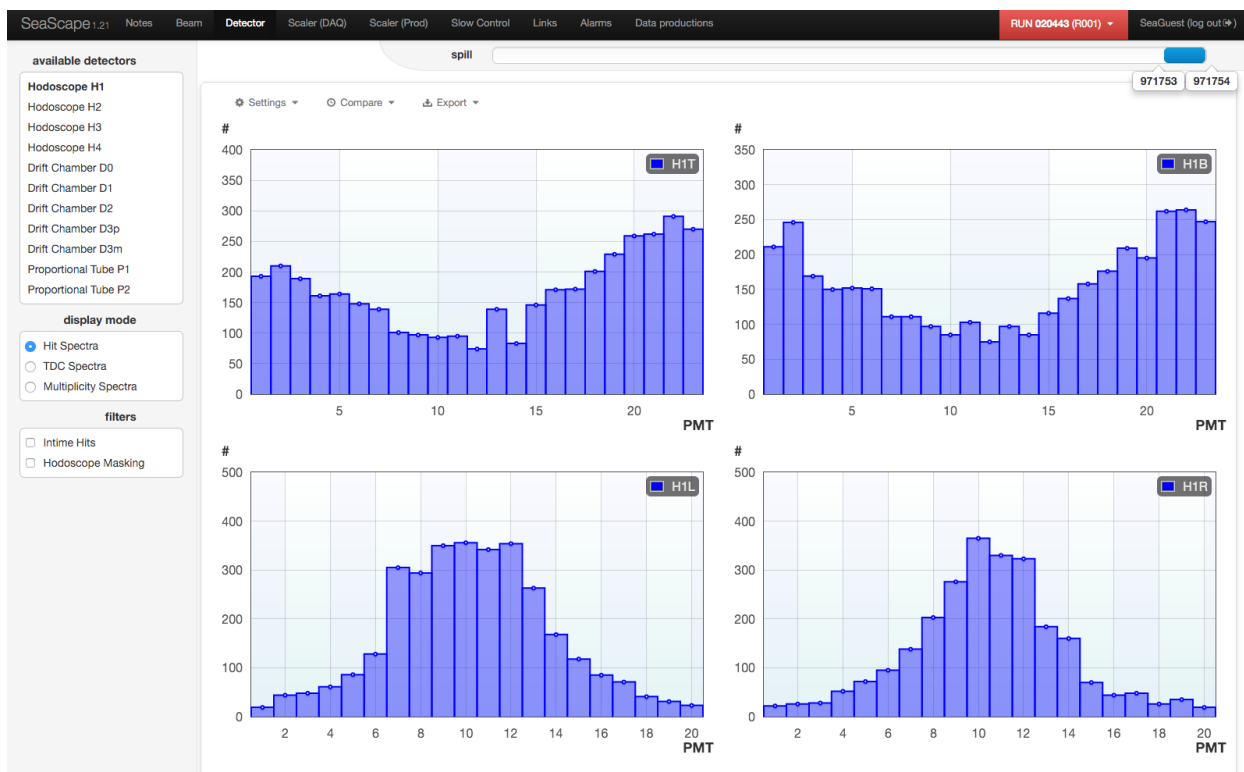


Figure 3.3: The SeaScape Online Monitor displaying the hit distributions for the Station 1 hodoscopes. Run selection is in the top red button selector, and the Spill selection is in the bar just below the top.




Figure 3.4: SeaScape displaying target position data for a sequence of spillIDs. Up to six different values can be shown at once against each other (scale-able for easier comparison) on each of the Slow Control, BeamDAQ, ScalerDAQ pages.

3.2.2 Offline Batch Processing

For off-line, or “batch”, productions, a large group of categorically similar runs is defined, and the chain of production processing is initiated. The steps of this process are generally emphdecoding, tracking, archiving, and merging. The logical blocks of runs are typically defined by “roadset”, or, the version of the set of trigger roads used in the L1 trigger matrix. Differences in these trigger matrices can affect the distributions to a degree commensurate with the difference in the set of roads that are triggered on by the FPGA-1 trigger. As such, each block of data using a given roadset should be analyzed individually and then have its results combined with those of other roadsets via a weighted average.

The decoding and tracking are performed on Fermilab Computing Service’s FermiGrid, which provides the computing resources necessary to process hundreds of runs simultaneously. The tracking has also expanded its capabilities to run on the Open Science Grid, which is a grid computing network that pools together computational resources from national laboratories and universities throughout the world[65]. This allows the intensive reconstruction algorithms to perform its tracking on thousands of grid nodes at once.

A single decoding job submission will output the processed data to one of the four available MySQL servers and also to a ROOT file. The processed data in ROOT form at this stage is called the “digit” data. Whether or not the decoding sends Hit and TriggerHit information to the MySQL schemas is an option that can be set at the time of batch processing. The current operating mode is to perform batch processing, sending Hit/TriggerHit data to MySQL for only 1-in-50 runs. This allows users to capture a glimpse of full hit distributions for runs within a roadset while preventing a single pass of the experiment’s data from filling up all available data storage space.

After this step, jobs will be submitted to run one or both of the two tracking programs based on the ROOT file and/or the MySQL data. Once the tracking is completed, the ROOT file is archived on the Fermilab BlueArc NAS backup system for storage and for tracking.

Upon the completion of decoding and tracking of a specified range of runs, all of their Run-, Spill-, and Event-level data, along with its tracked data, is combined into a single *merged* schema. These *merged* schemas are mirrored across all four of the MySQL servers for optimal redundancy and availability.

3.3 RDBMS Data Structure

The processed data is primarily stored in MySQL Server 5.1 databases. MySQL is an open-source **R**elational **D**atabase **M**anagement **S**ystem (RDBMS) developed by Oracle that is well-suited for the storage and responsive querying of hierarchical data.

Why was SQL chosen for the data storage? The ease of use in querying and the flexibility were important factors here. Querying using the Standard Querying Language (SQL) provides an English-readable (and thereby easier-to-learn) way to formally ask the database for what you're looking for. Flexibility stems from the ability to *update* existing data with best-known calibrations and the ability to modify the structure of the stored data *in situ*. *Why MySQL of all the SQL technologies?* The short answer there is speed and support. There are several RDBMS to choose from, but at the time of deciding in 2010, MySQL had exhibited a slight speed edge over most of the competition (such as PostgreSQL). What truly made the decision was the near-universal support that both Oracle and all other languages supported MySQL. Nearly every coding and analysis language out there supports an interface to MySQL.

Each run is decoded into its own schema and contains its own instances of all tables of a specified design. The tables are all *join*-able to each other by sharing *foreign keys* with each other in the form of the *runID*'s, *spillID*'s, and *eventID*'s. The contents of the tables are *indexed* in such a way that *joins* and queries gain a speed performance boost, but this comes at the cost of disk space.

3.3.1 Querying Language

The data on the server is world-wide accessible and can be queried using SQL, which is styled very much like a logical, readable sentence, and therefore has a relatively shallow learning curve as compared to other programming languages. For instance, in order to select the eventID, spillID, and dimuon vertex (*dz*) position from the *kDimuon* table for cases where $4.2 \leq M_{\gamma^*} \leq 10$ GeV, one need only execute the query:

```
SELECT eventID, spillID, dz
FROM kDimuon
WHERE mass BETWEEN 4.2 AND 10;
```

where whitespace is added only for readability and the query keywords are capitalized, but not case-sensitive.

The key benefit of using a relational database is in relating two tables. Any two tables can be *joined* on each other via overlapping fields, and that total result itself can be joined again. In the following query example:

```
SELECT dimuonID, dpz, SQRT(POW(dpx,2)+POW(dpy,2)) AS dpt, mass, costheta, dutyFactor53MHz
FROM kDimuon
INNER JOIN Spill USING(spillID)
INNER JOIN BeamDAQ USING(spillID)
WHERE
mass BETWEEN 4.2 AND 10 AND
```

```
Spill.dataQuality = 0;
```

exhibited are the following actions and features:

- The kDimuon table is joined to the Spill table using the common spillID field
- An “INNER JOIN” is one in which the joined field must have matching values in both tables in order to return a value
- The kDimuon-Spill joined result is, in turn, joined on the BeamDAQ table (again sharing the spillID field) to gain access to the associated duty factor field
- The SQL math library is used (SQRT() and POW()) functions) to calculate p_T from dpx and dpy
- On-the-fly renaming of fields (here, “dpt” is the name given to a calculated field)
- Use of the Spill data quality criteria discussed in depth in Section 3.4

And so with a modest familiarity with SQL and knowledge of what tables exist and what fields exist in what tables, an analyzer is able to very quickly put together a query to access a very *specific* set of data that they are looking for. For more on SQL, tutorials, how-tos [66], and very well-supported documentation from Oracle [67] are ubiquitous online for anyone wishing to pick up this language in a short amount of time.

What’s more, there is very broad support for interfacing with SQL servers in general (programming APIs), and so the resultset of the queries can be funneled directly into any analysis code in any programming language due to the large array of MySQL API’s available. This was one of the key factors in choosing SQL to host the analysis data. Those wishing to use ROOT, Excel, Mathematica, Matlab, R, or Python could do so, using a single interfacing language to access the data.

3.3.2 Atomic Data Schema

With a basic knowledge of the standard querying language, one then only needs to know where to find what data in order to query it. A summary of the atomic schema layout can be found in Figure 3.5. For every run that SeaQuest takes, when it is fully processed, a schema is created with all of the tables listed in this figure. While the raw data tables have already been sufficiently discussed in the previous sections, some new tables are the calibration tables (used to perform all the timing and mapping mentioned in Section 3.1.3), the BeamDAQ table which is directly uploaded from the BeamDAQ output, the Slow Control data which is the categorized data read out from the EPICS server, and the Tracked Data which is uploaded directly by the tracking software.

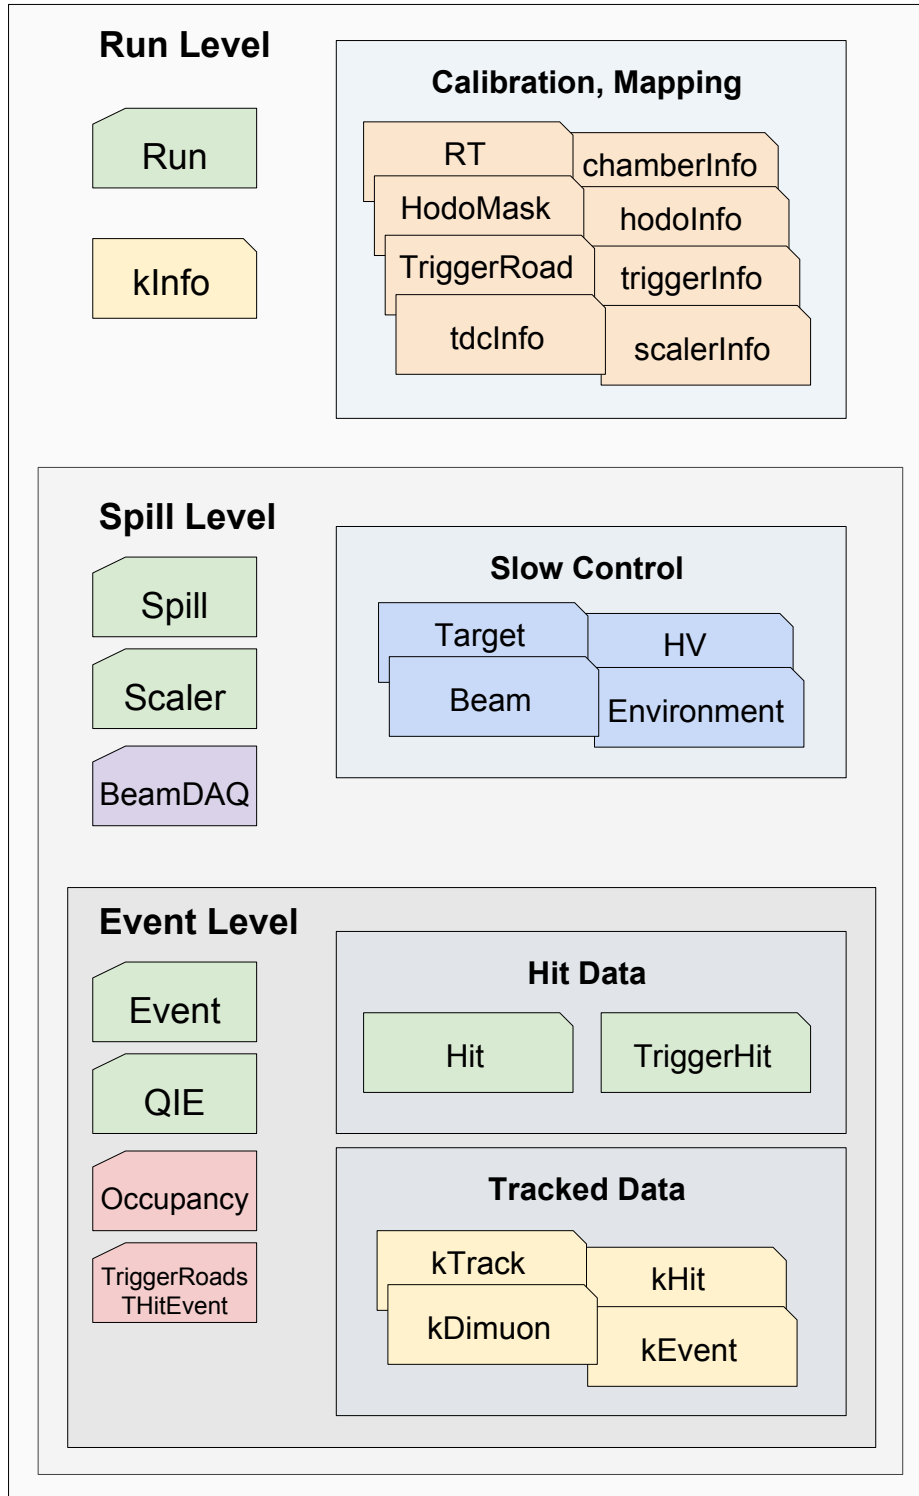


Figure 3.5: A visualization of the hierarchy and categorization of the data stored in the MySQL tables at SeaQuest. Data can be grouped into Run, Spill, and Event hierarchies, and categorized into raw data (green), calibrations (orange), BeamDAQ data (purple), slow control data (blue), reconstructed data (yellow), and aggregated metadata (red).

3.3.3 Production Management

Without a naming scheme, a database environment can become untenable as no one will know where to find what they're looking for. A set of prefixes and a template for standard run schema names were advised so that as little referential material would be needed as possible. The schemas kept on the servers are grouped into the following:

- **run_XXXXXX_RYYY**: *XXXXXX* is the 6-digit, zero-padded run number, *YYY* is the 3-digit, zero-padded production revision number. All runs decoded, processed, and tracked keep their data in these schemas.
- **merged_roadsetXX_RYYY_VZZZ**: *XX* is the 2-digit trigger roadset version number, *YYY* is the aforementioned production revision number, and *ZZZ* is the version of the merged production. This is the aggregated analyzable data from a large (~ 1000) number of runs combined into a single set of tables.
- **mc_***: The prefix used for Monte Carlo simulated data. Naming conventions defined and kept self-consistent by the Monte Carlo production manager.
- **geometry_***: The prefix used for experiment geometry definitions and survey numbers. Naming conventions defined and kept self-consistent by the geometry production manager.
- **user_***: A prefix reserved for users to create, store, and manipulate their own data freely.

As table definitions were revised and new standards were imposed for the standard *run*-level productions, the production revision number was incremented (e.g. *R003*, *R004*, *R005*). As new tracking results were generated and new *merged*-level definitions were imposed, the merged version number was incremented (e.g. *V001*, *V002*, *V003*). Users are encouraged to keep their own tables of signal dimuons for their own analysis in their own *user_username_descriptive_name* schemas for some analytical permanence that doesn't require re-running of their full analysis codes.

3.4 Data Quality

In Section 5.2, the number of data quality cuts is substantial. Over the course of analysis, many small issues were discovered that required a check for. Many of these values and criteria would change from roadset to roadset. As a result, performing a data quality check on the data became very cumbersome and error-prone for analysts. It was proposed to simply remove the bad quality data from the tables, but there were certain

circumstances where users might wish to *relax* some of the criteria. As a result, the use of a bit pattern, or bitmask, was implemented to summarize the data quality classifications.

3.4.1 Bit Pattern Fields

A single 32-bit integer can be represented as a string of 32 0s and 1s when represented in binary. With a bit pattern defined, each of these can act as a flag for 0 meaning “passed the quality check” and 1 meaning “failed the quality check.” With this being the case, it becomes a simple matter to select data which passes all data quality criteria by requiring `dataQuality = 0`, since any check failing will cause the entire integer to be non-zero.

This also allows users to relax criteria by use of bit-wise logic. If we have an 8-bit data quality word, and the 3rd data quality bit ($2^3 = 8$) is 1 (0b00001000), but we wish to relax the criteria for the 3rd bit, one could still select the data that passes all other checks with a bit-wise *and* and a *twiddle* operator:

```
~0b00001000 = 0b11110111
0b00001000 & 0b11110111 = 0b00000000
```

So, in practice, if we want to relax the 3rd bit criteria (0b00001000 = 8), then in place of a `dataQuality = 0` cut, one would add a `dataQuality & ~8 = 0`.

3.4.2 Spill Data Quality

Within the Spill table, there is a field named “dataQuality”, which is a 32-bit integer. Many of the 32 bits have been defined to reflect certain Spill quality criteria that were summarized in Table 5.3. Blocks of bits are reserved for certain categories of quality checks such as “Beam,” or “DAQ/Trigger,” though a data quality criteria have not been assigned to every bit. The Spill dataQuality bit pattern definition can be found in Table 3.1.

3.4.3 Other Data Quality Bits

The dataQuality bits paradigm is used on the Event- and Hit-level data, though to a lesser extent. At the event-level, there are cases where the QIE experiences a readout error, and there is no intensity information for that event as a result. There are also cases in which the v1495 TDCs experience an error, and there is no capability to perform the precise “RF in-time” flagging using the triggering hodoscopes and their RF timing. The 0th bit has been reserved for indicating events with excessively high occupancy in any one detector group, but no consensus could be reached regarding how high of an occupancy was too large. As

N^{th} bit	Category	Description	Roadset 57 & 59	Roadset 61	Roadset 62	Roadset 67	Roadset 70
0	Beam	Duty Factor	[15,60]	[15,60]	[10, 60]	[10, 60]	[10, 60]
1	Beam	G2SEM	[2e12, 1e13]	[2e12, 1e13]	[2e12, 1e13]	[2e12, 1e13]	[2e12, 1e13]
2	Beam	QJESum	[4e10, 1e12]	[4e10, 1e12]	[4e10, 1e12]	[4e10, 1e12]	[4e10, 1e12]
3	Beam	FMAG Current	> 1000 A	> 1000 A	> 1000 A	> 1000 A	> 1000 A
4	Beam	KMAG Current	> 1000 A	> 1000 A	> 1000 A	> 1000 A	> 1000 A
5	Beam	Undefined					
6	Beam	Undefined					
7	Beam	Undefined					
8	Target	Undefined Targ. Pos.					
9	Target	Targ. Pos.	∈ [1, 7]	∈ [1, 7]	∈ [1, 7]	∈ [1, 7]	∈ [1, 7]
10	Target	Undefined					
11	Target	Undefined					
12	Target	Undefined					
13	Target	Undefined					
14	Target	Undefined					
15	Target	Undefined					
16	DAQ / Trigger	Inhibit	[4e9, 1e11]	[4e9, 1e11]	[4e9, 2e11]	[4e9, 2e11]	[4e9, 2e11]
17	DAQ / Trigger	Busy	[4e9, 1e11]	[4e9, 1e11]	[4e9, 1e11]	[4e9, 1e11]	[4e9, 1e11]
18	DAQ / Trigger	AcceptedFPGA1	[1e3, 8e3]	[1e3, 12e3]	[1e2, 6e3]	[1e2, 6e3]	[1e2, 6e3]
19	DAQ / Trigger	AfterInhFPGA1	[1e3, 3e4]	[1e3, 1e6]	[1e2, 1e4]	[1e2, 1e4]	[1e2, 1e4]
20	DAQ / Trigger	Accepted/AfterInh	[0.2, 0.9]	[0.0, 0.9]	[0.2, 1.05]	[0.2, 1.05]	[0.2, 1.05]
21	DAQ / Trigger	TSGo	[1e3, 8e3]	[1e3, 12e3]	[1e2, 6e3]	[1e2, 6e3]	[1e2, 6e3]
22	DAQ / Trigger	BOS and EOS exist					
23	DAQ / Trigger	MATRIX1 Settings					
24	Decoding	Duplicate values					
25	Decoding	Missing values					
26	Decoding	Problematic Spill Range					
27	Decoding	Undefined					
28	Decoding	Undefined					
29	Decoding	Undefined					
30	Decoding	Undefined					
31	Decoding	Undefined					

Table 3.1: The Spill table dataQuality bitpattern definition.

Table 3.2: The Event table dataQuality bitpattern definition.

N th bit	Description	Comment
0	High Occupancy Check	*Not Implemented
1	v1495 Readout Problem	
2	QIE Readout Problem	
3	No in-time RF TriggerHit	No possible RF-based timing

Table 3.3: The Hit table dataQuality bitpattern definition.

N th bit	Description	Comment
0	In-time	
1	Hodoscope Masked	Taiwan and v1495 hodos used for masking prior to runID 11795, only v1495 hodos used from runID 11795 on.
2	Trigger Road Masked	Trigger roads reconstructed from in-time v1495 hodo hits only.
3	After-Pulse Check	Only the first in-time hit per element per event passes this check.

of now, the bit is reserved but unused.

For the Hit-level, requiring `Hit.dataQuality=0` will return all hits that are (0) in the in-time window, (1) are masked by a nearby hodoscope, (2) are masked by a hodoscope that could have possibly fired the trigger for that event, and (4) is the original pulse for that element for that event. This pares down the hits dramatically, which is useful when dealing with the Hit table, which can measure in the tens of millions of entries.

3.5 Discussion and Retrospective

In building the data storage servers, designing the tables and schemas, and building the code to process the raw data into an analyzable state, a healthy amount of insight has been gained regarding the abilities, limitations, and best practices. Here, some thoughts are shared on some of these matters after looking back on several years of development and use.

3.5.1 On Scalability

The largest obstacle experienced was the scalability of the full production processing chain. When the first MySQL servers were set up and the data format was established, it became quickly clear that *all* operations that were required to operate the data could be performed on the data *within* the database itself without having to *extract*, transform, and load it back up. With decoding and tracking, all steps could be executed using SQL commands, which was very practical and useful. When faster performance of said decoding and

tracking was sought, a parallelization approach was pursued. Once these operations were able to run in parallel on a MySQL server, scaling issues began to manifest.

Despite the specifications and hardware provided, even fully equipped servers are, at the end of the day, a single computer capable of handling only so much CPU and reading/writing with a certain I/O. When performing mathematical or transformative operations on datasets that are on the order of 10 GB, one will reach the CPU limits of most machines. Likewise, a single run will contain tables measuring 10+ GB in size. While this is not significant on the scale of the 10-25 TB servers used, it is certainly problematic as one *scales* the operation to the ~6000 runs that are typical of the SeaQuest corpus of data at the time of this writing. This corresponds to *at least* 60 TB of storage, which is untenable given the current storage capabilities of the four SeaQuest servers (currently 52 TB total).

So, what can be done about these limitations? This was the question that was asked and addressed as the analysis pipeline was developed and these realities manifested. Two simple answers are (1) to start using the servers for fetching and storage and not so much for large-scale transformation and calculations, and (2) the regular removal of unused data.

One of the strengths of using the MySQL server is that you can *very quickly* select the data you're looking for, pending good *indexing*. Extracting an exact slice of data and then performing higher-order steps locally is a paradigm that scales much better than every user and automatized job attempting to perform their tasks entirely on the server.

As time went on, it was found that, aside from online monitoring and some specially designated runs, the Hit and TriggerHit data was only used by the tracking software. And so, by only storing online *sampled* data and also offline 1-in-50 runs' Hit data in the database, this alleviates much of the data storage burden. The rest is archived in the *ROOT* TTree format specifically for the tracking software to use.

With these measures taken, the behavior of the data processing framework began to scale very well, even under heavy parallel loads. At the time of this writing, the existing framework can process 100 runs/server at 2 hr/run, and with four servers, can process the entirety of the dataset used in this thesis result in a matter of ~3 days (tracking not included). This has been demonstrated and repeated with the processing of production revisions R004 and R005.

It should also be noted here that, for added scalability and greater storage capacity, that the MyISAM storage engine of MySQL (used at SeaQuest) has a compression utility called `mysampack`. With this utility, a table and its indexes can be *reduced in size by 80%*. The only drawback to this is that the data becomes read-only once compressed. Flexibility has been a desirable attribute to keep up to this point, but perhaps for future productions when we know that the schemas will not change, the tables can be compressed with

this utility to bring added space and scalability to the experiment's data storage.

3.5.2 Key Practices for Success

If there is a chink in the armor of working with external databases, it is that one works over a remote connection with a large amount of data, and this can be error-prone. If you run a query that's somewhat complex and on a sizeable data set, one runs the risk of the operation not being able to run within the memory (RAM) of the server, in which case temporary files must be created, and the run/query time increases. With high enough complexity or large enough tables or poor table indexes, this can result in *frustratingly long* query times, sometimes with a `Lost Connection` type of error. There are a few habits one can adopt to reduce the likelihood of frustrating outcomes. These are completely the author's own opinion, and should be treated as such.

Temporary Tables. The temp tables are a great tool to consider when performing any tasks with several intermediary steps. These can be created by only adding the modifier `TEMPORARY` before the word `TABLE` in a `CREATE` statement. Temporary tables are persistent only within the existing session and cannot be seen or touched by any other connection. Once the user disconnects, the table is removed automatically. This type of mechanism is not revolutionary, but regular use of temporary tables gives a user the freedom to manipulate the data as if on a real table, and then once the work is done, there is no need to be concerned with the user space cluttering up, as the temporary tables will simply be gone once the session ends. An example of how the author uses temporary tables is to do the following:

```
-- Create a temp table that has the same structure as the target table
CREATE TEMPORARY TABLE temp1 LIKE merged_roadset62_R005_V001.kDimuon;
-- Fill it with a rough cut of the data you want
SELECT * FROM merged_roadset62_R005_V001.kDimuon
WHERE mass > 4.2 AND mass < 10 AND
chisq_dimuon < 15.0;
-- Perform more cuts on data in the temporary table, add new fields, etc.
...
-- Once all cuts and changes are made, create permanent analysis table and fill it with the data
CREATE TABLE my_table LIKE temp1
SELECT * FROM temp1;
```

The benefit of performing these tasks via this method is that you will only have a final analysis table if

all steps proceeded without error or disconnection. This allows one to avoid the hazards of half-processed data.

Perform tasks in piecemeal. Here, I mean piecemeal in the manner of “bit by bit” or “piece by piece.” This is highly context-dependent, but breaking down a large task into smaller ones can be sometimes faster, but always more stable. Just because you *can* execute all of your cuts in one comprehensive query joining on several large tables – it doesn’t mean you *should*. A common strategy (exhibited in the above temp table example) is to begin with a highly selective cut and then proceed with that highly diminished subset of the larger table.

Another measure that makes an analysis code more immune to scaling problems is to have the code first run something like a “`SELECT COUNT(*)FROM tbl_name;`” for whichever core table one uses to get a grasp on how large the dataset is. With this count in hand, one could go through that dataset one arbitrarily-defined *chunk* at a time. For example, the Roadset 67 data is quite voluminous. When applying advanced cuts on a large set of N dimuons from the Roadset 67 dataset, one could perform the cuts on X number of dimuons at a time using the “`LIMIT OFFSET, X`” clause at the end of the initial `SELECT` query, where `OFFSET` is the number to offset your selection by as you iterate through the whole set. This is generalizable to all the other roadsets and those to come, with some empirical testing to determine a suitable X .

All this is not to say that it is prohibitive to perform all tasks/cuts at once, which is why this is prefaced with context-dependent. Simply keep in mind that if size might be a problem, then consider this approach.

Develop using small sample sets. This is a small suggestion that applies widely to development in general. Too many times it has been observed analyzers still developing analysis code attempting to run it on large datasets only to find it return a resultset that has a problem with it. Choose a smaller test dataset or use the aforementioned “`LIMIT X`” clause to make the development more conducive to fast turnaround times and debugging.

3.5.3 Future Technologies

As recent as 2011, the so-called *Big Data* frontier has experienced an influx of new technologies that have allowed it to really take off. The term of Big Data is rather vague, as it means something different to various disciplines. As far as high-energy physics goes, it is a methodology of processing data in which **all** interesting data is available for **fast** analysis. To some degree, SeaQuest has attempted to provide this by making the **merged** production data available with indexing applied in such a way as to provide fast querying of data.

However, it certainly cannot be said that *all* data has been available at once since the Hit data proved to be too voluminous to simultaneously store all of it. It also cannot be said that retrieval for more complex queries has been exceptionally *fast*.

One of the new technologies that have emerged is an efficient, distributed data storage paradigm of NoSQL databases. With NoSQL systems, the ability to *join* tables together and use higher-level SQL queries is forfeited, but what is gained is the ability to scale data storage across an arbitrarily large number of nodes. This is achieved by storing all data in objects called *documents*, usually in the form of key-value pairs, and focusing on the feature of “horizontal scaling,” which is the ability to add/remove nodes to the system (as opposed to “vertical scaling,” which is extending the capability of a single node). In addition to speed and capacity scaling, the structure lends itself well to parallel read-writes.

In a proposal by Igor Mandrichenko, a services manager of the FNAL computing division, an outline of a possible future design for Fermilab Computing is described in detail [68]. The traditional approach is described as having the steps of collecting data via a DAQ system, writing that data to tape storage, read tape data in order to do reconstruction and data reduction, write this analyzed data to disk and then to tape. This data is then analyzed event by event for good analyzable data. In this general approach, magnetic tape is the primary storage with a “write once, read many” philosophy, and there is a lot of travel between disk buffers and tape. A diagram of such a paradigm is shown in Figure 3.6.

In the proposal, Mandrichenko suggests using a NoSQL storage engine to store *all* of the raw data from experiments at Fermilab combined with an SQL database to keep track of all of the indexing – keep tabs of *what* is in the mass storage and where to find it. The indexing engine can then be used by individuals/groups to index “interesting” data for their own analysis and can be summoned (even its raw data) at any time. In this paradigm, raw data is written to tape storage, but the philosophy here is “write once, read hopefully never.” Current estimates quote 5 PB with 2-3x replication (effectively 2 PB) storage) and an indexing server of ~100 TB. This approach is laid out in visual detail in Figure 3.7. This proposal is simply an example of the type of technologies that are going to inevitably emerge as new facilities and experiments are put together.

Further into the future, though possibly sooner than one might imagine, cloud computing and storage has become more and more of a presence in both industry and academia. Not too long ago, it would have been almost unthinkable for a national laboratory or large industrial companies to store their data on server farms hosted by someone else. With the proliferation and maturity of Infrastructure-as-a-Service (IaaS) paradigm, services like Amazon Web Services (AWS), Microsoft Azure, and Google have truly emerged as a viable solution for institutions of all walks to trust them with their valuable data. Such providers are able

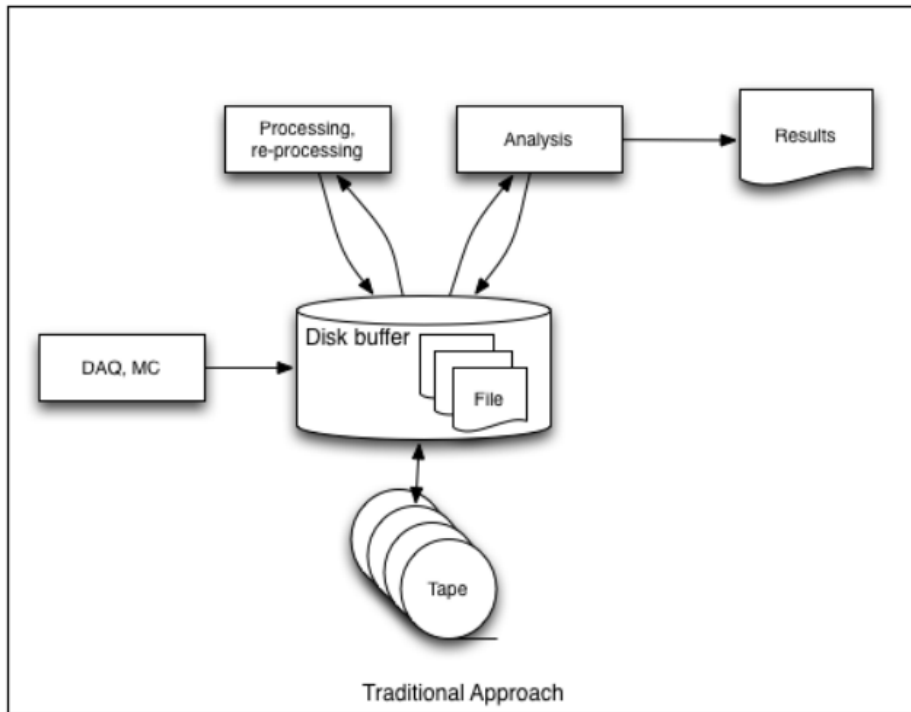


Figure 3.6: A traditional approach to HEP data handling at FNAL [68].

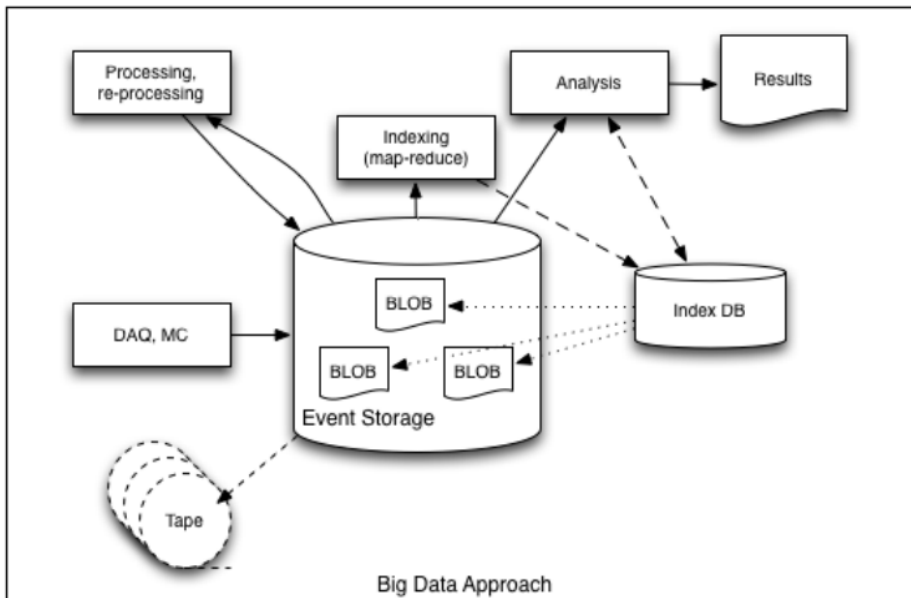


Figure 3.7: A proposal that suggests using a combination of horizontally-scalable NoSQL technologies to store all data [68]. Blobs are “binary large objects” that are application-dependent – they could be CODA or ROOT files, for instance.

to promptly provide scalable, “elastic” (expandable without downtime), redundant storage and processing using any storage engine, running on any operating system.

The matter simply comes down to the fact that assembling, administrating, and maintaining server hardware and systems has *significant* overhead. Additionally, as technology evolves and new capabilities become possible, smaller operations become less agile in being able to keep up with developments as they happen. As IaaS providers have gradually earned the trust of users in the way of data security and redundancy, they have suddenly, over the past two years, become a *leading* contender when new data infrastructure is considered.

When all costs are compared and needs are assessed, it would not surprise the author if mass data storage and handling are eventually outsourced to IaaS providers as new experiments are proposed in the next 5-10 years. The rationale behind this projection is that, since national labs and experiments are largely publicly funded, there will be pressure to spend the allotted funds in a responsible manner. As such services become cheaper, more effective, and more trusted, the number of reasons to keep data warehouses local to the laboratories will very quickly decline, and the pressure to adopt will increase.

There, however, will always be exceptions. In the cases of experiments like ATLAS and CMS, the data rates are beyond what exists elsewhere in the world at 20+ TB/day [69]. Clearly, for such situations, dedicated, specialized systems will likely need to be designed and maintained locally.

Chapter 4

Photomultiplier Tube Base Upgrade

During Run I of SeaQuest, observations of hodoscope wire maps (as in Fig. 4.1) suggested an apparent drop in expected performance in the y -measuring hodoscopes. While this performance was most obviously seen in the y -measuring hodoscope planes, the x -measuring planes were likely also affected. This effect was assumed to be due to high-intensity RF buckets that caused very high multiplicity in all of the detectors in the spectrometer for that event. The result of these intense events seemed to push the PMTs and/or their PMT base electronics past their operational capacity.

The understood cause of this “*sag*” in performance, as it came to be called, was due to a destabilization in the voltage divider in the PMT base. This critical component holds each dynode stage at a specific voltage, and when this destabilizes and is unable to maintain an appropriate voltage difference between dynode stages, inefficient performance of the PMT results.

During the Fall of 2012, prototyping and testing were performed with the goal in mind being to assemble a new base for the Philips XP-2008 PMTs [70] and compare its performance to the original PMT base and to some modern, high-performance Hamamatsu PMTs. Once a base design tested well, the new bases would be manufactured and installed in the existing frames of the original PMT bases.

4.1 PMT Basic Construction and Operation

Figure 4.2 shows a schematic design of a typical photomultiplier tube and base setup. It consists of a photocathode that is followed by an electron multiplier section (or dynode string) then an anode from which a final signal is delivered. During operation, a high voltage is applied to the photocathode, dynodes, and anode in such a way that there’s a potential “ladder” going from stage to stage. When an incident photon from the hodoscope scintillator paddle hits the photocathode, an electron is emitted via the photoelectric effect [71]. The voltage difference between the cathode and dynode stages draws the emitted electron to the dynodes, and each time an electron hits a dynode, some of that electron’s energy is transferred to other electrons in the dynode. These electrons then are emitted and become accelerated towards the next dynode

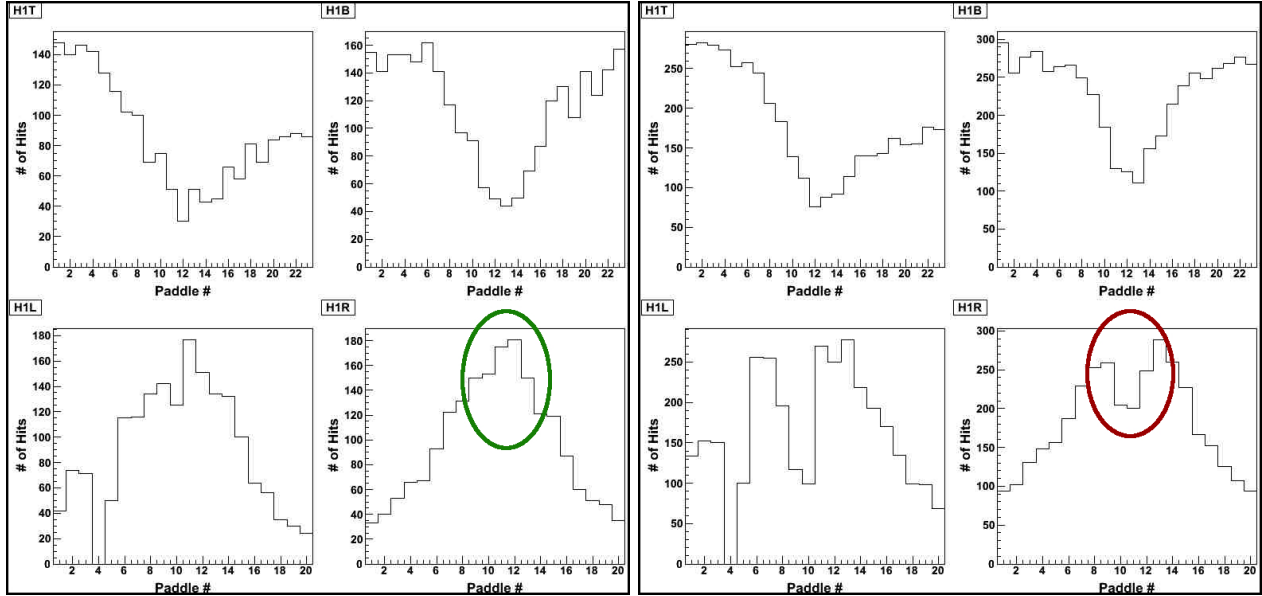


Figure 4.1: (Left) Histogram of hodoscope ‘hits’ in a typical event; (Right) Histogram of high-intensity event, with marked sagging most noticeably in the middle of the y-measuring hodoscopes




Figure 4.2: A diagram of typical PMT operation. The circuit controlling the voltage-dropping resistors is the part that was upgraded in this chapter.

stage. This process is called secondary emission, and by the time the process is repeated, there is a cascade or avalanche of electrons that land on the anode, resulting in a signal that can be amplified and analyzed.

It is the case that the voltage divider ultimately supplies the electrons that are emitted in this signal cascade. If too many photons and resultant electron cascades occur, the dynode stages' voltage divider will destabilize as they attempt to resupply the dynode stages with electrons. The problem that was experienced at SeaQuest was that these high-intensity events were flooding the PMTs with photons, causing this “saturation” which caused this destabilization and the inefficient performance that was observed. The goal specifically was to test out modern base designs that provided for added stability to the performance of the voltage divider, even under high rates.

In general, each base divides around a -1500 V potential total over the photocathode (K), ten dynode stages (D1-D10), and the anode (A). There are two currents that are referred to here:

- Signal Current: This is the signal that passes over the anode, which is the end-result of the cascading secondary emission electrons from each dynode stage.
- Bleeder Current: This is the current through the voltage divider. It is termed the “bleeder” current since the compounding electrons in the signal current must be “bled” from the current through the voltage divider.

Throughout these voltage base designs, capacitors are commonly implemented in the latter dynode stages where the most electrons are emitted. These capacitors, when charged, are able to replenish the lost charge on its corresponding dynode stage in the event that an intense light pulse induces a large signal current. As the capacitor is able to hold its own charge, this resupply can occur without requiring the charges to be drawn from the bleeder current, thereby keeping the voltage across the dynode stages more stable.

4.2 PMT Base Design Iterations

There were several iterations of base design to determine which was best to approach for full base production and installation at SeaQuest. The core addition was the inclusion of transistors between dynode stages, according to the improvements suggested by C.R. Kerns in his paper regarding high-rate PMT bases [72]. Common solutions to destabilization in PMT bases have been to have (1) very large capacitor banks with charges $> 10^3$ times greater than the time-averaged dynode current and/or (2) miniature on-board, separately-powered Cockcroft-Walton power supplies for the final dynode stages. A Kerns-style transistorized base allows for a light-weight, small size, and simple base that does not require extra power supplies or voluminous energy storage capacitors.

In general, there are three important features that were tuned in this set of prototypes that affected the performance of the phototubes:

- Lower resistance
- Transistors (with protective diodes)
- Higher capacitance between dynode stages
- Distribution of voltage division

The lower overall resistance of the voltage divider increases the bleeder current. This means that the base will be more capable of handling high-intensity, as it will be better able to replenish the charges on each dynode stage in the case of a large signal. Typically, the larger the bleeder current, the larger the signal current can be without destabilizing the voltage divider. Higher rates usually put higher demand on the signal current, so by reducing the overall resistance, one can easily increase the rate capability. The shortfall here is that with voltage constant and resistance decreased, according to Ohm's Law ($V = IR$), the current will increase. As a result, the power dissipated by the circuit ($P = I^2R$) will go as I^2 , and the PMT base may heat up to critical temperatures faster than it can dissipate the heat as there is no significant ventilation in the PMT base enclosure. The Philips XP-2008 manual quotes that for continuous usage and storage the ambient temperature should not exceed $50^\circ C$ ($122^\circ F$). In addition to heat concerns, there's typically a power rating for the class of small on-board resistors that were planned to be used. Approaching or exceeding that power limit would run the risk of burning out a resistor and rendering the base inoperable.




Figure 4.3: MOSFET showing gate (G), body (B), source (S) and drain (D) terminals. The gate is separated from the body by an insulating layer (white) [73]

The use of a specific type or micro-transistor, a metaloxide semiconductor field-effect transistor (MOSFET), is introduced here to maintain the proper voltage division. In general, MOSFET transistors have an *source*, a *drain*, and a *gate*, where current flows freely through from the source to the drain, gate permitting (Fig. 4.3). If at any point a certain voltage across the gate of the transistor is not supplied (here, the voltage across dynode stages), then the source-to-drain current through the transistor is stopped until the proper gate voltage is restored. This helps greatly to “intelligently” regulate the voltage across the dynodes. Wherever transistors are used,

diodes are also implemented to prevent the unlikely case of a current moving across the transistors' gate

Recommended voltage divider

Type A for maximum gain

K	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	A	
2	1	1	1	1	1	1	1	1	1	1	1	(total :12)

Type B for better timing / linearity compromise

K	D1	D2	D3	D4	D5	D6	D7	D8	D9	D10	A	
2	1	1	1	1.25	1.25	1.5	2.25	2.25	2.5	2.25		(total :18.25)

K: photocathode Dn: dynode A: anode

Figure 4.4: Suggested voltage division schemes for gain vs. timing/linearity compromise [70].

in the wrong direction. This protects the transistors from being damaged particularly when powering the circuit on and off.

Having capacitors along the higher dynode stages allows for a quick resupply of charges to the dynodes, thus maintaining proper voltage division. This, however, is only a stop-gap measure and is only effective in cases of high instantaneous currents. It is the case that, should there be a constant too-high rate of operation, the capacitors will not be able to recharge themselves. Typical recharge time for the capacitors discussed in this section can range from 0.1-1 ms.

Finally, the specific division of voltage across each stage, from D1 to A, has an influence on the behavior of the PMT operation. As we see from the operations manual of the phototube in Fig. 4.4, in the case of a progressively increasing voltage division, there is a good compromise between timing and linearity. With respect to phototube operation, “linearity” is the quality that the amount of charge deposited on the anode is linearly proportional to the energy of the incident photon. “Timing”, on the other hand, is the quality that the time it takes for a high-energy photon signal and a low-energy photon signal to progress through the stages should be the same. For the purpose of optimization at SeaQuest, we wish to optimize the amount of signal (i.e. amplification or *gain*) that the phototube can accommodate. For this, the recommended voltage partitioning is flat from D1-A [70].

It should be noted that the prototype iterations of PMT base design changes were not intended to cover the entire phase space of these tunable parameters. The purpose of these tests was to get a sense of how changing one or more of these parameters could affect the PMT high-rate capabilities. The decision for a final base design had to be decided on with relative speed to get them manufactured, built, and installed before Run II of the experiment. As such, true optimization of all parameters could not be achieved within the scope of these tests.

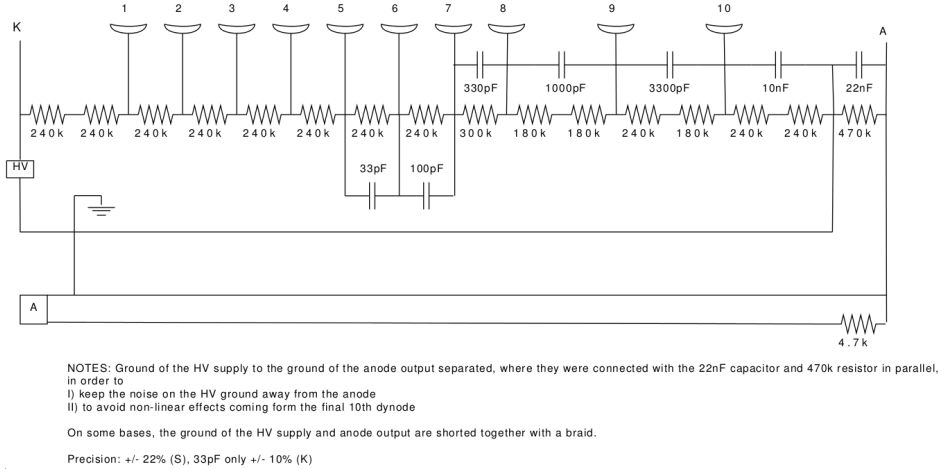


Figure 4.5: The original PMT base inherited from the ARGUS and HERMES experiments.

4.2.1 Original Base

The base that came attached to the PMTs were manufactured specifically for use by the ARGUS experiment, which was a relatively (by SeaQuest standards) lower-rate collider experiment that used e^+e^- annihilation at the *DORIS II* ring at DESY. After their tenure at ARGUS, they were handed down to the HERMES experiment located at the *HERA* polarized electron accelerator at DESY.

Though no actual circuit diagram was documented for the original PMT base, it was dissected and each component was measured. The results can be found in Fig. 4.5, and its voltage division seen in Fig. 4.6. It features a simple string of resistors with capacitors of increasing capacitance along the last six stages. The voltage division can be recognized to be similar to the timing-linearity compromise scheme described in Fig. 4.4. With a total resistance of approximately $3.95 M\Omega$ and the operational voltage of $-1500 V$, the expected standing (bleeder) current even when sitting in the dark is expected to be $0.38 mA$. With this in mind, we would expect the voltage divider to destabilize when the signal current approaches this value.

4.2.2 Prototype Base v1

Once the task was set to update the PMT base design, the Fermilab Particle Physics Division was consulted on the matter. In 2010 a base design with similar goals for the exact same PMT model was designed by Sten Hansen [74]. The circuit diagram for the new base can be found in Fig. 4.7.

Here, the resistance was significantly reduced by a factor of about 2.9, allowing for much more bleeder current, without exceeding or closely approaching the on-board resistors' power rating. Also, the voltage division was designed to be relatively "flat" (Fig. 4.8) across stages from D1 to A, which is stated to be recommended for optimal gain. With a total resistance of approximately $1365 k\Omega$ and the operational voltage




Figure 4.6: The voltage division between subsequent stages for the original PMT base design when supplied with -1500 V.

of -1500 V, the expected standing (bleeder) current even when sitting in the dark is expected to be 1.1 mA. Already here, we see that this design parameter alone suggests its ability to withstand $\sim 3x$ more signal current as compared to the original base.

The introduction of MOSFET transistors is seen between each stage from D7 to D10. Sending the current in parallel over a $1 M\Omega$ resistors allows the gate of the transistor to measure the voltage without drawing much current. The Zener dynodes are there in place before each transistor gate to ensure that current only goes one way across the sensitive gate channels.

It is also notable that there are banks of capacitors in parallel across the higher dynode stages. Since there is higher current through this circuit than the original base under the same voltage, there will be a greater demand placed on the capacitors to resupply the dynode stages with spent charge. The two 10 nF capacitors in parallel across each stage (which amount to 20 nF total) is significantly greater than the capacitance across the stages of the original base.

4.2.3 Prototype Base v2

The first modification made to the prototype board was to keep everything identical except for the total resistance of the circuit. This was accomplished by halving the resistance of each of the first six stages (R6-R13 on Fig. 4.7) from to increase the bleeder current. The resulting current of the base at -1500 V is at around 2.2 mA. The voltage division retained the same values as described in Fig. 4.8.

4.2.4 Prototype Base v3

In the case that destabilization was occurring prior to the dynode stages with the added capacitors and transistors, the third prototype was decided to take the Prototype v1 design and add more “transistorized”

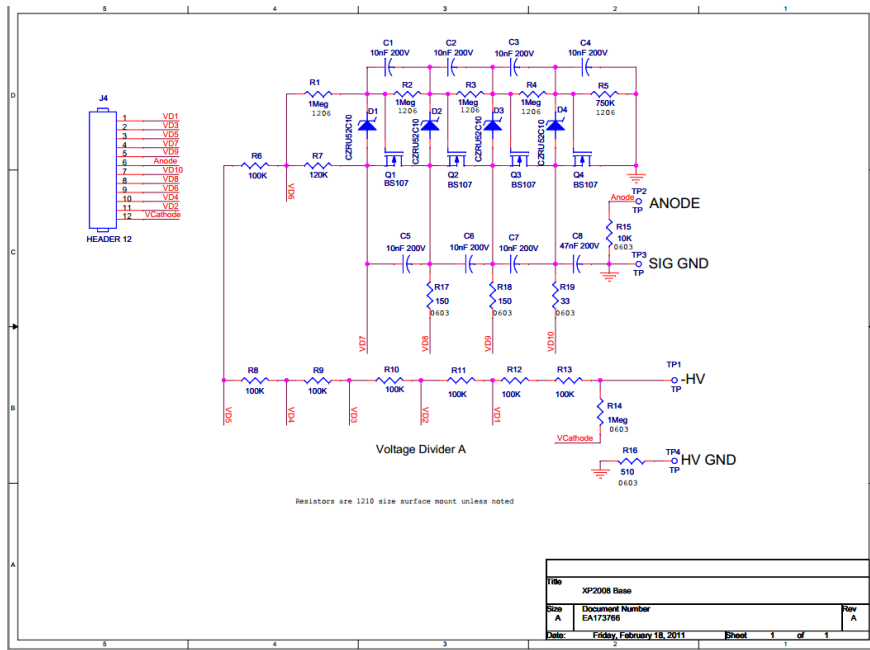


Figure 4.7: The Prototype v1 board circuit diagram received from Fermilab Particle Physics Division [74]. The parts are denoted as R: resistor, C: capacitor, Q: MOSFET transistor, D: Zener diode.




Figure 4.8: The voltage division between subsequent stages for the Prototypes v1, v2, and v3 PMT base designs when supplied with -1500 V.

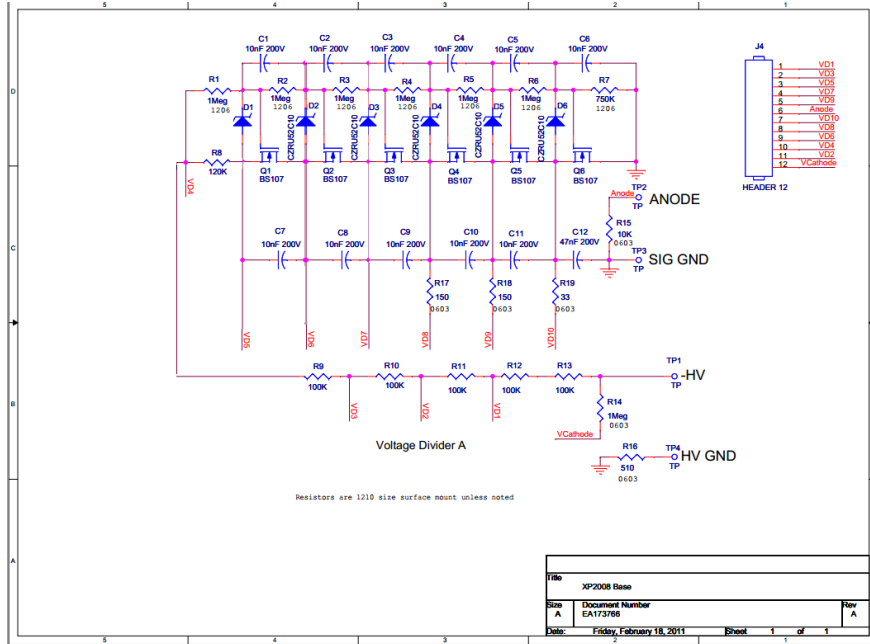


Figure 4.9: The Prototype v3 board: 3 more transistorized stages than the Prototype v1 design.

stages earlier on. This entailed extending the parallel configuration of capacitors, transistor, diode, and $1\text{ M}\Omega$ to the D5-D6 and D6-D7 stages. This prototype configuration can be seen in Fig. 4.9). The voltage division was not significantly altered by this change and remained relatively the same as what is described in Fig. 4.8.

4.2.5 Prototype Base v4

The final modification arose from a suggestion from a Fermilab collaborator. It was suggested that it may significantly extend the dynamic range of the tube/base by increasing the voltage drop in, specifically, the last stage relative to the other stages. This reduced to simply replacing R5 (of Fig. 4.7), a $1\text{ M}\Omega$ resistor, with a $1.5\text{ M}\Omega$. All of the rest would remain unchanged from Prototype v1. The premise for this modification was in the case that the final batch of electrons needed help being “swept” to the anode with a higher voltage difference. The change applied resulted in the voltage distribution according to Fig. 4.10.

4.3 PMT Base Comparisons

There is a specific difficulty with the objective to increase the rate capability of our PMT’s. This difficulty is that there was not a known instantaneous intensity or target rate capability to attain. The intensity that caused the original PMT performance to sag is unknown, and if it was known, it would be difficult to




Figure 4.10: The (negative) voltage between subsequent stages for the Prototype v4 PMT base.

match the intensity with an experimental setup. As a result, the objective of these tests was to *compare* the performance of the same PMT under controlled conditions using the original and various prototype bases.

Due to the effects of using different PMTs and due to temperature and humidity fluctuations, the PMT behavior can be somewhat variable from test to test. For this reason, one can only reasonably compare results within each base comparison test, and not across different comparison tests. Each was performed on different days, and sometimes with different PMTs.

4.3.1 Testing Apparatus, Measurements, and Procedure

In this experiment, a PMT attached to a PMT base was placed into a light-tight box facing a fast-pulsing 470nm wavelength LED, which provides the driving photonic signal.

PMT Test Setup

The LED light source assembly consisted of a machined aluminum block that housed the LED, which was a fast-pulsing, narrow beam, water-clear blue LED, peaking in the $\lambda = 450$ nm region. This block had a sliding aluminum insert which had a 1/4" depression used for inserting a 1" diameter NDF of arbitrary optical density [75]. The LED, according to specifications, has a N ns rise time and a N ns fall time, which is important when considering pulsing frequencies on the order of 30 MHz.

The LED itself was driven by an Agilent 33520 Function / Arbitrary Waveform Generator [76] capable of generating signals up to 30MHz. The LED's unattenuated intensity was far too much light to perform any useful test with such photosensitive hardware. Its intensity was attenuated by use of a neutral density filter (NDF), with a rating $D = 3.0$, where the NDF allows 1 in 10^D photons through (1 in 1000 for $D = 3.0$). When testing began, a $D = 4.0$ NDF was used, but the amount of light reaching the PMT was so small that no decrease in PMT performance for any of the bases was observed all the way up to the highest LED




Figure 4.11: Inside of the light box, with the prototype board (left) wired up to a Philips XP-2008 PMT (middle), facing a fast-pulsing LED source (right).

frequencies. As such, the $D = 3.0$ NDF was chosen for this study.

These were all kept within the light-tight box¹ loaned to the SeaQuest collaboration from the Daya Bay group. This light box was equipped with a patch panel that had both BNC and HVBNC connectors by which to power the PMT base and read out its signal. The light box and all of the installed components can be seen in Fig. 4.11.

A simple data acquisition was assembled from an amplifier, discriminator, and scaler in order to observe that the PMT was functioning properly and firing off at the rate that the LED was set to pulse at.

Measured and Calculated Quantities

The PMT base was powered by a high voltage supply, and an ammeter was connected between the two in order to measure the amount of current drawn, or bleeder current, from the HV power supply. The PMT signal was processed by an oscilloscope, averaging the pulse over 300 pulses. The primary measurement was measuring the area of the averaged pulse ($V \cdot s$).

We calculated the signal current from the anode as:

$$Q_{pulse} = \frac{\int V dt}{R} \quad (4.1)$$

$$I_{signal} = f Q_{pulse} \quad (4.2)$$

¹These light-tight boxes are paradoxically referred to as both “light boxes” and “dark boxes”, as they’re used for testing *light*-sensitive equipment and they’re made to be very *dark* inside.

where f here is the driving frequency of the pulsing LED, R is the termination resistance of the signal ($50\ \Omega$), and $\int V dt$ is the integrated area of the averaged PMT pulse. The amplitude of the pulses was also measured, as it was important in determining if it were feasible to remove the typically noisy hodoscope amplifiers from the SeaQuest stations 1 and 2 DAQ setup.

The measured and calculated quantities of interest are the bleeder current, the averaged signal amplitude, averaged signal area, and signal current over the anode.

Due to the limitations of the instruments used in these tests, there is some systematic uncertainty to consider. The ammeter used to measure the bleeder current between the HV supply and the PMT base used an analog gauge that had closely spaced ticks at $0.2\ \text{mA}$ intervals. As such, the systematic uncertainty for the bleeder current is considered to be $\pm 0.05\ \text{mA}$. Also, due to pulse-to-pulse variations which cause a mild amount of smearing in the average, the measurement of the amplitudes and signal areas have an uncertainty of $\pm 5\ \text{mV}$ and $\pm 0.1\ \text{nVs}$, respectively. The uncertainty in the signal area translates to an uncertainty in the calculated signal current of around $\pm 0.04\ \text{mA}$. These fluctuations are considered negligible with respect to the quantities measured.

Test Procedure

Most of the prototyping was performed over the span of two weeks. As new suggestions were made at base improvements and the iterations v1-v4 progressed, more comparison tests were performed. The procedure for the each test was the same and is as follows:

1. Connect the base to the PMT and to the HVBNC and BNC connectors.
2. Close the lid to the light box.
3. Gradually power the PMT ($-250\ \text{V}$ every $5\ \text{s}$) until it is as $-1500\ \text{V}$.
4. Watch the scaler counter on the DAQ to see if it is rapidly counting. No (or very few) counts on the scaler. means the box is light-tight.
5. Let the PMT and base warm up to near-operational temperatures with its standing current for a minute or two.
6. Power on the function generator and LED, setting the function generator to create a $16\ \text{ns}$ square wave pulse to the LED circuit (the minimum achievable pulse width [76]) at a frequency of $10\ \text{Hz}$.
7. For each of several frequencies from $10\ \text{Hz}$ to $30\ \text{MHz}$, perform the following:
 - (a) Observe the scaler to ensure that the count is increasing at the same rate as is set by the function generator. This confirms that both the PMT, base, function generator, and LED are generally working as planned.

- (b) Trigger on the signal from the PMT output.
 - (c) Average the signal over 300 triggered pulses and then freeze the frame.
 - (d) Record the current draw from the HV supply (bleeder current).
 - (e) Measure and record the pulse amplitude with the oscilloscope cursor.
 - (f) Measure the pulse area using the signal integration feature, setting vertical cursors at the zero-intercept on each side of the pulse.
 - (g) Return to triggering mode and increase the function generator frequency as warranted.
8. Power off the function generator and LED circuit.
 9. Step the PMT voltage down slowly to 0 V.
 10. Open the light box and replace the base for subsequent tests as necessary, repeating the steps above.

4.3.2 Original vs. Prototype v1

The first test was to compare the original base with the first iteration of the prototype base. It was generally expected that this new design with more modern technology would exceed the performance of the original base, but the question was twofold: to what extent does it perform better and what is the baseline by which to compare future improvements?




Figure 4.12: Measurements of the (a) signal amplitudes and (b) bleeder and signal currents in the original and prototype v1 PMT bases. The HV was inadvertently set to -1600V V for this test instead of the intended -1500 V.

The first important observation is the baseline amplitudes of the signals from the two bases, as seen in Fig. 4.12(a). While the absolute voltage itself is not as useful since the LED doesn't necessarily reflect actual operating conditions, we do see that the amplitude from the new base design is increased by a factor of more than two. This in itself is significant in when considering the limits of the discriminators that were used in the hodoscope arrays for Stations 1 and 2. With even a modest boost to gain and signal amplitude,

the PMT signals should be large enough that an amplifier is not required in order to get a discriminator to trigger on a PMT pulse. Seeing as the amplifiers used had often caused troubles with the amount of electronic noise they tended to imbue, this was an encouraging observation.

Next, we look at the behavior of the amplitudes as the frequency progresses. It is safe to assume that the unchanging behavior in the lower frequencies indicate that the voltage divider is stable and that there is no voltage breakdown yet. At higher frequencies, there is a qualitative peak in amplitude for both bases, followed by a steep decline. This rise and then sharp fall would seem to indicate the end of stable PMT base operation, with the fall indicating the onset of degradation of performance.

To better identify the conditions that cause this onset, we look to currents applicable to both bases (Fig. 4.12(b)). In this particular test, we see the original base peak in amplitude at ~ 4 MHz and prototype base v1 peak at ~ 10 MHz. Looking to the currents, we see that these particular points mark a certain point for both where the signal current reaches almost exactly 50% of the bleeder current. The baseline improvement with this base over the original one was estimated to extend the operational range by a factor of three – the point where the amplitude drops below the low-frequency baseline.

Around the time of these tests, Sten Hansen conducted a SPICE (Simulation Program with Integrated Circuit Emphasis) simulation of the prototype circuit and was able to reproduce what was observed. In particular, it was observed that the gain began to change dramatically when the anode (signal) current reaches a significant fraction of the total current drawn from the HV supply; a current that is much less than 100%. It was concluded that 50% was consistent with his findings.

4.3.3 Original vs. Prototype v1 vs. Prototype v2

Several suggestions were made in altering the prototype base in order to maximize this improvement, including lowering the overall resistance of the circuit. If the signal current grows at the same pace, then by increasing the bleeder current, there is a chance that a higher frequency signal will be required to bring the signal current up to the point where the divider breaks down, which we estimated above to be $\sim 50\%$ of the increased bleeder current.

Looking to Figure 4.13(a), we see that this basic premise is technically sound. The prototype v2 base exceeds the dynamic range of both the original base and the prototype v1 base by factors of approximately a factor of four and two, respectively. This is a significant improvement, and this, at first, seems to be an avenue to pursue. However, an issue arose after increasing the LED frequency from 13 MHz to 14 MHz when a resistor burned out. The test continued with the rest of the bases, but the results paint a clear picture of what happened.

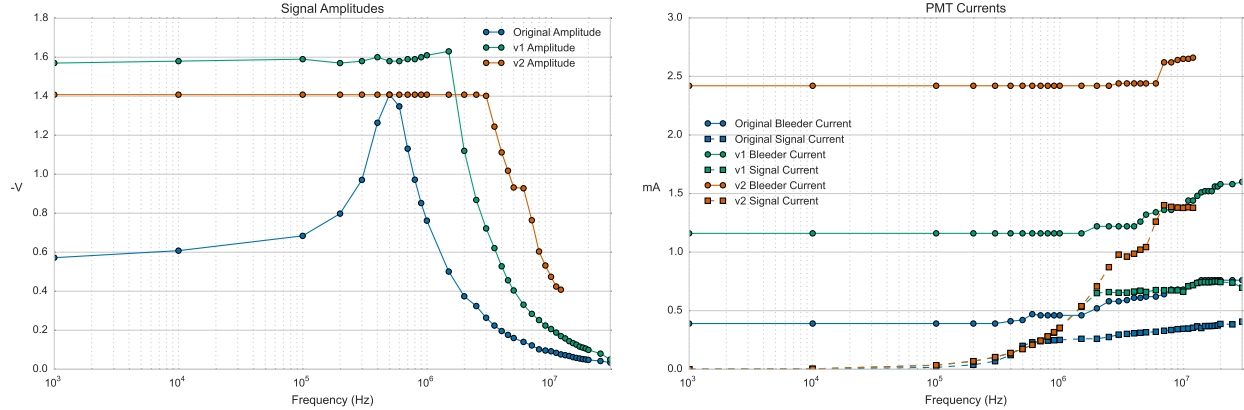


Figure 4.13: Measurements of the (a) signal amplitudes and (b) bleeder and signal currents in the original, prototype v1, and prototype v2 PMT bases.

By halving the resistance, the bleeder current doubled (Figure 4.13(b)). This caused the power being dissipated by the resistors to exceed their power rating. The prototype bases were built with 0.25 W resistors, and as the bleeder current rose well above 2.5 mA, one of the resistors failed. This did not bode well for how a smaller on-board resistor might fare with the same current.

To further exacerbate matters, it was found that the LeCroy 1440 HV supplies that were to be used could only supply a maximum of 2.5 mA per channel [77]. Seeing as these power supplies had proven themselves sensitive and prone to communication and performance failures, it was deemed inadvisable to demand anywhere near the maximum current from these HV supplies.

For these reasons, it was decided that the new PMT base design of prototype v1 not be modified in this manner which might suffer from overheating, component failure, and a too-high current demand.

4.3.4 Prototype v1 vs. Prototype v3

Perhaps it was the case that the “transistorized” steps were doing their job well and the destabilization was occurring at prior stages. Since the only downside of adding additional steps was added cost, a base with more transistorized stages was created as depicted in Fig. 4.9.

As is quite evident in Figure 4.14, there was no significant difference in the performance of the two prototype bases. As such, this modification was not adopted into the final design.

4.3.5 Prototype v1 vs. Prototype v4

The final test modification of increasing the resistance and thereby the voltage drop across the final D10-A stage was tested against the prototype v1 base.




Figure 4.14: Measurements of the (a) signal amplitudes and (b) bleeder and signal currents in the prototype v1 and prototype v3 PMT bases.




Figure 4.15: Measurements of the (a) signal amplitudes and (b) bleeder and signal currents in the prototype v1 and prototype v4 PMT bases.

This modification turns the voltage divider slightly away from the “optimum gain” configuration and more towards the “optimal timing / linearity compromise” configuration. This is seen to have a direct effect on how a loss in gain can be observed between the v1 and v4 base in Fig. 4.15(a). While the rate capability does seem to slightly increase from 4 MHz to 5 MHz, the value of having a higher gain was deemed more desirable. Higher gain, again, means that the then-used noisy amplifiers would no longer be needed, and the marginal loss in rate capability was considered a reasonable tradeoff.

4.3.6 Base Comparison Conclusions

After these tests were conducted and the information regarding each of the bases’ performances was evaluated, it was decided that the prototype v1 base design was the optimal base configuration to move forward with. This base design yielded a clear improvement in both signal amplitude ($\sim 2x$) and rate capability

($\sim 3x$) at the expense of more complex circuitry and an increased supply current draw. It then became the task to fit this circuit onto small printed circuit boards (PCBs), and then manufacture, install, and test them.

4.4 Base Manufacturing and Installation

It was decided that the original base structure was well-packaged and well-shielded with its μ -metal canister. As such, a two-component board was designed by Sten Hansen that consisted of a 2.5" diameter board and a small rectangular "daughter board". By June of 2013, the mother and daughter boards (attached) had been manufactured and delivered to UIUC for assembly into the rest of the base.

Since the sockets that coupled the board to the PMT pins was not commercially available, they had to be salvaged by deconstructing the original bases. This consisted of cutting the components between the two boards of the original base and then using a solder vacuum to separate the board from the socket pins. The socket and the other structural elements of the PMT base were separated and prepared for construction with the new base PCB (Fig. 4.16).




Figure 4.16: (Left) The components of the disassembled original base and the manufactured new base PCB with attached daughter board. (Right) The original and new bases fully assembled, side-by-side.

The deconstruction and re-soldering of the bases took place over the span of about three weeks. Each base was tested to provide the same voltage division across stages and was tested in operation inside the light box apparatus by observing its signal characteristics and the rate of the scaler counting (which should match up with the LED frequency).

After all bases were constructed and tested, the final task was to find the optimal operating voltage for a PMT assembly that would be suitable for detecting minimum ionizing particles passing through a

scintillator paddle similar to the ones used at Stations 1 & 2 at SeaQuest. We also want the minimum possible voltage, as higher voltages correspond to increased noise. This was performed by assembling three stacked scintillator paddles, each with a PMT attached to their light guide. Their outputs were passed through a discriminator and then to a coincidence unit, which in turn had its output sent to a scaler counter. The top and bottom PMTs were powered to -1500 V, and their coincidences were counted on one channel of the scaler counter. Simultaneously, the middle paddle, set to a variable voltage setting, had its coincidences with the top and bottom paddles counted on the other channel of the scaler counter. The ratio of these two counts ($\epsilon = (\text{triple coincidence})/(\text{top} - \text{bottom coincidence})$) described the efficiency of the middle hodoscope. The voltage of the middle PMT is varied until a plateau in the efficiency is identified and recorded to be the ‘nominal voltage’ of this reference PMT.

All of the constructed PMT bases and the “reference” PMT and base were then sent up from UIUC to FNAL to be installed into Stations 1 & 2 and used for “gain matching”. The gain matching is done by exposing a scintillator paddle to a γ -ray source while connected to a PMT and adjusting its voltage until it produces the same signal amplification, or *gain*, as the rest of the PMTs. The original PMT that the rest are matched to is the reference PMT from the previous efficiency study.

Matching the gain is not a simple task, and the typical approach is to take advantage of a physical phenomenon called the Compton Effect occurs in scintillator material when a γ -ray scatters off of it. Often times, the γ -ray of energy E will not deposit the whole of its energy, but only a portion of it. The spectrum of the photons that scatter off will have energy

$$E' = \frac{E}{1 + \frac{(1 - \cos \theta)E}{m_e c^2}} \quad (4.3)$$

where θ is the angle of deflection and m_e is the electron mass. This, in turn, deposits transferred energy $E_T = E - E'$ into the scintillator which has a maximum value of

$$(E_T)_{max} = E \left(1 - \frac{1}{1 - \frac{2E}{m_e c^2}} \right) \quad (4.4)$$

This energy is known as the Compton Edge, as it is characterized by a very sharp falloff in the energy spectrum of energy deposited by a source of a specific energy. The particular energy of this edge is not important for the purpose of gain matching. Instead, it’s used as a qualitatively sharp feature that can be clearly identified when examining the energy spectrum of photons in scintillator material, as seen in Figure 4.17. As such, it can be used as a landmark to use for adjusting the voltages of each PMT in a way such that the gain of each is the same.




Figure 4.17: The γ -decay energy spectrum of Cs-137 as seen through a scintillator-PMT-qVt setup [78]. Here, the Compton edge can be seen at channel 110.

Once installed onto a scintillator paddle, the reference PMT was powered to nominal voltage and a radioactive Cs-137 source was pointed at the paddle. Since the PMT signal has some *linearity* to its signal, the amount of energy of incoming photons will translate approximately linearly to the amount of signal out of the PMT. The signal from the PMT (with the Cesium source pointing at the scintillator paddle) is passed through a qVt (charge-voltage-time) module in “Q-mode” to histogram the total integrated charges of the PMT signals. The qVt’s channel location of the Compton edge for this reference PMT at this nominal voltage is recorded. Then, every other PMT is powered and measured through the qVt with the Cesium source, and its voltage is adjusted until the Compton edge falls on the same channel. This procedure is generally known as “gain matching”.

4.5 New Base Performance

All of this work is all for naught if the new bases do not hold up to even the highest of beam intensities at SeaQuest. After all of the aforementioned testing, installation, and gain matching, data taking resumed. In Figure 4.18, the hit per event distributions are shown for x - and y -measuring hodoscopes as for different intensity ranges. It is difficult to compare these plots to those in Fig. 4.1, as there was no measure of per-event intensity at the time, and the beam back when that measurement was made was less intense (though




Figure 4.18: The number of hits per event for the Station 1 hodoscopes for several intensity ranges. Shown is (left) the top x -measuring array and (right) the left y -measuring array. I_p is in terms of number of incident protons for the event.

just as non-uniform).

Regardless, the new readouts can be interpreted positively. One can see in the y -measuring array that the shape does flatten out at the peak for higher intensities. It can be noted that this flattening is much less severe than the “sag” that was initially observed. More importantly, when observing the absolute shape of the distribution of the x -measuring array, the shape remains intact and uniformly increasing all the way up to the highest intensities. Seeing as the x -hodoscopes are the only arrays that are currently being used for triggering and analysis, this is the most important behavior to observe.

Along with this confirmation, the amplitude output by the PMT signals was above the minimum threshold of the discriminators being used. As a result, the noisy amplifier modules were able to be removed from the signal pipeline and thereby reduced the number of false hits due to electronic noise.

Chapter 5

Tracking and Data Selection

5.1 Reconstruction of SeaQuest Dimuons

Once the raw hit and timing information has been gathered and processed by the production level software, the reconstruction software is run on it in order to find tracks that pass through the spectrometer. The ultimate purpose of the detectors in the spectrometer is to identify the 4-momenta of particles that occur as a result of high energy collisions and particle decays. Individual particle tracks and their 3-momenta are identified by tracking their paths through the known magnetic fields. These tracks are then matched up to find possible dimuon vertices. Reconstructed tracks, dimuons, and their constituent hits along with their full defining set of kinematics are provided as output for analysts to use in their studies. This tracking software, known as “kTracker”, was developed and is maintained by Kun Liu of Los Alamos National Laboratory. In this section, an overview of the tracking algorithm will be given in how the kTracker steps through the mass quantities of raw data in order to provide the reconstructed set of data used for SeaQuest’s analyses. Details of data reduction, track finding, track fitting, and vertex fitting will be discussed broadly.

5.1.1 Data Reduction

As the number of hits in the tracking chambers increases, it becomes more and more difficult to successfully perform pattern recognition in identifying real tracks. As a result, a litany of methods are performed in order to reduce the hits registered down to hits that are likely to have come from physical tracks and not unwanted noise and radiation. The sum of these data reduction measures thereby serve to make events that were once un-track-able more likely to be track-able. For events that would already be able to be tracked, these cuts still serve to reduce the run time of the tracking algorithm and curtail the number of erroneous tracks that could be reconstructed purely due to combinatorics and coincidences. These measures can be grouped into two groups: hit removal and event multiplicity cuts.

Timing and Masking

The first of the hit removal measures is the **in-time cut**. When an event is triggered by the trigger supervisor, each TDC issues a “stop time” by which each hit’s timing is measured against to derive its TDC time. Numerous “trigger timing” studies of the detectors’ and the TDCs’ behaviors have identified a certain range of values known as an “in-time window.” Any hits that occur outside of this timing window can be eliminated from tracking, as there is no possibility that the hit could have come from the span of time when the detectors were triggered. It should be noted, as it will later be relevant, that the drift chambers have in time windows as large as ~ 600 ns, which is equivalent to ± 15 beam pulses (RF buckets). This means that this alone will not do a very good job of paring the hits down to the relevant hit data.

Due to the behavior of the electronics of the drift chambers and their communications with the ROCs, there is an *echoing* effect that causes a signal to bounce back and forth along the signal line. The real signals that result from this echoing are called “afterpulses,” and they should not be used for reconstruction. Measures have been taken to reduce this effect in the form of adding ferrite cores at the ends of the cable ribbons connecting the electronic readout boards to the TDC cards, but afterpulses are still not absolutely avoided. Building upon the “in-time” classification, an the **“afterpulse removal”** criteria is defined such that only the first in-time hit for each element of each detector is accepted. One minor flaw with this method of data reduction is that It is possible that the first in-time hit for an element is actually the afterpulse of a hit outside of the in-time window. In these cases, the actual desired in-time hit will be removed.

The hodoscopes and their relatively high-precision timing can be used to eliminate spurious hits with great efficacy by a method called **“hodoscope masking.”** The hodoscopes have two qualities: the speed at which they operate, yielding a precise in-time window, and the fact that they are required to have fired in order for the trigger to have fired. The conclusion that one can arrive at is that if a chamber wire fires off and there is no adjacent hodoscope paddle fired off within the same event, then these chamber hits can be removed. From the detector specifications and the geometry of the experiment, a lookup table can be made associating a hodoscope paddle with a swath of chamber plane elements (and vice versa). This many-to-many relationship is used to eliminate all drift chamber hits that don’t have at least one associated in-time hodoscope hit.

Taking this even further, there is also **“trigger road hodoscope masking”** that is applied as a superlative of standard hodoscope masking. In this data reduction method, the set of in-time hodoscope hits is matched against the set of trigger roads (H1-H2-H3-H4 paddle combinations) used to trigger on the data was triggered with. First, the finite set of trigger roads that could have fired the trigger are determined. Then, hodoscope masking is performed once more, but the masking considers *only* the hodoscopes that are




Figure 5.1: (Left) A flow chart mapping out the decisions in how to treat clusters. (Right) A histogram of time differences of hits within a cluster. The peak below 10 ns is from electronic noise, and is the deciding factor in electronic noise classification.

a part of those reconstructed trigger roads.

Cluster Removal

The final measure taken to reduce the number of background hits is the **cluster removal**. A cluster is defined as any grouping of two or more adjacent hits within a drift chamber plane. The focus here is to either determine which wire is the best to use of a cluster or determine if the whole cluster should be removed. To decide on this, the three known sources of clusters are evaluated: electronic noise, cell-edge hits, and delta rays. Specific criteria are imposed in order to determine the cluster type and the action to take, as described by the flowchart in Figure 5.1.

Attached to the drift chambers, there are electronics boards (ASDQ) that relay the signals from the sensory wires to the TDCs. Each of these boards are connected to eight or sixteen wires. Similar to the afterpulse source, and it's possible that several adjacent wires can fire off due to noise in a single board. In these cases, none of these hits are likely due to any actual physical hits and should be completely remove where found. These types of clusters of hits are classified by the criterion of having very similar timing ($\Delta t \leq 10$ ns), which is logically to be expected if noise from the electronics board was at fault for the signals. This cut is applied only to clusters of > 3 hits for stations 1, 2, and 4 detectors. At station 3, there is a marked proclivity for electronic noise, and so the threshold for the number of hits is lowered to > 2 hits while narrowing the timing difference to $\Delta t \leq 8$ ns.

The next type of clusters to classify is “cell-edge” clusters, which is defined as cases where there are two adjacent hits in a single plane where both of them have a drift distance that is greater than 90% than the half of the cell width. In these cases, it is very likely the case that an energetic muon passed right between

two sense wires, and the hit was registered by both wires. In these cases, it is advantageous to reduce the number of hits and thereby reduce the combinatoric complexity for the tracker to simply reduce the number of hits here from two down to one. The hit that is retained is the one with the smaller drift distance.

The final known type of clusters that can occur are from *delta-rays*, which are (relatively) low-energy knock-on electrons that scatter away along the $x - y$ -plane of the detector. While these electrons are lower energy, they can still ionize the gas and incur hits in the detector. These clusters differ qualitatively from the electronic noise clusters in that, due to the speed of the traversing delta rays and the drift speed of the chamber gas, their timing difference will be $>10\text{ns}$ and will be missed by the electronic noise cluster removal. If a such a cluster is found, it is desired to keep the original muon hit and remove the rest of the hits incurred by the delta ray. In order to do so, the two hits on the edges of the cluster are kept while the rest are discarded.

Event Removal

After imposing all of these data reduction methods, there is still a chance that the event's contents would make it prohibitive for tracking. The primary criteria for such classification is the *multiplicity cut*. If there is an extremely large number of hits within one of the detector groups, the tracking algorithm will end up spending a very long time working through all possible combinations of hits in order to find good tracks – too long to be considered reasonable. As such, certain multiplicity limits (shown in Table 5.1) are imposed on the data to abort tracking on certain events. If the occupancy of any of the detector groups exceeds the limit imposed on it, the event tracking will be aborted regardless of the occupancy of the rest of the detector groups.

Finally, there is an additional multiplicity cut on the number of reconstructed possible *trigger roads* for an event. If the in-time hodoscope hits can match five or more trigger roads of a single sign (positive or negative muon roads denoted with positive or negative roadID, respectively), then event reconstruction is aborted. In these cases, most, if not all, chamber hits in that detector half will be masked by the road masking, and it's likely the combinatorics will be too much for the tracking to reasonably handle.

Det. Group	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

Table 5.1: The multiplicity limits imposed on detector groups by the tracking.




Figure 5.2: (Left) The definition of the sagittas s_1 and s_2 , with the red asterisk denoting the actual hit positions in the stations. (Right) The ratio of the *true* ratios of s_1/s_2 across muons of all kinds of kinematics (within acceptance) shows a single mean value with a defined variance.

5.1.2 Track Finding

Once the set of hits for a single event have been narrowed down to only the most plausibly useful hits, the next step is to begin the track finding. This consists of testing out combinations of hits between stations in order to arrive at a set of likely track candidates. The procedure for doing so can be broken down into four phases: local triplet reconstruction, St. 2-3 tracklet reconstruction, St. 1 projection of tracklets, and global track reconstruction.

The local triplet reconstruction can be thought of as the process of turning the hits of several wires in a drift chamber into a single set of $x - y$ position points within the chamber. Even further than that, a “stub” of a track is formed since a rudimentary slope can be estimated from the information garnered from the hits. The process begins with the combination of hits from the primed-unprimed planes. Having a pair of adjacent primed and unprimed hits allows for the elimination of the wire chamber left-right ambiguity, and better constrains the estimated position and slope. Then, all hits in the X-view is paired with all overlapping U-view hits. Those overlapping points from the X-U combinations are matched with the possible V-view hits. Each of these X-U-V combinations is then fit to an $x - y - z$ -position and a slope. In forming these triplets, a single hit may be used in multiple triplets.

Once this has been done for all Station 2 and 3 hits, the set of their track stubs are matched up to form straight “tracklets” between the two stations. Tracklets are then kept according a certain set of criteria: if their constituent stubs sufficiently point towards each other, if the tracks vaguely point back towards the target, if the tracklet points at hodoscopes (H2, H3, and H4) that fired, and if they point to proportional tubes (3 out of 4 planes) that fired. The absorber wall between Stations 3 and 4 provide a good requirement with looking to the corresponding prop tube hits, as only muons should have made it past the iron wall and fired the prop tubes.

With a set of reasonable tracklets, they are each projected through KMAG and onto Station 1 to find

possible matching hits there. The method by which this is done makes use of an analysis of the *sagitta ratio* if known DY muon tracks. The magnetic field of KMAG applies a p_T -kick of 0.404 GeV on the muon, and the influence on its deflection is based purely on the p_z of the muon at hand (assuming charge is known). As such, a large set of DY muons within the spectrometer acceptance are generated via Monte Carlo in order to measure good track behavior and the relationship between hits in Stations 1, 2 and 3. In Figure 5.2, the track path of a single muon is described with the influence of the magnets being treated as a single-impulse p_T kick at its center. With a line drawn between the tracklet position at Station 3 and the target ($z = x = y = 0$), the quantities s_1 and s_2 are defined, which are the distances between this line and the actual hits at Station 1 and Station 2, respectively. These studies show that the ratios of these two values across various muon characteristics is distributed as a gaussian with a defined mean and variance ($\mu = 1.77, \sigma = 0.055$). Since this ratio and s_2 is known for a given tracklet, s_1 can be approximated and used to constrain where the hit would actually be on Station 1 within a range of ± 5 cm. With this window established, triplets are matched up at Station 1 with the tracklet in order to find candidate *global tracks*.

Finally, with this set of candidate global tracks is narrowed down based on certain characteristics. This is useful as the track fitting phase which comes next is relatively intensive, so if there is any way to eliminate unlikely global tracks it should be done. The first step towards this end is the removal of hits with large *residuals*, where hit residuals are defined as the difference between where the hit triplet is and where the track is projected to pass through at that z position. Tracks with residuals greater than three times the resolution of the chambers are removed one at a time, with the track being refitted after each removal. Once this is done and all hits are within this acceptable residual range, then quality cuts are applied to the tracks. These cuts are (1) the hodoscope paddles along the track path must have fired, (2) the track momentum must be $5 < p_{tot} < 100$ GeV, (3) the track must have at least four hits in each station, (4) there must be at least one hit in each view (prime-unprime plane pairs) at each station, and (5) the multiple scattering in the iron absorber wall must not result in a scattering angle of more than 30 mrad. With these quality cuts in place, the remaining tracks are passed on to the Kalman filter track fitting algorithm.

5.1.3 Kalman Filter Track Fitting

For a given track, each positional measurement throughout the spectrometer has a well-defined uncertainty to them. When these spatial information is used to reconstruct the kinematics of a muon, the sum of these *fuzzy* measurements can be combined to render a better overall fit than any single measurement could make on its own. The tracking procedure has adopted an established algorithm to do just this: the *Kalman Filter*[80]. The Kalman filter was developed as a signals processing algorithm to clean up a true signal out




Figure 5.3: Shown is the application of the Kalman filter technique to track fitting [79]. The vertical lines correspond to the detector planes with indicated on it the measurement points and their errors. The cones represent the uncertainty in the reconstructed track parameters. As can be seen the knowledge of the track parameters is step wise updated with each measurement (Kalman filter technique applied from right to left).

of some noisy measurements, and it has since been adapted to many different uses, including particle physics track fitting.

Before the algorithm must be applied, one must first set up a physical model that will propagate one physical state to another; in the case here, behavior of linear kinematic motion that takes a particle from Stations 3 to 2 to 1 are defined, including the process of passing through KMAG and its fringe field. Once this is set up, the model is used to propagate the the state of the muon from one station to the other, evaluating at each step the kinematics, along with their associated uncertainties. This estimation of the measured parameters is combined with the state vector of the muon at that point to create a weighted average used to scale the measurements by their expected accuracy, which is largely dominated by the resolutions of the drift chambers. After the weighted average is calculated, the model propagates on to the next state. The algorithm repeats, ending up with a final state with a defined covariance.

Compared to the generalized Kalman Filter procedure where the propagation step is done analytically, the tracking algorithm uses Geant4 simulation to perform a dynamic state propagation from one station to another. This method is preferable for our purposes since Geant4 is well-tested and robust in its ability to handle relativistic interactions with matter and properly simulate the effects of fringe magnetic fields. Also, due to the procedure by which SeaQuest reconstructs tracks and the backwards-forwards time invariance of particle kinematics, the Kalman Filter is used to step *backwards* in time. It is primarily done in this way since the downstream end of the experiment is generally “cleaner” than the upstream end, and the straight path between Stations 3 and 4 are generally well-defined. The Kalman algorithm is iterative, and starting

out with a less uncertain starting state leads it to converge to a stable final track state faster (fewer epochs).

In the fitting of tracks, a state vector consists of the momentum and the position (p_x, p_y, p_z, x, y, z) . The tracklets, which allow for a good starting state, are then propagated upstream towards Station 1, with the measurements being updated as more steps are taken. Figure 5.3 provides a depiction of stepping (from right to left) through the stations with the uncertainty in the observables updating at each step. The uncertainties of the measurements shown stem from known chamber resolutions. The propagation between Stations 2 and 1 and between Stations 1 and the target area are performed using measured magnetic field maps and Geant4 simulations of particle interactions through matter. At the projected target area, the final observables are estimated. At this point, the process is reversed, and the particle is propagated forward in time, finely adjusting the parameters to minimize the χ^2 of the track. If the χ^2 of the global track converges, then it is passed on to the vertex finding/fitting stage of the tracking algorithm.

5.1.4 Vertex Fitting

With the fitted global tracks, they are matched up with each other to create a two-particle vertex with all its kinematics (combined from constituent tracks) estimated, along with a goodness of fit of the tracks to the vertex itself. The vertex-finding and -fitting algorithm used in the kTracker is adapted from the approach taken by Gorbunov and Kisel at the ALICE collaboration[81]. This approach, again, makes use of the Kalman Filter algorithm.

Before tracks can be combined to vertices, they are propagated upstream from Station 1 through FMAG via a method called a magnet ‘swim.’ In this process, the solid iron magnet is considered as many small slices in z , and a commensurate amount of energy loss and p_T -kick is applied at each z -slice. Once it emerges from the upstream face of FMAG, the track is then projected as a straight line to the target region, and it is these straight segments which are used as inputs to the vertex fitting process.

In the vertex fitting, a state is defined as being the position of the (presumed) dimuon vertex (x_d, y_d, z_d) . This is first initialized to $(0, 0, \bar{z})$, where \bar{z} is the average of all tracks’ points of closest approach to the z -axis. Pairs of tracks are matched up, and the vertex is then allowed to move around in order to minimize the χ^2 of this vertex fit. If the vertex position converges, then the reconstruction is complete and the vertex is stored. This can be done for any combination of tracks in a given event.

5.1.5 Monte Carlo Validation

With the tracking algorithm established, the next logical step is to measure its accuracy and efficiency, given a sample of perfectly clean Monte Carlo generated events. These events use a geometry definition




Figure 5.4: The residuals of reconstructed track kinematics against MC truth values.




Figure 5.5: The residuals of fitted track vertex kinematics against MC truth values.

that represents the latest survey numbers describing the hall, and they also simulate the ‘digitization’ of the hits (e.g. spatial coordinates in detectors into wire hits). Tracking is then run on these clean events to see if the developed algorithm can extract what is generated. In Figures 5.4 and 5.5, the residuals are shown, describing the reconstruction resolution of the track and vertex momenta. Overall, the tracking can reconstruct tracks and dimuon p_{tot} within a $\sigma = 1.8\%$ uncertainty. The residuals on track p_T/p_z is very good ($\sigma = 0.005$ mrad), but the effects of multiple scattering through FMAG and the vertex fitting smears out this to an residual uncertainty of $\sigma = 3.5$ mrad.

Regarding efficiencies on clean MC, the tracking was able to find and fit two muon tracks (i.e. χ^2 converged in fitting for both) at an efficiency of 92.2% of thrown events. Once the effects of FMAG and vertex fitting is factored in, the rate at which the tracking is able to find one muon pair vertex is 69.2%. This tracking efficiency and it’s target-, kinematic- and intensity-dependence will be investigated in great depth later on in the Rate Dependence section.

5.2 Analysis Cuts

The tracking software is concerned only with reconstructing possible tracks of relativistic charged particles through the SeaQuest spectrometer. It is, however, completely agnostic as to whether or not the charged particle was a muon that originated from a Drell-Yan interaction between the beam and the target. Though the tracker is agnostic as such, it should be noted that certain measures are in place in the design of the experiment that serve to increase the chances that a reconstructed track is a true DY muon track. Such examples are the choice of trigger matrix and the construction of an absorber wall before Station 4. Despite these measures, there are many cases in which we would wish to discard some signals, events, or whole spills in order to keep the data from becoming contaminated by anything but true $p - A$ Drell-Yan target dimuons.

As such, a set of cuts must be imposed to the tracked set of dimuon yields

- Drell-Yan Selection Cuts: Exclude kinematic regions where the signal is dominated by background
- Target Selection Cuts: Dimuons can be created in our beam dump. These cuts seek to remove the signals that likely originate from the dump.
- Trigger Selection Cuts: Many triggers were used to take data, and we only wish to analyze data triggered by a certain trigger configuration.
- Intensity Cuts: Data taken at unusually high intensities where components of the experiment can perform abnormally should be excluded.
- Quality Cuts: Most cuts fall under this umbrella. These include, but are not limited to: hardware settings, DAQ quirks, dimuon/track goodness-of-fits, etc.

In this section, these cuts are discussed in the order of tiers of data taking: Run, Spill, Event, Dimuon, Track. The intent with all of these cuts is to be as conservative as possible to ensure that all possible contamination is removed from the data despite the inevitable reduction in statistics.

5.2.1 Data Selection

The first cut on analyzable data is the choice of dataset to use. At SeaQuest, the blocks of data are broken up into what are called **Roadsets**. A Roadset is technically the version of the firmware for the triggering system that defines what set of hodoscope combinations can fire the trigger. In practice, a Roadset is a contiguous block of data for which experimental conditions were stable.

This document presents and discusses the EMC ratio results for SeaQuest roadsets #57, #62, and #67. Each of these roadset settings were in place for the span of months and contain varying amounts of statistics.

Table 5.2: The experiment-specific data sources and their details.

Schema	Date Range	# Dimuons	Roadset Notes
merged_roadset57_R005_V001	June 25 2014 - Aug. 20 2014	10.8 M	
merged_roadset62_R005_V001	Nov. 8 2014 - Jan. 14 2015	13.3 M	Hot roads removed RF-clocked firmware
merged_roadset67_R005_V001	Jan. 25 2015 - June 19 2015	40.8 M	Magnet-flipped roads hot roads re-evaluated

In total, many millions of tracks and dimuons of varying quality are found through the track reconstruction and vertex fitting procedure discussed in the previous section. Several cuts regarding quality of track/dimuon fit and experimental conditions are discussed below and applied to these datasets in order to isolate signal dimuons from the *target*.

The names of the SeaQuest MySQL merged productions along with their date range, magnet polarity, respective statistics, and notes regarding roadset differences are presented in Table 5.2.

5.2.2 Run Level Cuts

It is an outspoken goal of the SeaQuest collaboration to absolutely keep as much data as possible. This means, in general, that the largest chunk of data that we want to allow to be thrown away is on a spill-level. That withstanding, there are a certain set of factors that could lead an entire run’s (~ 1 hr’s) worth of data to be excluded.

One factor is the amperage settings of the focusing and spectrometer magnets. If these are not set nominally (2000 A and 1600 A, respectively), the tracking will improperly reconstruct tracks, as it depends on a specific p_T -kick value from each magnet in order to discern sensible tracks from the data. For this reason, if the averaged magnet currents are below 50% of their nominal values, the run is excluded from analysis.

Another crucial characteristic for a run is if it contains Spill values. This can be attributable to a malfunction in the Beginning of Spill / End of Spill triggers or a malfunction in the Spill Counter function. Regardless of cause, without this spill label, it becomes impossible to relate data taken to any other critical values such as target position or beam intensity.

The rest of the run-level cuts can be summarized by the umbrella term of “DAQ issues”. If the Slow Control failed to be recorded for any reason and there are zero Beam, Target, or Environment table values in a run, it is excluded. The same goes for the BeamDAQ and Scaler readouts. Perhaps most prevalent of the DAQ issues is the case of a run that failed upon starting. These resulted in empty 32 kB files with no usable data.

Beyond these issues, there are also ranges of runs for which there was some larger issue at hand. The best example of this is the period of data taking during Roadset 62 where the automatic target control was malfunctioning. The shift crew controlled the target position manually, but the actual target that was in the path of the beam during the spill could not be verified. As such, these runs were excluded. There is very little indication from the data itself that would incriminate itself as being untrustworthy, and so these are specifically excluded by analyzers. In practice, the runs are removed from the data set via an excluded range of spillIDs.

5.2.3 Spill Level Cuts

Seeing as most of the experimental readouts pertaining to the condition of the experiment are read out only once per spill, it is a natural result that a large variety of quality cuts that are applied are evaluated at the spill level. A summary of all such cuts are summarized in Table 5.3, but I will discuss a few key conditions here briefly.

One feature of a spill is the “dutyfactor53MHz”, which is the aggregated duty factor ($\frac{\langle I \rangle^2}{\langle I^2 \rangle}$) measurement over the whole spill as calculated by the Beam Intensity Monitor discussed in an earlier chapter. This an important measure that directly describes the quality of the beam provided to the experiment. A value that is too low corresponds to a spill that had the majority of its protons in just a small number of RF buckets. A duty factor value that is too high exceeds the capability of the main injector, and must therefore be erroneous.

The ratio of “acceptedMatrix1” to “afterInhMatrix1” is another quality factor that may not be intuitively obvious. The acceptedMatrix1 value is the number of accepted triggers recorded by the DAQ, and the afterInhMatrix1 value is the total number of FPGA1-triggered events that were not inhibited by the BIM. Putting bounds on the ratio of these two allows the exclusion of two things: (1) spills with high DAQ dead time and (2) spills where, for whatever reason, only part of the spill was recorded. The latter of these could be the case if a spill was being partially recorded after data taking for a run had just begun.

Other bounds such as the “S:G2SEM” and “QIESum” are measures of total delivered beam in units of protons and QIE, respectively. By putting bounds on these values, we avoid spills that are empty (no beam delivered) and spills that have far too much beam delivered. The lower bounds are non-zero due to a pedestal in each measurement.

Other sweeping criteria for qualifying a spill as good are that spillID \neq 0 and that there must be no duplicates of any of the values listed in Table 5.3. It is necessary due to the relational nature of the data that there must be a real spillID by which one can determine all other spill-relevant values; cases where

Table	Value	Roadsets	Good range
Spill	dataQuality	all	0
	targetPos	all	$\in \{1, 2, 3, 4, 5, 6, 7\}$
Target	TARGPOS_CONTROL	all	Match Spill tbl.
Scaler	TSgo	57, 59	[1e3, 8e3]
		61	[1e3, 1.2e4]
		62, 67, 70	[1e2, 6e3]
	acceptedMatrix1	57, 59	[1e3, 8e3]
		61	[1e3, 1.2e4]
		62, 67, 70	[1e2, 6e3]
	afterInhMatrix1	57, 59	[1e3, 3e4]
		61	[1e3, 1e5]
		62, 67, 70	[1e2, 1e4]
	$\frac{\text{acceptedMatrix1}}{\text{afterInhMatrix1}}$	57, 59	[0.2, 0.9]
		61	[0.0, 0.9]
		62, 67, 70	[0.2, 1.05]
Beam	S:G2SEM	all	[2e12, 1e13]
	F:NM3S (FMAG Current)	57, 59, 62, 67, 70	> 1000
		61	[200, 500]
F:NM4AN (KMAG Current)	all	> 1000	
BeamDAQ	QIESum	all	[4e10, 1e12]
	trigger_sum_no_inhibit	all	[4e9, 1e11]
	inhibit_block_sum	57, 59, 61	[4e9, 1e11]
		62, 67, 70	[4e9, 2e11]
	dutyfactor53MHz	57, 59, 61	[15, 60]
62, 67, 70		[10, 60]	
kTrack	#Tracks/Spill	all	> 0

Table 5.3: Some spill data quality criteria.

Min SpillID	Max SpillID	Details
371870	376533	Trigger timing shift
378366	379333	Trigger timing shift
394287	409540	Run 3 commissioning period
394287	414555	LD2 flask filled with LH2
416207	424180	Manual target control
482574	484924	Magnet flipped before roadset changed
526201	526364	Improper QIE inhibit setting
551030	610059	Track reconstructing bug distorting kinematics
581369	582460	KMag off
675814	684041	Drift Chamber 1 gas flow problem
684663	689443	Incorrect D3p in-time window

Table 5.4: Specific spill ranges of excluded data for all roadsets.

spillID=0 are cases in which no Spill Counter event was available, and these should be excluded. Regarding duplicate values, if such exist for one spillID, it is almost impossible which is the correct one, so the spill should not be considered.

Finally, certain periods of data taking are known by the collaboration to be problematic and not suitable for analysis. There exist periods of time where the trigger timing was off by a certain margin and degraded the efficacy of the trigger’s ability to trigger on dimuon events. There’s also a range of spills for which the magnets were flipped before a new trigger roadset could be devised for such a configuration. These are defined as a series of “bad spill ranges” and can be found summarized in Table 5.4.

In the Appendix on database design, there is a discussion regarding the use and implementation of the “dataQuality” bitpattern field (3.4). Here, this field is used as a bitmask in order to identify whether or not a spill passes or fails a given cut by flagging a bit as 0 (pass) or 1 (fail). This allows for analyzers to easily select only spills that pass all criteria by simply requiring $dataQuality = 0$.

5.2.4 Event Level Cuts

The bare minimum requirement for the Event level data is that it must have an entry in all event-level tables: Event, kEvent (the kTracker’s event meta-data), and QIE. If any of these is missing, then the event is discarded. More specific criteria can be imposed to consider an entry to be “missing,” such as when a row in the QIE table reports all “-1” values for the BIM RF readouts. In these instances, there was a malfunction in the BIM, and the “-1” acts as a *sentinel value* to indicate that the information does not exist for that event.

In addition to missing data, certain criteria can be placed on the Event data in order to exclude poor-

quality or unwanted events. The simplest cut that is imposed is the requirement that the event that was triggered must have satisfied the FPGA-1 trigger. The different trigger configurations are designed on optimized for different purposes; the FPGA-1 trigger is designed to select high-signal, low-background Drell-Yan (opposite sign muon pair) events, and by removing other types of triggers such as the NIM-3 (random RF) and FPGA-4 (singles) that exist in the data, it is ensured that there will be no effect due to different triggering systems when performing analysis.

The cut on beam intensity is another that ensures that we exclude data that would introduce undue volatility and uncertainty into the measurements. Due to the varying beam intensity at SeaQuest, there are inefficiencies which are corrected via weighting of events (this process is described in great detail later). Cutting on very high intensities removes the region where uncertainties in the rate dependence fits can cause large uncertainties in correction weights.

The cuts described here and more are summarized in Table 5.5.

5.2.5 Dimuon Cuts

Dimuons are selected based on many criteria: the known acceptance of the spectrometer, physical constraints, goodness of vertex fits, and more. The precise description can be summarized by the following: we are looking for two tracks of particles with opposite signs that can be *well-fit* to originating from the target table. Stepping through the cuts and their justifications:

- **dx, dy, dz:** These define the fitted point of origin (“dimuon (x, y, z)”) for the dimuon pair. If the vertex was outside the (dx, dy) region about the beam path, then this is probably not a real dimuon. If it is outside the desired dz range, then it probably did not come from the target.
- **dpx, dpy, dpz:** These define the total momentum of the dimuon (“dimuon momentum (p_x, p_y, p_z)”). If $|p_{x/y}| > 2 \text{ GeV}$, then one or both muons would pass outside of the acceptance of the experiment. Likewise, if $|p_z| < 30 \text{ GeV}$, one or both of the muons would have such low momentum as to end up bent out of acceptance of the spectrometer by the magnets. Finally, since a 120 GeV proton beam is used for SeaQuest $p_z > 120 \text{ GeV}$ must only come from unrelated muons that are reconstructed to this high momentum value.
- **sgn(px1), sgn(px2):** These are terms for the momentum of the positive and negative track, respectively. Requiring a specific sign for the x -momentum for each track acts as a safety check that the tracking identified the charge of the particle correctly.

- **sgn(roadID1), sgn(roadID2)**: This is another check that ensures that the muons came from opposite sides of the experiment (top-bottom). Such should be the case when triggered by a FPGA-1 event, but this enforces it.
- **trackSeparation, chisq_dimuon**: These are measures of the goodness of fit that two tracks came together to form at a single vertex. The track separation is the distance between the points of closest approach from each track to the beam axis (z-axis). The χ^2_{dimuon} is a measure of the quality of tracks combined with how well the two tracks converge to a single point.
- **mass**: There is a di-lepton range of invariant masses that are book-ended by the J/Ψ and the Υ in which Drell-Yan should be the only source. A cut on this range excludes any background that would be introduced by real J/Ψ s and Υ s.

5.2.6 Track Cuts

Going even deeper to remove signal contamination, if one of the two tracks for even a *good* dimuon is found to fail certain checks, the entire event is excluded from analysis. Again, stepping through each cut:

- **numHits**: A track passes through 18 drift chamber planes. Even if the chambers are somewhat inefficient, it is unlikely that the number of hits corresponding to a single track would be less than 15.
- **roadID \neq 0**: It is possible that a reconstructed track passes from the top half of the experiment to the bottom (and perhaps bends back up to the top). By the roadID naming convention, these are assigned a roadID of 0, and since we are looking to only analyze cases where there is one “top” track and one “bottom” track, these should be excluded. By imposing this requirement, such tracks are removed.
- **chisq/(numHits-5)**: Each track is described as a muon with five physical parameters ($x_1, x_2, p_T, \theta, \phi$). This expression described the χ^2 per degree of freedom, and can be cut on to select only good quality tracks that fit the hit pattern well.
- **pz1**: This is the momentum of the track at Station 1. In studies of signal background sources, one of the primary contributors is combinatorial background, which can be analyzed by looking to like-sign triggered events and studying their kinematics. As can be seen in Figure 5.6, depending on the number of hits, applying a cut on pz1 can remove the strong peak of combinatoric background from unrelated muons.

Table	Value	Good range	Type
kTrack	numHits	> 14	Quality
	roadID	$\neq 0$	Quality
	z0	$(-400, 200)$	Quality
	chisq/(numHits - 5)	< 5	Quality
	pz1 (Where numHits $\neq 18$)	> 18	Quality
	chisq_dump - chisq_target	> 10	Target
kDimuon	dx	$(-2, 2)$	Quality
	dy	$(-2, 2)$	Quality
	dz	$(-300, 200)$	Quality
	dpx	$(-3, 3)$	Quality
	dpy	$(-3, 3)$	Quality
	dpz	$(30, 120)$	Quality
	xB	$(0, 1)$	Quality
	xT	$(0, 1)$	Quality
	xF	$(-1, 1)$	Quality
	trackSeparation	$(-250, 250)$	Quality
	chisq_dimuon	< 15	Quality
	px1	> 0	Quality
	px2	< 0	Quality
	dz	$(-300, -60)$	Target
	mass	> 4.2	Drell-Yan
Event	MATRIX1	1	Trigger
kEvent	status	0	Quality
QIE	$\forall(\text{RF}+08, \text{RF}+07, \dots \text{RF}-08)$	$< \text{Beam.Inh.thres}$	Quality
	$\forall(\text{RF}\pm XY)$	> 0	Quality
	triggerCount	≥ 0	Quality
	intensity_p	< 60000	Quality

Table 5.5: A list of all quality and kinematic cuts for selecting good quality dimuons from the target.

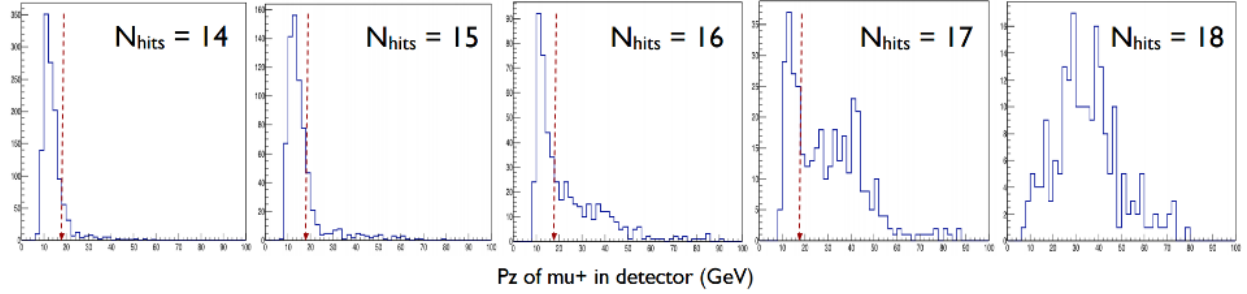


Figure 5.6: The p_z distribution of muons that form like-sign dimuon pairs, which is a source of background for DY signal events. In cases where $14 \leq \text{numHits} < 18$, imposing a cut on $p_z > 18$ GeV removes a majority of this background.

- **chisq_dump-chisq_target:** These are the χ^2 s of a track when forced to go through the beam dump and the target, respectively. By requiring this quantity to be above a certain value means, roughly speaking, that the track is very likely to have come from the target and not the beam dump.

5.3 Analysis Data Set

After enforcing all of the aforementioned data reduction cuts for both physics and cautious quality reasons, the resulting data set is, in total, 61,667 dimuons over three run periods (Roadsets 57, 62, and 67). The breakdown of raw statistics can be found in Table 5.6. This number does not include the empty flask and “none” targets, which will be used in correcting for background signals in the next chapter.

Table 5.6: Total raw dimuon yields after all spill, track, dimuon, target, and quality cuts. Not shown: empty flask target and “none” target (for background study).

Target	x_2	Roadset 57	Roadset 62	Roadset 67
LH2	(0.1, 0.13]	240	378	813
	(0.13, 0.16]	662	1000	2274
	(0.16, 0.195]	779	1247	2710
	(0.195, 0.24]	591	987	2211
	(0.24, 0.29]	320	568	1148
	(0.29, 0.35]	171	265	593
	(0.35, 0.45]	58	106	246
	(0.45, 0.58]	8	18	54
LD2	(0.1, 0.13]	277	384	1014
	(0.13, 0.16]	750	1011	2672
	(0.16, 0.195]	870	1263	3303
	(0.195, 0.24]	667	952	2513
	(0.24, 0.29]	351	471	1435
	(0.29, 0.35]	184	268	682
	(0.35, 0.45]	87	121	291
	(0.45, 0.58]	15	14	45
C	(0.1, 0.13]	143	109	357
	(0.13, 0.16]	358	349	990
	(0.16, 0.195]	439	470	1147
	(0.195, 0.24]	326	374	853
	(0.24, 0.29]	167	213	507
	(0.29, 0.35]	84	100	219
	(0.35, 0.45]	42	41	127
	(0.45, 0.58]	2	8	10
Fe	(0.1, 0.13]	112	145	358
	(0.13, 0.16]	279	398	1075
	(0.16, 0.195]	345	509	1284
	(0.195, 0.24]	262	411	1020
	(0.24, 0.29]	125	212	564
	(0.29, 0.35]	90	94	274
	(0.35, 0.45]	32	41	138
	(0.45, 0.58]	5	8	16
W	(0.1, 0.13]	113	169	442
	(0.13, 0.16]	309	500	1239
	(0.16, 0.195]	420	609	1618
	(0.195, 0.24]	288	488	1262
	(0.24, 0.29]	153	223	659
	(0.29, 0.35]	95	136	346
	(0.35, 0.45]	31	67	144
	(0.45, 0.58]	5	6	26
Total:		10255	14733	36679

Chapter 6

Analysis

There are many steps to get the data from this point to a state suitable for analysis. Here, the procedure is outlined for how to get the measured and pared-down dimuon yields into the desired ratio measurement. The greatest challenge of the analysis is the handling of what is known as the *rate dependence*, which is gone into with great detail. In addition, there is discussed target normalization, target contamination corrections, and isoscalar corrections for asymmetric nuclei. A summary of analysis steps (in the order performed) can be found in the final section.

6.1 Normalization of Relative Yields

The base measurement for the studies presented in this paper is, fundamentally, the number of dimuon events on a given target. This is known as the *raw yield* (Y). The ratio of these yields for a given target is used to study the kinematic dependence of the *per nucleon cross section ratio* that defines these phenomenon, but there must be a target-to-target normalization for such measurements.

$$R \equiv \frac{F_2^A}{F_2^D} \approx \frac{\sigma_{\mu\mu}^A}{\sigma_{\mu\mu}^D} \frac{2}{A} = N_{A/D} \frac{Y^A}{Y^D} \quad (6.1)$$

Here, A is the mass number of the target nucleus, and $N_{A/D}$ is the normalization factor, which can be expressed as:

$$N_{A/D} = \left(\frac{N_0^D \cdot lt^D \cdot T^D(\xi)}{N_0^A \cdot lt^A \cdot T^A(\xi)} \right) \cdot \left(\frac{2 \cdot n_D \cdot L^D}{A \cdot n_A \cdot L^A} \right) \cdot \frac{\varepsilon^D}{\varepsilon^A} \cdot \frac{\bar{\Omega}^D}{\bar{\Omega}^A} \cdot \frac{\epsilon^D}{\epsilon^A} \quad (6.2)$$

where N_0 is the number of incident protons, lt^A is the % live time on the target, T^A is the target transmission (or beam attenuation) factor, n_A is the number of nuclei per unit volume, L^A is the target thickness, ε^A is the detector efficiency, $\bar{\Omega}^A$ is the spectrometer acceptance, and ϵ^A is the tracking reconstruction efficiency. One of the key benefits of performing a ratio measurement is that, to good approximation, the terms regarding chamber efficiency, tracking efficiency, and acceptance should be equal across targets and should therefore cancel each other out. How well this assumption holds up will be assessed and evaluated.

Live Proton Sums (in $\times 10^{16}$ protons)			
Live Proton Target	Roadset 57	Roadset 62	Roadset 67
LH2	4.027	5.437	16.184
Empty	0.448	1.130	3.677
LD2	2.024	2.447	7.729
None	0.664	1.146	3.671
C	1.292	1.085	3.521
Fe	0.417	0.535	1.723
W	0.431	0.539	1.755

Table 6.1: Live proton sums for each target for each roadset.

With the normalization applied, this ratio R can be accurately studied by being binned in one or more kinematics in order to observe interesting behaviors and compare them against previous measurements and various models. Once the cuts described in the previous section are applied, the surviving dimuons are all considered to be *good quality target Drell-Yan dimuons*, and all analyses and corrections described from here on refer to them only.

This section will discuss the several components of the target-to-target relative normalization from the determination of beam intensity, target thickness, live time, chamber efficiency, and acceptance. Beam intensity and live time information was obtained from the BIM and scaler data. Acceptance was calculated with Monte Carlo simulation. Finally, chamber efficiencies and their ratios per target were studied with NIM3 random trigger data embedded onto clean MC. The tracking efficiencies are highly rate-dependent, and will be discussed at length in another section.

6.1.1 Live Proton Count and Beam Attenuation

The measurements of the number of incident protons, live time, and beam attenuation combine to give a figure for the total number of protons on target for which data was recorded. The important sources to consider here are the upstream SEM detectors (primarily the “G2SEM” readout) that record the number of incident protons, the QIE recorded sums of protons (in QIE units), the number of QIE units for which there was either a trigger inhibit or DAQ deadtime. In practice, SeaQuest calculates a value called “live protons” for each spill as:

$$N_0^A = \text{G2SEM} \quad ; \quad \text{lt}_A = \frac{\text{QIESum} - \text{inhibit} - \text{deadtime}}{\text{QIESum}} \quad ; \quad \text{liveProton} = N_0^A \cdot \text{lt}_A \quad (6.3)$$

this *liveProton* value is integrated, good spill by good spill, for every target in each data set.

It is also important to factor in the effects of inelastic scattering of the beam within the target, causing

Target	Target Position	Length[cm]	n_A [mol/cm ³]	ρ [g/cm ³]	NIL [g/cm ²]	λ [cm]	ξ : #Inter. Len.'s	T(ξ)
H_2	1	50.80	0.0702	0.0708	52.0	734.46	0.0691	0.966
D_2	3	50.80	0.0811	0.163	71.8	439.41	0.1156	0.944
Fe	5	1.91	0.141	7.874	132.1	16.78	0.1136	0.945
C	6	3.32	0.150	1.802	85.8	47.61	0.0698	0.966
W	7	0.95	0.105	19.300	191.9	9.94	0.0958	0.954

Table 6.2: Details pertaining to the measurables and physical properties individual nuclear targets.

the effective amount of beam-on-target to decrease with every step dz through the target. This is different for each target based on target nuclear densities and cross sections. With this, the liveProton value can be corrected. Consider incident proton number N_0 , where dN is the number of protons attenuated when passing through target slice dz , and L is the total length of the target. The function for the number of protons remaining at a given position z is

$$N(z) = N_0 e^{-\frac{\rho_A}{\text{NIL}_A} z} = N_0 e^{-z/\lambda} \quad (6.4)$$

where ρ_A is the density of the material, $\text{NIL}_A \text{ g/cm}^2$ is the nuclear interaction length of the target, and $\lambda = \frac{\text{NIL}_A}{\rho_A}$. When integrating over the length of the target, one arrives at the total [protons \cdot cm] exposed to the target. Assuming that each proton interacts independently, one can divide this integral by the length of the target to arrive at *effective* number of protons incident on the target

$$N_{\text{eff}} = \frac{1}{L} \left[N_0 \lambda (1 - e^{-\frac{L}{\lambda}}) \right] = N_0 \left[\frac{1 - e^{-\xi}}{\xi} \right] = N_0 T(\xi) \quad ; \quad \xi = L/\lambda. \quad (6.5)$$

The normalization of the transmission or beam attenuation factor T^D/T^A is therefore only a function of the length of the target, its density, and its nuclear interaction length. In total, the product of $N_0 \cdot \text{lt} \cdot T(\xi)$ represents the *effective protons on target* which was recorded by the data acquisition systems. Values of ξ and $T(\xi)$ can be found in Table 6.2.

6.1.2 Target Density Ratios

The number of nucleons that the effective protons will interact with should be considered when performing the target-to-target normalization. Aside from the background “targets” and liquid hydrogen, there are four nuclear targets in this A-dependence study. The nucleons in liquid deuterium are viewed as essentially free particles, and therefore proved an isoscalar reference by which the heavier targets are compared to. In all, the targets used are H_2 , D_2 , C, Fe, and W, and their attributes pertinent to target density ratios can

be found in Table 6.2. The yield ratios are normalized to the number of nucleons in each target by the normalization component $(2 \cdot n_D \cdot L^D)/(A \cdot n_A \cdot L^A)$.

An analysis of the pressures and temperatures were performed to conclude that the density of 0.0708 g/cm^3 for LH2 and 0.1634 g/cm^3 for LD2 with a $< 0.1\%$ uncertainty. In practice, since the deuterium target is contaminated by the presence of hydrogen, the density, mass number, and interaction length is adjusted to reflect the contamination. More on the contamination in Section 6.3.

6.1.3 Detector Efficiency

The key cause of detector inefficiency is high occupancies from high intensity events. Also, the total number of tracks that fired off elements of the many detectors is dominated by events generated from the 503 cm iron beam dump. The difference in interaction lengths of the targets should have only a small effect on the absolute chamber and hodoscope efficiencies. Further, any longer term time-dependent drifts in detector efficiencies will largely cancel out due to the frequent minute-by-minute interchange of targets. It is therefore justifiable to assume that there is no target-to-target chamber efficiency dependence to normalize to. The overall ratio of detector efficiencies between targets is therefore taken to be $\varepsilon^D/\varepsilon^A = 1$.

It is however possible that pions created in the targets have more of a chance to decay into muons that proceed to add to the background hit rates of the experiment; pions created in the beam dump are much more likely to get re-absorbed before they decay. To investigate the possible magnitude of this contribution to occupancy, detector occupancies were studied for the different targets in unbiased NIM3 random events as a function of chamber intensity. We see in Figure 6.1 that they are not significantly different from each other, though there is a slight difference in occupancies between the deuterium target and the rest. However, since the occupancy distributions at a given intensity bin do not change much for a given intensity, it is safe to assume that the efficiencies likewise do not change by much from target-to-target.

To override any concern of chamber efficiencies having an effect on any absolute or ratio measurements, a study was performed to measure the effect of chamber efficiencies on the track reconstruction algorithm. This was done by embedding clean Monte Carlo-generated Drell-Yan dimuons into real NIM3 random event hit data. Then, a constant 94% chamber efficiency was applied by removing a random 6% of the chamber hits, and tracking was performed on this sample. The 94% figure reflects the best estimate of overall chamber efficiencies at SeaQuest. Following that, a rate-dependent chamber efficiency was imposed of the form

$$\varepsilon(I) = 0.95 - a \cdot I \tag{6.6}$$




Figure 6.1: The Station 2 drift chamber occupancy for the four targets as a function of intensity. The error bars show a systematic deviation of occupancies from the mean for a given intensity bin. The inset plot shows a KDE plot of distributions for the intensity bin in which the values differ the most. The triangle markers represent the modes of the distributions.

where a was set to $2e-6 \text{ protons}^{-1}$. This resulted in a linear drop in efficiency with it dropping to as low as 75% at the highest intensity regions considered (100000 protons). The tracking algorithm was then run on this data sample.

For even the constant 94% efficiency case, the tracking is expected to decrease in *tracking efficiency* (a very different topic to tackle, as we will approach later in this chapter) as intensity increases and pattern recognition becomes more difficult as overall occupancies increase. The important observation to make is how this intensity-dependent tracking efficiency compares to the case of decreasing efficiency. For both cases, the ratio of successful reconstructions to all thrown dimuons (before even the trigger acceptance is imposed) can be seen in Figure 6.2. It is perhaps surprising to see that even at high intensities, the linear drop in tracking efficiencies does not change between the two cases. The reason for this is due to the fact that the tracking algorithm is very *robust* in the way it is able to handle missing hits, as it only requires a few hits in each station (4 out of 6) to successfully find a track.

Putting this together:

- Differences in chamber efficiencies between targets could only arise from different chamber occupancies as a result of the varying targets.
- Chamber occupancies are only slightly different for deuterium when compared to the other nuclear targets.




Figure 6.2: The ratio of successful reconstruction to all thrown by MC for two chamber efficiency models. The linear decreases in efficiency are not significantly different from each other.

- Even in the case of vastly different efficiencies, the yields from tracking remain unchanged.
- As a result, the effect of chamber efficiency on the yields will be the same for all targets.
- $\varepsilon^D/\varepsilon^A$, a measure of how the chamber efficiency effects the deuterium target yields as compared to a nuclear targets, can therefore be considered to be 1.

6.1.4 Detector Acceptance

The *true* acceptance as a function of any set of kinematic variables \bar{x} can be expressed as:

$$A(x) = \lim_{N \rightarrow \infty} \frac{Y_{accepted}(\bar{x})}{Y_{thrown}(\bar{x})} ; N \equiv \int Y_{thrown}(\bar{x}) dx \quad (6.7)$$

where *thrown* denotes Drell-Yan simulated events from Monte Carlo which generates every kinematic value by sampling over the distributions of that variable with a random number generator. The simulation includes such processes as multiple scattering and energy loss that occur due to particle passage through materials (such as our targets and beam dumps). Survey numbers of the experimental hall are used to build an accurate representation of the apparatus and its sensitive detector regions. Muon pairs over much of the kinematic phase space and was weighted according to the known Drell-Yan cross section. This does not factor in any higher order nuclear effects or anything like \bar{d}/\bar{u} asymmetries.

The pairs that qualify as *accepted* are those that satisfy the following two criteria: (1) were able to trace their way through the spectrometer without leaving the sensitive range of the detectors at all stations, and

(2) they followed paths that would have satisfied the un-like-sign dimuon trigger (FPGA1). Such data was written to file in a format that matches real experimental data and is then ready for analysis. These events are weighted according to the kinematics and DY cross-section as mentioned above, but for analysis, these weights when analyzed were divided by N , the number of thrown events.

Approximately 800,000 dimuons were thrown for each target to have the acceptances analyzed. For an absolute acceptance, a study of a full 4π (or 2π) simulation is required, where 4π is the solid angle of an entire (semi-)sphere. In such simulations dimuons are thrown in every physically possible direction, and then one can derive the absolute acceptance of the target and/or spectrometer when taking into account this source of *thrown* events. However, for the purposes of these studies, we are only concerned about $\bar{\Omega}^D/\bar{\Omega}^A$, or the relative acceptance correction. As we see in Fig. 2.6, we see that though the spatial distribution of the liquid deuterium flask and the solid targets are similar, they are fundamentally differently distributed. The density of the deuterium in the flask is constant throughout the length of the flask while the disks are discrete and separated. As a result, it perhaps may be the case that due to beam attenuation, there is a spatial difference between where most dimuons are generated between the different targets. The cause of any acceptance difference, however, is moot; the goal here is to estimate the relative acceptance correction factor between deuterium and the nuclear targets.

The target dependence of the acceptance was studied via yield ratios of heavy targets to that of deuterium for the valid kinematic regions of the experiment. These regions are $4.2 \text{ GeV} < M_{\gamma^*} < 10 \text{ GeV}$, $0.1 < x_2 < 0.58$, $0.35 < x_1 < 0.85$, $0.0 < x_F < 0.8$, and $0.0 \text{ GeV} < p_T < 2.5 \text{ GeV}$. Since no nuclear effects are simulated by the Monte Carlo generator, this ratio should lie at 1; any deviation of this must be purely from the acceptance effects due to the slight differences in the geometries and densities of the targets. The results can be seen in Figure 6.3.

In order to determine an acceptance correction, the ratio values shown are compared to the line of unity. The parameter α in the expression $1 - \alpha \cdot R_{DY}(k) = 0$ is optimized and is used as this correction factor in the target-to-target normalization. The acceptances of the liquid deuterium is found to be $\alpha = 0.9989 \pm 0.0012$ smaller than that of the solid targets. As such, a correction factor of

$$\bar{\Omega}^D/\bar{\Omega}^A = 0.9989$$

is included in the target-to-target normalization. The systematic uncertainty of this correction is approximately 0.1%.

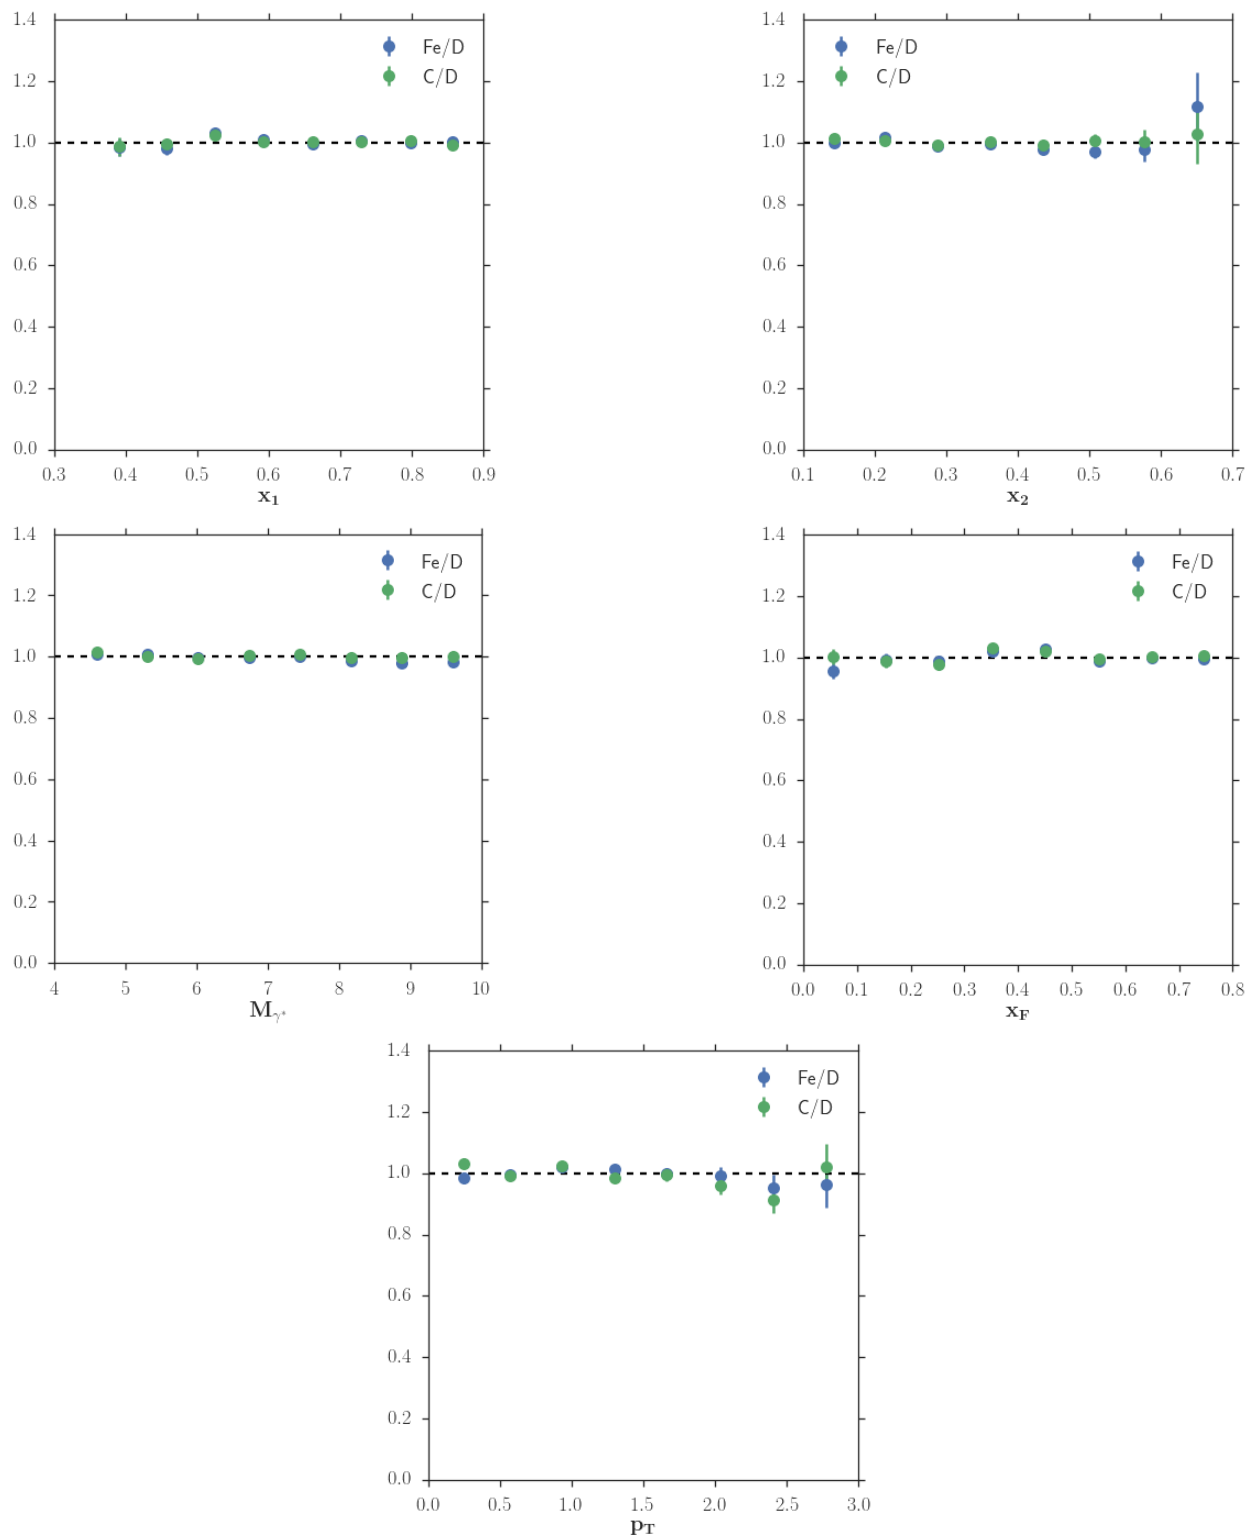


Figure 6.3: Ratios of accepted dimuons from heavy targets to deuterium for all kinematic regions of interest.




Figure 6.4: The rate dependence, depicting the decline in yields per triggering proton as a function of the overall environmental chamber intensity.




Figure 6.5: The rate dependence has a different behavior for different kinematics. Shown is the same plot as in Fig. 6.4, but in three statistically equal x_2 bins.

6.2 Rate Dependence

From the many generations of high energy experiments preceding SeaQuest, it was known that there might be an issue of *Rate Dependence* in the results of the data. This happens when conditions that lead to good reconstruction of signal events are affected by intensity or fluctuations in intensity of the beam delivered to the experiment. This was the case with E-772 [82] and E-866 [60], and so given the highly erratic nature of the beam from the slow spill extraction, such an effect was anticipated at SeaQuest.

It was uncovered in early 2015 that SeaQuest indeed had a rate dependence problem as well. This was done by investigating the number of reconstructed dimuons per trigger proton and plotting those yields against the chamber intensity (Figure 6.4). Here, we define the terms **trigger intensity** as referring to the number of protons in the RF bucket that triggered the event, and **chamber intensity** as a weighted sum of the number of protons ± 13 RF buckets about the triggered RF bucket. The former gives a measure of the number of protons on which the trigger could have fired while the latter gives a measure to the general intensity *environment* of the triggered event. It was observed that as the chamber intensity increases, the yields per trigger proton would drop precipitously, where the value is desired to remain stable at any intensity. The goal in this section is to apply a series of corrections to bring these plots show a flat series of values vs. chamber intensity.

To complicate the matter further, the rate dependence has a kinematic-dependence – or perhaps it is more accurate to say that the rate dependence affects some signal events with certain kinematics more severely than others. The kinematic dependence can be seen by simply showing the plot from Figure 6.4, but with the data split up into a few bins. The x_2 dependence of the rate dependence can be seen by observing the shape of the distributions in Figure 6.5. What this implies is that a single intensity-dependent

correction factor (as used in the two aforementioned experiments) would be inappropriate for correcting the data, and thus the correction becomes rather complicated.

The issue of rate dependence is discussed here, and as the sources of this effect are identified, a correction is proposed, calculated, and applied for each source in the form of an event weight. The two known sources addressed in this paper are the reconstruction efficiencies and empty target background corrections.

6.2.1 Weights as Corrections

Weighting of events is a standard procedure for mapping one distribution onto another. The most common example is with Monte Carlo (MC) simulation weighting. A simple physics event MC meant to simulate a well-defined process will ‘throw’ an event with certain kinematics randomly drawn from known distributions via *inverse transform sampling*. By doing this, the simulation is already close to being physical, but according to the cross section of a given process, a certain combination of kinematics will be decidedly more or less probable. A weighting subroutine calculates this likelihood for each event and assigns a weight with respect to how likely that event is based on the kinematics thrown. When binning this weighted data, the adjusted number of events in a bin is given by

$$N_{bin} = \sum_{i \in bin} w_i \ ; \ \sigma_{bin} = \sqrt{\sum_{i \in bin} w_i^2} \quad (6.8)$$

which reduces to the Poissonian statistical picture when all weights are 1.

Weighting in this manner can be used in any case where it is desirable to map one distribution to another, as in the case of applying the corrections discussed in this section. The most intuitive connection between the use of weights and corrections is in the case of efficiency corrections. Let us say that a dimuon with a certain set of kinematics, for whatever combination of effects, has a reconstruction efficiency of 80%. This means that four out of five actual dimuons with these kinematics will be reconstructed. If you know this efficiency ϵ , you can apply a weight as $w = 1/\epsilon_{recon}$ to each dimuon, as each observed event represents, in part, a larger population that is not fully represented. In the case of our example, a single dimuon represents $1/0.80 = 1.25$ dimuons in order to make up for the ones that are missed due to that specific inefficiency, and so this weight would be applied to the event.

6.2.2 Tracking Efficiency

As the intensity increases, the number of background hits on all of the detectors increases. As the number of hits in all the detectors increases, the harder it become for the tracking to perform its pattern recognition to

discern the actual dimuon(s) from the whole event. This is regarded to be an intensity-dependent tracking efficiency. Seeing as the primary tracking algorithm for SeaQuest is kTracker, this referred to at SeaQuest as the “**kEfficiency**”. Its intensity, target, and kinematic dependence makes it a factor to model and correct for.

The procedure developed by Evan McClellan of UIUC for calculating this kEfficiency is to compare the performance of the tracking on a clean MC sample of dimuons to an analogous sample that’s mixed in with real intensity-dependent background. To do this, a *single* sample of MC-generated dimuons and all of their hits is generated; this sample is denoted as the “*clean*” sample. Then, this *same* sample is mixed with all of the hits from a NIM3 event from the experimental data. The NIM3 trigger, as is discussed in Chapter 2, is the “random” trigger, and by using these events we avoid any effects due to trigger selection bias. This mix of MC dimuons and NIM3 triggered data is denoted as the “*messy*” sample. These NIM3 events all have an intensity, and by embedding a clean MC dimuon into a situation typical of an event at a given intensity, we get a basis by which to judge how efficiently the tracking can reconstruct a dimuon in the clean sample versus the messy sample as a function of intensity (and other kinematic variable).

It is important to note that the DY cross section is so small that there is a negligible chance that a given randomly-triggered NIM3 event will contain an actual dimuon, nonetheless a dimuon matching the characteristics of the embedded MC dimuon. As such, it is assumed that the dimuon embedded in the NIM3 event is the only dimuon to find in the event. It is also justified to assume that the kinematics of a dimuon is in no way correlated to the intensity, and the intensity of an event relates only to the amount of background hits observed. The relation to the intensity and the number of background hits can be observed in Figure 6.6.

The tracking software is run both clean and messy samples, and the two outputs are compared. The first step is to bin the resulting reconstructed dimuons into intensity bins. This binning uses the weights of the dimuons that is assigned by the GMC that created them. The weights are used instead to avoid undue influence on the efficiency from dimuons that are highly unlikely to occur. The value of the bin and its uncertainty are calculated according to Eq. 6.8.

For each matching bin, the ratio of $\frac{\#messy}{\#clean}$ values is used to calculate the efficiency of the tracking for that bin. The uncertainty of that efficiency is more complicated, as the relation between the messy and clean samples is of a *binomial* nature. Each dimuon successfully reconstructed in the messy sample represents a *positive* outcome from the number of possible outcomes which is represented by the clean sample. With this kind of relationship, binomial uncertainty calculations must apply. For large enough N number of trials and




Figure 6.6: The linear tendency of detector multiplicity to increase with intensity is observed by looking to unbiased NIM3 events. Here, the occupancy per event of the drift chambers at station 2 are shown as a function of intensity.

efficiency $\epsilon \notin \{0, 1\}$,

$$\epsilon = \frac{N_+}{N} \quad ; \quad \delta\epsilon = \sqrt{\frac{N_+ N_-}{N^3}} = \sqrt{\frac{\epsilon(1-\epsilon)}{N}} \quad (6.9)$$

but when weighted trials are involved, things become more complicated in calculating the statistical error.

The calculation of this statistical error [83] is as follows:

$$\epsilon = \frac{\sum_{i \in +} w_i}{\sum_i w_i} \quad (6.10)$$

$$\delta\epsilon = \frac{\sqrt{\sum_{i \in +} w_i^2 \left(\sum_{i \in -} w_i\right)^2 + \sum_{i \in -} w_i^2 \left(\sum_{i \in +} w_i\right)^2}}{\left(\sum_i w_i\right)^2} \quad (6.11)$$

where the sum of the weights (and sum of squares of the weights) for non-positive samples is approximated as

$$\sum_{i \in -} w_i = \sum_i w_i - \sum_{i \in +} w_i \quad (6.12)$$

$$\sum_{i \in -} w_i^2 = \sum_i w_i^2 - \sum_{i \in +} w_i^2. \quad (6.13)$$

For a set of efficiencies, an exponential function is fit to the efficiencies. This fit function would then be




Figure 6.7: The ratio of “messy” to “clean” dimuon samples from the Deuterium target in Roadset 67 data. This shows the linear relationship of the tracking efficiency to the intensity of the events. The shaded band indicates the 95% confidence band of the fit.

used to calculate the efficiency for any given dimuon that occurs at any given intensity. The single-parameter exponential function,

$$f(x) = e^{-Ax},$$

is chosen for its behavior at zero and infinity. The function of best fit along with a confidence interval for this can be seen in Figure 6.7.

Further, these fits have a statistically significant kinematic dependence, so it becomes important to unfold the kinematic space and perform this procedure for several kinematic bins in one or more dimensions to assure a robust efficiency calculation. Of the six defining kinematics of the dimuon, we examine the dependence of the fit on five of them: $(x_1, x_2, p_T, \theta, \phi)$, neglecting the azimuthal production angle ϕ_{γ^*} . Each of these five kinematics is divided into three statistically equal bins (low, medium, high) in order to observe the efficiency dependencies. The results are shown in Figure 6.8. The curves produced seem to indicate that there only exists a strong kinematic dependence on the x_1 and x_2 kinematics. If this is the case, then the clean and messy samples can be split two-dimensionally in these two kinematics, and a fit can be made to both variables. When analyzing dimuons, its kinematics can indicate which fit to use to calculate the tracking efficiency based on the intensity of the event, and a weight can be calculated.

But before we come to any clear conclusions, let us investigate two more kinematic phase spaces. We know from the discussion in Chapter 1 that $\{x_1, x_2\}$ can be used almost interchangeably with $\{x_F, M_{\gamma^*}\}$ to describe the same process. As such, it’s possible that if only one and not the other shows to influence the

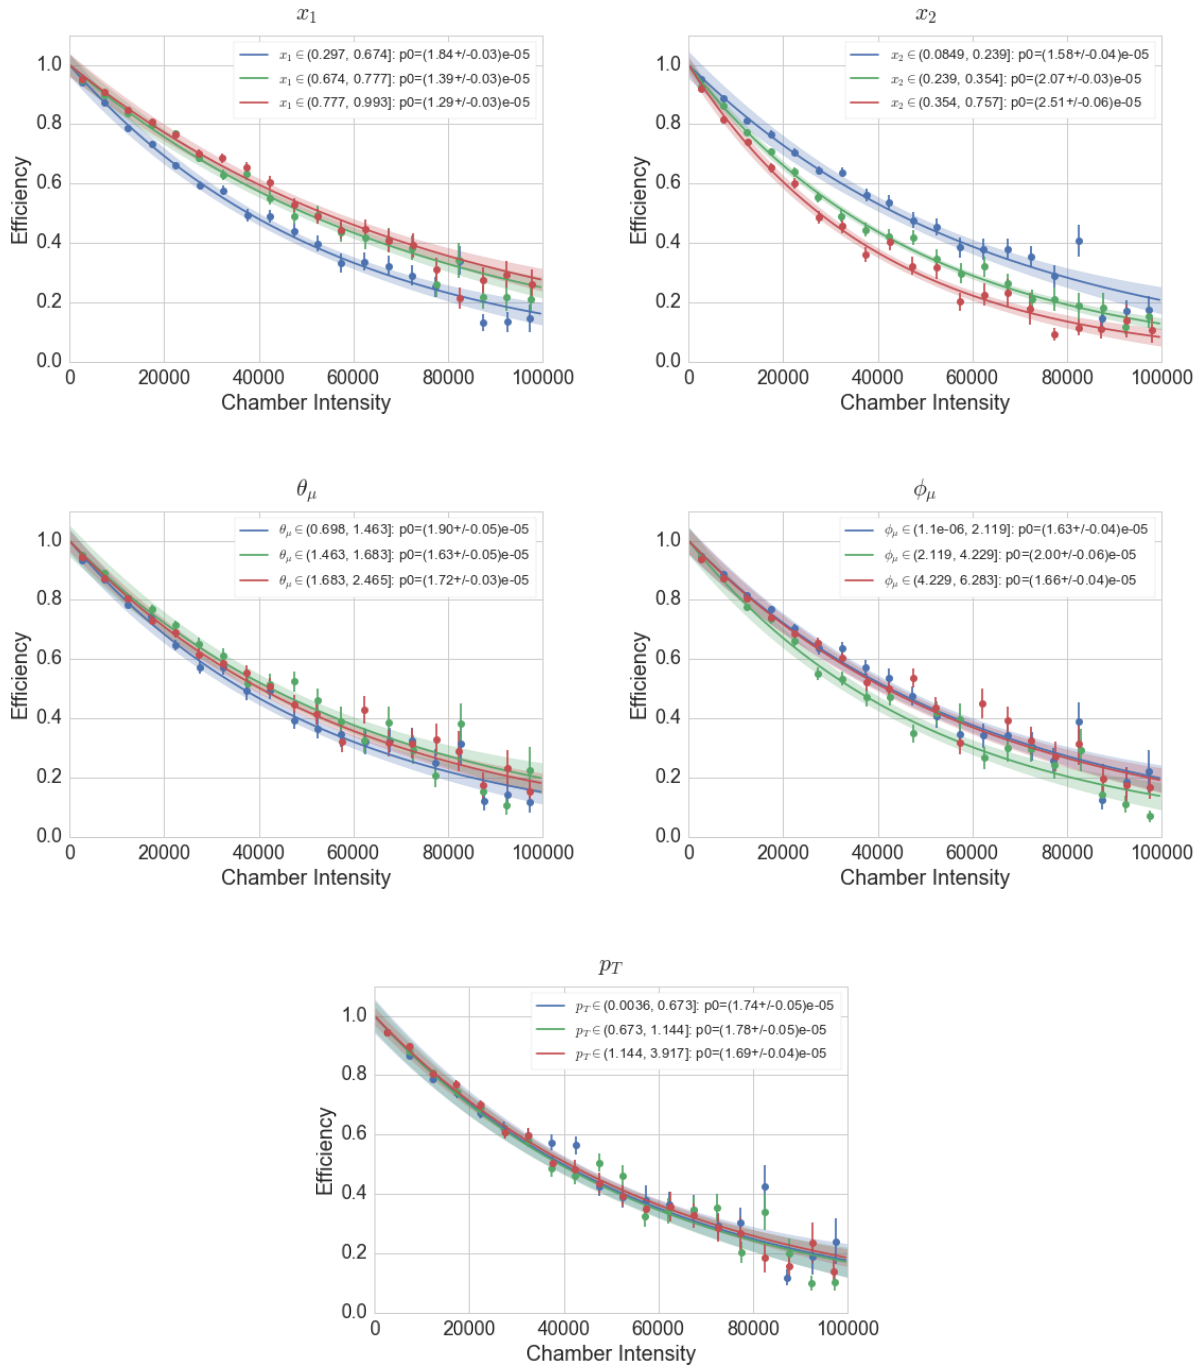


Figure 6.8: Tracking efficiency with the data broken into three statistically equivalent bins in the five primary kinematics, $(x_1, x_2, \theta_\mu, \phi_\mu, p_T)$. Significant kinematic dependence is seen in x_1 and x_2 .

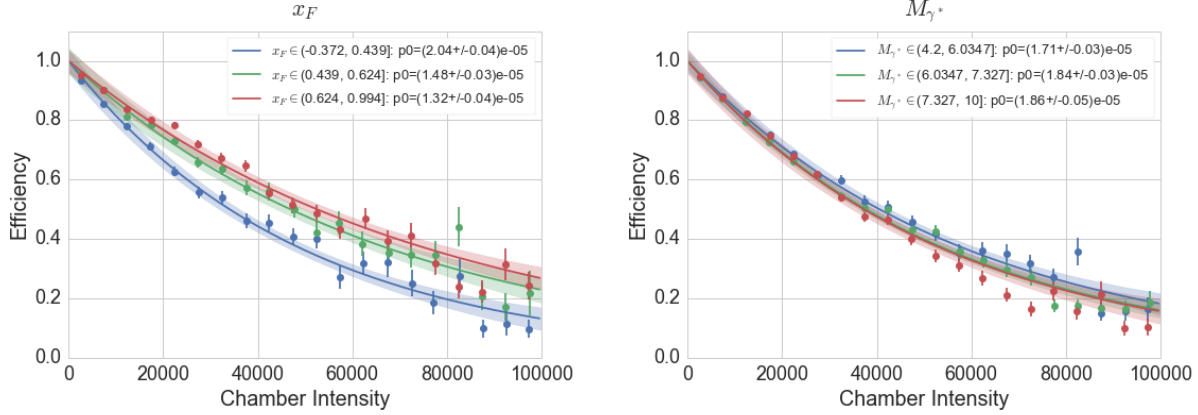


Figure 6.9: The kinematics x_F and M_{γ^*} are investigated. A clear kinematic dependence exists in x_F while M_{γ^*} is largely consistent.

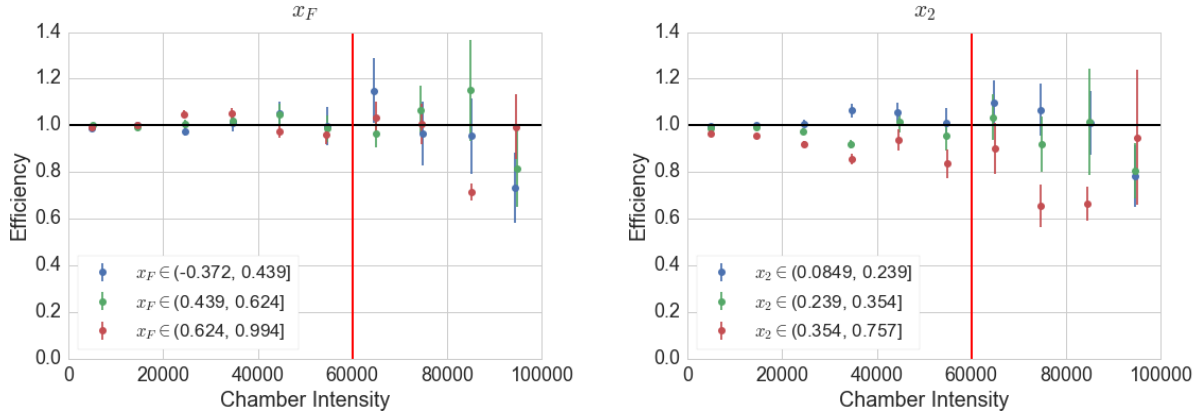


Figure 6.10: The efficiencies after re-weighting using kEfficiency values calculated when binned in x_F .

tracking efficiency curves, then only that one might be needed for the tracking efficiency correction. We see the behavior of the efficiency curves as a function of x_F and M_{γ^*} in Figure 6.9. It can be concluded that since there is no substantial mass dependence observed, the tracking efficiency has a kinematic dependence may be corrected by solely correcting on x_F .

In order to evaluate the effectiveness of a correction, the messy sample was reweighted with the kEfficiency factors, with the efficiencies calculated in bins of x_F , and then recalculate these efficiency plots. This is shown in Figure 6.10 for x_F and also x_2 to see if by correcting in one, the other is also corrected. It becomes evident that even though the efficiency certainly seems to be fixed in x_F , any analysis in x_2 may not be properly corrected. One can come to the conclusion from this that *if one performs an analysis on SeaQuest data in certain kinematics, one must correct for kEfficiency based on those kinematics*. It should be noted that above approximately 60,000 protons of chamber intensity, the tracking efficiency is low, and the weights become quite large. This combined with the fact that there are low statistics at high chamber intensity result in

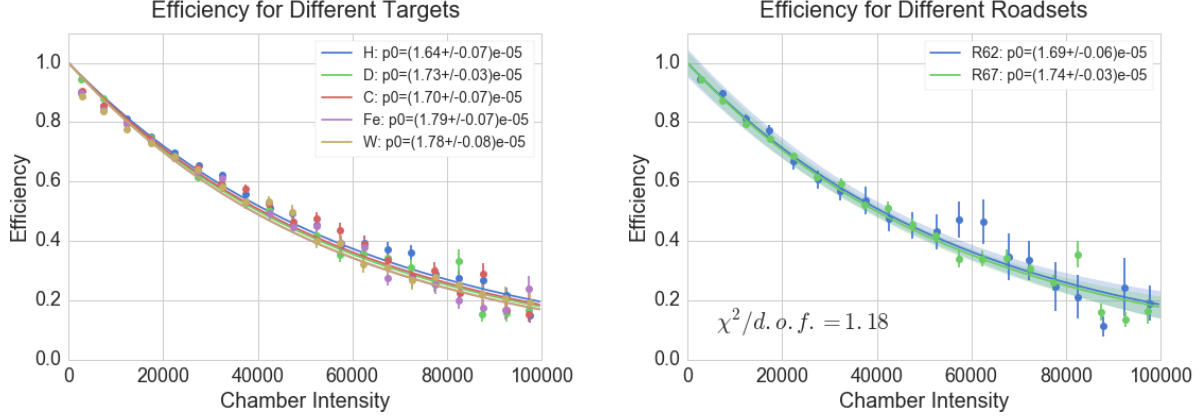


Figure 6.11: Comparing the tracking efficiency curves across (left) the five targets and (right) two temporally different subsets (different roadsets) of data. Confidence bands removed from the target plot for the sake of clarity.

those events possessing a greater influence *and* instability with respect to the yields. It is for this reason that events above this intensity threshold are removed from analysis. This boundary can be seen visually on Figure 6.10.

Another consideration is with respect to whether or not there is a target dependence to factor into the correction. For all of the above “kEfficiency” plots, only deuterium data is used. The kEfficiency curves for deuterium, hydrogen, carbon, iron, and tungsten can be found on Figure 6.11. It can be concluded that there is enough of a difference between deuterium, iron, and the rest to justify calculating and applying this correction on a target-by-target basis.

Also, as a sanity check, there should be no time-dependence of this function. A quick check comparing different sets of data from two different roadsets quickly confirm that there is no such time dependence (Fig. 6.11).

As a result, the final calculation of a tracking efficiency for a given dimuon event should depend on (1) target, (2) specific analysis kinematics, and (3) intensity. In order to apply this as a weight, we first define $kEff_{t,\bar{x}}(I)$, a set of exponential fit functions: one for each target (t) and for each kinematic bin (\bar{x}), which

	A	δA	$\chi^2/p.d.f.$
x_2 -bin			
(0.0856, 0.137]	1.39e-05	4.08e-07	0.80
(0.137, 0.153]	1.50e-05	9.09e-07	2.32
(0.153, 0.168]	1.58e-05	5.67e-07	0.74
(0.168, 0.183]	1.45e-05	7.32e-07	1.62
(0.183, 0.202]	1.63e-05	6.77e-07	1.36
(0.202, 0.227]	1.69e-05	7.01e-07	2.00
(0.227, 0.269]	1.92e-05	7.71e-07	3.32
(0.269, 0.679]	2.24e-05	3.78e-07	1.65

Table 6.3: LD2 kEfficiency fit parameters with the χ^2 evaluation of the fit to the data.

may be binned in any number of kinematic dimensions. The procedure in weighting each dimuon will then be to take each dimuon event i and its chamber intensity I_i , originating from target t_i in \bar{x}_i bin x_{F_i} , create a weight as

$$w_i = \frac{1}{\text{kEff}_{t_i, \bar{x}_i}(I_i)}. \quad (6.14)$$

When re-weighted in x_2 bins, the set of yields shown in Table 5.6 can be updated to what is seen in Tables 6.8, 6.9, and 6.10 at the end of the chapter.

6.2.3 Empty Target Background Correction

Upon inspection, the Empty/None target data has shown to exhibit a rate dependence, with many more events per proton being found at higher intensities. This can be corrected for, but the events should first be weighted by their tracking efficiencies. Unfortunately since there are no *messy/clean* embedded data sets for the empty flask, the “kEfficiency” fit must be extrapolated from the fits for the deuterium and hydrogen targets. If we are to assume that the difference in fits is a function of the beam’s interaction with the target, then we can arrive at the *kEff* exponential fit constants for the empty/none background targets (A_{bg}) for each kinematic bin from the exponential fit constants from the hydrogen and deuterium targets ($A_{D/H}$). Here, we only additionally need the number of interaction lengths in the targets (ξ) to find the proportion of the beam that did not interact with the target (X_t) The values used can be found in Table 6.2.

$$X_t = 1 - e^{-\xi t} \quad (6.15)$$

$$A_{bg} = A_H - X_H \cdot \frac{\Delta A}{\Delta X} \quad (6.16)$$

$$= A_H - X_H \cdot \frac{A_D - A_H}{X_D - X_H} \quad (6.17)$$

With this fit constant in hand and the kEfficiency correction applied, one can plot the yields of empty flask and none events as a function of chamber intensity. This plot and its shape can be found in Figure 6.12. Due to the low-statistics nature of the empty and none events, they are combined across all roadsets and kinematics instead of being broken down. As such, this correction is applied solely based on chamber intensity with no other consideration for roadset or kinematics. The fit that is applied is of the form: $f(I) = p_0 \cdot (1 + p_1 \cdot I_c^2)$. This form is chosen because the the scale of the final function will be a calculated from the data used, and as such, p_1 is the only parameter that is needed. The fit parameter p_1 was found




Figure 6.12: The intensity-dependence of the Empty target and "None" target, with fit to $p_0(1 + p_1 \cdot I_c^2)$ for each. 95% confidence band shown.

to be:

$$p_1^E = (1.03 \pm 0.23) \times 10^{-9} \quad (6.18)$$

$$p_1^N = (1.08 \pm 0.19) \times 10^{-9}. \quad (6.19)$$

The difference between the two backgrounds are interpreted to be due to interactions of the beam with the aluminum walls of the empty flask target.

The fit function used, $C_{norm}(I_c) = B \cdot C(I_c) = B \cdot (1 + p_1 \cdot I_c)$ can be interpreted as **the probability that a given target dimuon actually arose from something other than a target interaction** (beam dump, upstream interactions, flask container). Here, B is the normalization factor, where the condition should stand that the sum of $C_{norm}(I_c)$ over all target dimuons should be equal to the number of events from the empty and none targets. Both sums need to be corrected for kEfficiency and scaled by the number of live protons.

$$\sum_T \frac{B \cdot C(I_c)}{P_T \cdot K_T(I_c)} = \sum_E \frac{1}{P_E \cdot K_E(I_c)} \quad (6.20)$$

$$B = \frac{P_T}{P_E} \cdot \frac{\sum_E \frac{1}{K_E(I_c)}}{\sum_T \frac{C(I_c)}{K_T(I_c)}} \quad (6.21)$$

$$\delta B = \frac{P_T}{P_E} \cdot \frac{\sqrt{\sum_E \frac{1}{K_E(I_c)^2}}}{\sum_T \frac{C(I_c)}{K_T(I_c)}} \quad (6.22)$$

Here, P_T is the live protons incident on target T described in Table 6.1, and $K_{(T/E)}(I_c)$ is the kEfficiency correction function. Calculated B-values for all targets and roadsets can be found in Table 6.4.

This normalization constant is calculated for each target and each roadset. With this in hand, a weight for correcting for the empty target correction can be applied as $(1 - C_{norm}(I_c^2))$, which can be interpreted as **the probability that the given dimuon arose from an actual target DY event**. One can then combine this with the weighting for the kEfficiency correction (as the weighting should combine multiplicatively):

$$\epsilon_{tot} = \epsilon_{kEff} \cdot \epsilon_{empty} = \frac{1 - C_{norm}(I_c)}{K_T(I_c)} \quad (6.23)$$

It is worth noting that this model does not take into account the attenuation of beam throughout the different target materials, which could alter the scale of the correction slightly. However, without having a figure as to how much of the empty target signal comes from the beam dump and how much of it comes from upstream of the target area, it is currently infeasible to appropriately account for the degree to which a beam attenuation factor should be worked in.

Uncertainties that arise as a result of the uncertainties in the $C(I_x)$ fit and the calculation of B are difficult to propagate analytically. What is performed instead is an independent variation of p_1 and B by -1σ , 0 , and $+1\sigma$ and observing the magnitude of which the per nucleon cross section ratio changes as a result. This systematic uncertainty due to this correction is calculated in the Systematic Uncertainties section of this chapter.

		B	δB
Roadset	Target		
57	LH2	0.119	0.014
	LD2	0.057	0.006
	C	0.046	0.004
	Fe	0.019	0.002
	W	0.017	0.002
62	LH2	0.107	0.013
	LD2	0.051	0.006
	C	0.042	0.004
	Fe	0.019	0.002
	W	0.015	0.001
67	LH2	0.095	0.011
	LD2	0.039	0.005
	C	0.043	0.004
	Fe	0.018	0.002
	W	0.015	0.001

Table 6.4: The calculated normalization constants for the background proportion function, $C(I)$.

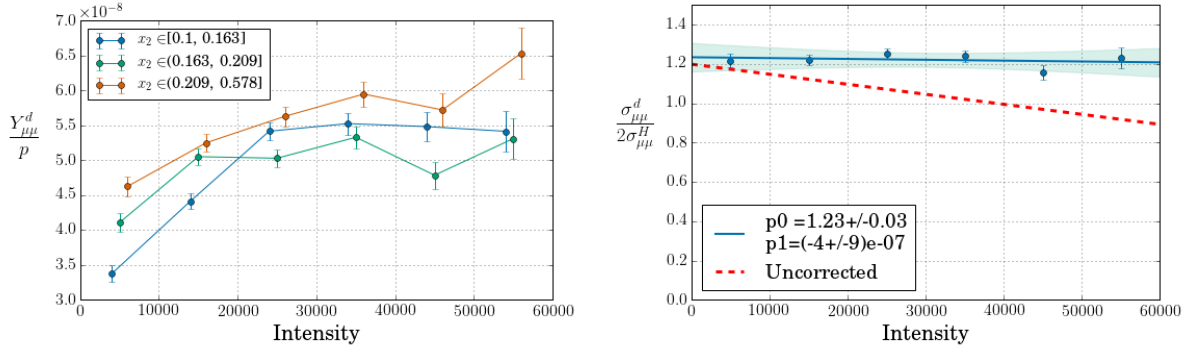


Figure 6.13: (Left) The remaining rate dependence behavior after corrections are made and (right) the rate dependence of a cross section ratio measurement with the linear fit and the 95% confidence band shown.

6.2.4 Residual Rate Dependent

It should be emphasized that *the issue of rate dependence is not a closed one*. As shown in this section, efforts have been put forth to make clever use of MC-generated data, tracking, and background information in order to identify and correct for as much of the rate dependence as possible. Due to the early nature of these approaches, there are bound to be factors which have yet to be accounted for, namely combinatoric backgrounds (uncorrelated single muons identified as signal dimuons). Plotting the yields per trigger proton again after both corrections (Fig. 6.13, left), one can see that there is indeed still some rate dependence, but overall, the yields per proton no longer significantly decrease at higher intensities. Now the question is how to handle the residual rate dependence – should an additional correction be applied, or should a systematic uncertainty be defined?

To decide, the cross section ratio measurement is taken and binned in chamber intensity. If the target-to-target and kinematic-dependent rate dependence is corrected, then this should be a flat line. The results as seen in Figure 6.13 (right) show that the ratio measurement is consistent with a line of zero slope. That is to say, the uncertainty in the slope is larger than the slope itself. Since this is the case, no further correction at this point will be applied, and a systematic uncertainty will be defined. Assuming that the rate dependence of the ratio measurement is a line, this is done by calculating the uncertainty of the fit σ_f over the value of the fit at the mean chamber intensity value. It should be noted that σ_f benefits from negative off-diagonal covariance elements for the fit parameters.

The systematic uncertainty here is calculated to be 1.07% of the cross section ratios.

		3230-11143		11144-11477		11478-14652	
Species:	Runs: Test:	1	2	1	2	1	2
H_2		1.23	1.225	0.25	0.25	4.603	4.591
HD		16.187	16.171	8.393	8.391	6.983	6.982
D_2		79.78	79.729	90.871	90.808	80.504	80.443
He		2.20	2.188	0.018	0.014	<0.01	<0.01
CH_4		<0.009	<0.009	<0.007	<0.007	<0.008	<0.008
H_2O		<0.009	0.073	0.038	0.099	0.033	0.061
HDO		0.015	0.026	0.024	0.033	0.021	0.04
D_2O		<0.007	<0.009	<0.007	<0.008	<0.009	<0.012
Ne		0.007	0.065	<0.027	<0.027	<0.03	<0.03
N_2		0.428	0.419	0.323	0.323	6.395	6.382
C_2H_6		<0.006	<0.006	<0.005	<0.005	<0.006	<0.006
O_2		0.085	0.098	0.068	0.068	1.389	1.416
Ar		0.005	0.006	0.004	0.003	0.063	0.064
CO_2		<0.002	<0.002	0.011	0.009	0.01	0.01
C_3H_8		<0.006	<0.006	<0.005	<0.005	<0.006	<0.006

Table 6.5: Results of the isotropic analyses of the liquid deuterium samples. Two tests are applied to each sample, and all numbers are in percentages.

6.3 LD2 Contamination

During Run III of the SeaQuest experiment, analysis of the deuterium source used indicated that it was impure to a significant degree. This calls for a significant correction since most measurements, yields, ratios, and analyses assume that the targets are pure. An impure deuterium target changes many factors including the averaged mass number of the fluid and its overall density. In this section, the contamination measurements and the correction procedure is discussed.

6.3.1 Contamination Measurements

The SeaQuest deuterium target was emptied and filled several times during data taking. At each of these times, a sample was taken from either the target or from the deuterium gas cylinder used to fill the target. These samples were sent to Los Alamos for isotopic analysis where each sample was analyzed twice. The results of these analyses are shown in Table 6.5. In each case of parallel analyses on a single gas sample, the two results were in agreement. Upon examination of the deuterium compositions, some problems are immediately clear. Specifically, the ratios of H_2 , HD , and D_2 are different between the three samples, but the bottles were supposed to have been filled from the same source many years ago. Second, some of the samples had a significant contribution of other gases, including N_2 and O_2 . This indicates an additional contamination of air. A sample of the first H_2 target was analyzed and found to be $\sim 98\%$ pure with negligible contributions from other species.

	Run	3230-	11144-	11478
	Range:	11143	11477	14652
Species:				
H_2		1.26	0.25	4.99
HD		16.65	8.44	7.59
D_2		82.08	91.31	87.42

Table 6.6: Renormalized compositions (in percentages) when considering only the hydrogen isotopes.

The source of the deuterium used at SeaQuest is from a bubble chamber experiment at Fermilab. At some undocumented point in the past, the contents of the bubble chamber were transferred into a tube trailer for storage. At a later time, the D_2 was transferred into cylinders for storage. Discussions with gas storage experts indicate that there was likely no special treatment of the cylinders other than being evacuated when they were filled with D_2 , with the desired special treatment being to “bake” the cylinders, or heat them up to release residual hydrogen that tends to diffuse into the metal surface of their inner walls. In addition to that, the containers used to take samples from the target or cylinders were also not “baked,” so existing estimates of the contamination may *only* be from the sampling containers – or the contamination could be doubly worsened by contamination in the sampling containers. As such, effects of this contamination are treated as a systematic error with asymmetric uncertainties. After run 14652, a commercial deuterium gas has been used which has an analyzed purity of $> 99.9999\%$, and no systematic uncertainty will need be applied.

Contamination from air in the sample gases is unlikely to have come from the targets. Prior to filling, the targets undergo several cycles of evacuation and refilling with deuterium before the final fill is done. This cycle should flush out all residual air from the target flask. Further, if there was air supply being fed into the target, all but the lowest temperature gases should have been removed by the liquid nitrogen “cold trap.” Aside from hydrogen isotopes, helium is the only other component found in the samples. Since the target temperature were approximately 23 K, any helium would have been in a gas state and not in the bulk liquid in the flask. Thus, the gases from the analysis shown in Table 6.5 can be renormalized to consider only the hydrogen isotopes, which can be seen in Table 6.6.

Until the true compositions can be ascertained, it has been recommend that analyses use the atomic percentages provided by the Los Alamos analysis as shown in Table 6.7, which represents the best understanding of the deuterium at this point. To assign a systematic uncertainty to these numbers, the recommended approach is to use asymmetric uncertainties, assuming that one the extreme values for H and D. Due to the fact that there is a complete inverse correlation between hydrogen and deuterium an uncertainty is assigned to the deuteron percentage only.

	Run	3230-	11144-	11478-	14653-
	Range:	11143	11477	14652	Present
Atom:					
H		9.6	4.5	8.8	0.0
D		90.4^{+5}_{-0}	95.5^{+0}_{-5}	91.2^{+4}_{-1}	100.0

Table 6.7: Atomic percentages of the deuterium targets with asymmetric uncertainties in the d percentage.

6.3.2 Contamination Correction

With the assumption that the heavier atoms are caught in the cold trap and the gaseous helium stays at the top of the flask and away from the path of the proton beam, a contamination correction is implemented based on the atomic percentage composition detailed in Table 6.7. The nuclear interaction length, density, and atomic mass of the liquid deuterium target is corrected based on the cotamination. With A as the atomic mass, λ as the nuclear interaction length, ρ as the density, and $F_{H/D}$ as the proportion of H or D in the target, the following are assumed to be approximately true:

$$A = A_H \cdot F_H + A_D \cdot F_D \quad (6.24) \quad \frac{1}{\lambda} = \frac{F_H}{\lambda_H} + \frac{F_D}{\lambda_D} \quad (6.26)$$

$$\rho = \rho_H \cdot F_H + \rho_D \cdot F_D \quad (6.25) \quad \lambda = \frac{\lambda_H \cdot \lambda_D}{F_H \cdot \lambda_D + F_D \cdot \lambda_H} \quad (6.27)$$

Due to the contamination, some of the signal events with the liquid deuterium target in place will actually come from hydrogen instead of a deuteron. With $\hat{\sigma}_{pH}$ and $\hat{\sigma}_{pD}$ representing the *measured* target cross section, and σ_{pp} and σ_{pn} representing the proton-proton and proton-neutron cross section,

$$\hat{\sigma}_{pH} = a \cdot \sigma_{pp} + b \cdot \sigma_{pn} \quad (6.28) \quad \sigma_{pp} = \frac{b \cdot \hat{\sigma}_{pD} - d \cdot \hat{\sigma}_{pH}}{b \cdot c - a \cdot d} \quad (6.30)$$

$$\hat{\sigma}_{pD} = c \cdot \sigma_{pp} + d \cdot \sigma_{pn} \quad (6.29) \quad \sigma_{pn} = \frac{a \cdot \hat{\sigma}_{pD} - c \cdot \hat{\sigma}_{pH}}{a \cdot d - b \cdot c} \quad (6.31)$$

Here, the coefficients a, b, c , and d denote the number of protons and neutrons per atom for the different targets. We know that $a = c = 1$, and since the hydrogen target is considered to be pure, $b = 0$. The remaining value, d , is the fraction of deuteron atoms in the sample (F_D above). With this, we can arrive at

an expression for $\sigma_{pD}^{\text{pure}}$.

$$\sigma_{pH}^{\text{pure}} = \sigma_{pp} = \hat{\sigma}_{pH} \quad (6.32)$$

$$\sigma_{pn} = \frac{\hat{\sigma}_{pD} - \hat{\sigma}_{pH}}{d} \quad (6.33)$$

$$\sigma_{pD}^{\text{pure}} = \sigma_{pp} + \sigma_{pn} \quad (6.34)$$

$$\sigma_{pD}^{\text{pure}} = \frac{\hat{\sigma}_{pD} - \hat{\sigma}_{pH} \cdot (1 - d)}{d} \quad (6.35)$$

In practice, this correction is applied bin-by-bin, utilizing the weighted dimuon yield for each target in place of $\hat{\sigma}$. The d used depends on the run ranges applicable to the roadset being analyzed, and the values are derived from those in Table 6.7. It's the case that the deuterium was changed in the middle of the roadset – twice for Roadset 62 and twice again for for Roadset 67. As a result, the contaminations are averaged, proportional to the number of runs with a given sample, and generous systematic uncertainties are allowed with the understanding that the deuterium compositions will be re-analyzed and known more exactly in the future. The values of $91.3 \pm 4\%$ is used for Roadset 62 and $95.3 \pm 3\%$ is used for Roadset 67. In order to arrive at a systematic uncertainty due to the uncertainty in the knowledge of the contamination, the analysis is performed at the with values at nominal, $+\sigma$, and $-\sigma$ are performed, with the largest % change in the ratio value to be valued as the systematic uncertainty.

6.4 Isoscalar Corrections for ^{184}W and ^{56}Fe

The proton and neutron are similar in many ways, but they do not share the same F_2 structure function. This is primarily due it being a charge-weighted summation of its quark constituents. So, taking the ratio of cross sections for asymmetric nuclei ($N \neq Z$) to deuterium ($N = Z = 1$) will not result in a proper extraction of the ratio of F_2 . to correct for this, following Aubert *et al.* and Bodek *et al.* [23, 24], an additional isoscalar correction factor (R_{ISO}) is applied to the yield ratio, approximating the ratio to a hypothetical nucleus with equal numbers of protons and neutrons ($N = Z = A/2$):

$$\frac{\sigma_{DY}^A(x, Q^2)_{ISO}}{\sigma_{DY}^D(x, Q^2)_{ISO}} = \frac{\sigma_{DY}^A(x, Q^2)}{\sigma_{DY}^D(x, Q^2)} \cdot R_{ISO}(x, A, Z) \quad (6.36)$$

where $R_{ISO}(x, A, Z)$ is expressed as

$$R_{ISO}(x, A, Z) = \frac{A}{2} \cdot \frac{F_2^p(x, Q^2) + F_2^n(x, Q^2)}{Z \cdot F_2^p(x, Q^2) + N \cdot F_2^n(x, Q^2)} = \frac{A}{2} \cdot \frac{1 + R_{np}(x, Q^2)}{Z + N \cdot R_{np}(x, Q^2)} \quad (6.37)$$




Figure 6.14: (Left) The ratio of free neutron to proton F_2 structure functions following a detailed parametrization [84] and the isoscalar correction factor [85], $R_{ISO}(x)$, for iron and tungsten.

and $R_{np}(x)$ is defined as

$$R_{np}(x, Q^2) = \frac{F_2^n(x, Q^2)}{F_2^p(x, Q^2)}. \quad (6.38)$$

The ratio of the free nucleon structure functions for neutrons to protons is extracted from world data on DIS scattering [84] off of deuterium and hydrogen with corrections accounting for Fermi motion of nucleons in deuterium. The behavior of $R_{np}(x)$ and $R_{ISO}(x, A, Z)$ for the case of iron can be seen in Figure 6.14.

6.5 Summary of Analysis Steps

Due to the dynamic nature of the SeaQuest analysis at the time of the writing of this document, it is important to lay out the precise procedure of the application of all of these analysis steps. This ensures an ability to reproduce results, understand differences between methods, and allows those new to analysis to have a roadmap for getting started with analysis. Pragmatically speaking, the process taken in this analysis going from the raw dimuon yields to the final extraction of $F_2^A(x)/F_2^D(x)$ can be outlined in several finite steps:

1. Get raw yields after all analysis cuts (Chapter 5).
2. Get live proton count (Section 6.1.1).
3. Calculate the kEfficiency and apply it for all (non-BG) targets (Section 6.2.2).
4. Extrapolate kEfficiency for Empty/None target. (Section 6.2.3)
5. Calculate $C(I_c)$ and B_T empty target correction fit and parameters (Section 6.2.3).

6. Calculate weighted yields:

$$Y'_T = \sum_T \frac{1 - B_T \cdot C(I_c)}{K_T(I_c)} \quad (6.39)$$

7. Adjust the deuterium yields for the hydrogen contamination (Equation 6.35).

8. Normalize each targets' yields for each roadset according to the number of live protons (P_A) on target (Section 6.1.1).

9. Take ratios of these adjusted, normalized yields:

$$\frac{Y'_A/P_A}{Y'_D/P_D} \quad (6.40)$$

10. Apply the remaining components of the target-to-target normalization factor to the ratio (Section 6.1):

$$N_{A/D}^{rem.} = \left(\frac{2 \cdot n_D \cdot L^D \cdot T^D(\xi)}{A \cdot n_A \cdot L^A \cdot T^A(\xi)} \right) \cdot \frac{\varepsilon^D}{\varepsilon^A} \cdot \frac{\bar{\Omega}^D}{\bar{\Omega}^A} \quad (6.41)$$

11. With asymmetric nuclei ($^{56}Fe, ^{184}W$), apply $R_{ISO}(x)$ isoscalar correction factor using the mean x_2 value in that ratio bin.

This procedure is general enough that it can be performed with any kinematic binning.

6.6 Corrected, Aggregated Yields

After all corrections are implemented, presented here are the per-live proton yields for the three analyzed roadsets, 57, 62, and 67, binned in x_2 (Tables 6.8, 6.9, and 6.10). The yields here are for the case of kEfficiency corrections in the x_2 kinematic. With this data in hand, by the above steps, we can arrive at our cross section ratio measurements. Also shown here (Table 6.11) is the effect of applying the target contamination correction for the nominal value of the contamination for each roadset.

Before we do so, it is important to ensure that these results are able to be combined. This can be confirmed by investigating the *time-dependence* of the distributions. In this case, this would be comparing the kinematic distributions for different roadsets. This comparison is made in Figures 6.15, 6.16, 6.17, and 6.18 where the normalized distributions for x_2, x_1, M_{γ^*} , and x_F for different pairs of roadsets are shown against each other. The difference between distributions are overlaid on each plot, with the red line showing the location of zero with respect to the difference. Overall, the distributions possess the same shape across all roadsets and the $\chi^2/p.d.f.$ values are acceptably low.

Roadset 57: All Yields Normalized to Live Protons

Target	x_2	Raw Yields	kEff-Corrected Yields	kEff+BG Corrected Yields
LH2	(0.1, 0.13]	60±4	89±6	69±4
	(0.13, 0.16]	164±6	240±10	188±7
	(0.16, 0.195]	193±7	281±10	221±8
	(0.195, 0.24]	147±6	225±10	176±7
	(0.24, 0.29]	79±4	132±8	102±6
	(0.29, 0.35]	42.5±3.2	71±6	56±4
	(0.35, 0.45]	14.4±1.9	24.9±3.5	19.1±2.6
	(0.45, 0.58]	2.0±0.7	3.8±1.4	2.8±1.0
LD2	(0.1, 0.13]	137±8	194±12	186±12
	(0.13, 0.16]	370±14	545±20	522±20
	(0.16, 0.195]	430±15	611±21	589±21
	(0.195, 0.24]	329±13	503±20	482±20
	(0.24, 0.29]	173±9	282±16	271±16
	(0.29, 0.35]	91±7	161±13	154±12
	(0.35, 0.45]	43±5	74±8	71±8
	(0.45, 0.58]	7.4±1.9	12.7±3.5	12.4±3.4
C	(0.1, 0.13]	111±9	161±14	146±12
	(0.13, 0.16]	277±15	429±23	392±21
	(0.16, 0.195]	340±16	500±25	458±22
	(0.195, 0.24]	252±14	391±22	357±20
	(0.24, 0.29]	129±10	224±18	203±16
	(0.29, 0.35]	65±7	111±13	102±12
	(0.35, 0.45]	33±5	53±8	49±8
	(0.45, 0.58]	1.5±1.1	2.1±1.5	2.0±1.4
Fe	(0.1, 0.13]	271±26	(3.9±0.4)e+02	(3.8±0.4)e+02
	(0.13, 0.16]	(6.7±0.4)e+02	(1.03±0.06)e+03	(9.9±0.6)e+02
	(0.16, 0.195]	(8.4±0.5)e+02	(1.31±0.07)e+03	(1.26±0.07)e+03
	(0.195, 0.24]	(6.4±0.4)e+02	(9.9±0.6)e+02	(9.6±0.6)e+02
	(0.24, 0.29]	304±27	(5.0±0.5)e+02	(4.8±0.5)e+02
	(0.29, 0.35]	219±23	(3.7±0.4)e+02	(3.6±0.4)e+02
	(0.35, 0.45]	78±14	135±25	130±24
	(0.45, 0.58]	12±5	26±12	25±12
W	(0.1, 0.13]	267±25	(4.1±0.4)e+02	(4.0±0.4)e+02
	(0.13, 0.16]	(7.3±0.4)e+02	(1.12±0.07)e+03	(1.09±0.06)e+03
	(0.16, 0.195]	(1.00±0.05)e+03	(1.53±0.08)e+03	(1.48±0.07)e+03
	(0.195, 0.24]	(6.9±0.4)e+02	(1.12±0.07)e+03	(1.09±0.07)e+03
	(0.24, 0.29]	368±30	(6.1±0.5)e+02	(5.9±0.5)e+02
	(0.29, 0.35]	230±24	(4.1±0.4)e+02	(3.9±0.4)e+02
	(0.35, 0.45]	76±14	142±27	137±26
	(0.45, 0.58]	12±6	22±10	21±10

Table 6.8: Total yields for each target for each x_2 bin normalized to the live protons on target. The LD2 has also been corrected for contamination in the last column.

Roadset 62: All Yields Normalized to Live Protons

Target	x_2	Raw Yields	kEff-Corrected Yields	kEff+BG Corrected Yields
LH2	(0.1, 0.13]	70±4	102±5	83±4
	(0.13, 0.16]	184±6	272±9	218±7
	(0.16, 0.195]	229±6	341±10	274±8
	(0.195, 0.24]	182±6	281±9	226±7
	(0.24, 0.29]	104±4	174±8	139±6
	(0.29, 0.35]	48.7±3.0	86±6	69±4
	(0.35, 0.45]	19.5±1.9	33.2±3.4	26.7±2.7
	(0.45, 0.58]	3.3±0.8	5.5±1.4	4.5±1.1
LD2	(0.1, 0.13]	157±8	230±12	221±12
	(0.13, 0.16]	413±13	612±20	590±19
	(0.16, 0.195]	516±15	763±22	737±22
	(0.195, 0.24]	389±13	603±20	580±20
	(0.24, 0.29]	193±9	328±16	313±16
	(0.29, 0.35]	110±7	204±13	195±13
	(0.35, 0.45]	49±4	86±8	83±8
	(0.45, 0.58]	5.7±1.5	10.0±2.7	9.6±2.8
C	(0.1, 0.13]	100±10	146±14	134±13
	(0.13, 0.16]	322±17	511±28	470±26
	(0.16, 0.195]	433±20	653±31	603±28
	(0.195, 0.24]	345±18	539±29	497±26
	(0.24, 0.29]	196±13	316±22	293±21
	(0.29, 0.35]	92±9	159±17	147±15
	(0.35, 0.45]	38±6	66±11	61±10
	(0.45, 0.58]	7.4±2.6	11±4	11±4
Fe	(0.1, 0.13]	273±23	415±35	399±34
	(0.13, 0.16]	(7.5±0.4)e+02	(1.15±0.06)e+03	(1.11±0.06)e+03
	(0.16, 0.195]	(9.6±0.4)e+02	(1.53±0.07)e+03	(1.48±0.07)e+03
	(0.195, 0.24]	(7.8±0.4)e+02	(1.22±0.06)e+03	(1.17±0.06)e+03
	(0.24, 0.29]	402±28	(7.1±0.5)e+02	(6.8±0.5)e+02
	(0.29, 0.35]	179±18	314±34	304±33
	(0.35, 0.45]	78±12	138±23	133±22
	(0.45, 0.58]	15±5	26±9	25±9
W	(0.1, 0.13]	319±25	(5.0±0.4)e+02	(4.8±0.4)e+02
	(0.13, 0.16]	(9.5±0.4)e+02	(1.47±0.07)e+03	(1.42±0.07)e+03
	(0.16, 0.195]	(1.16±0.05)e+03	(1.83±0.08)e+03	(1.78±0.07)e+03
	(0.195, 0.24]	(9.3±0.4)e+02	(1.51±0.07)e+03	(1.47±0.07)e+03
	(0.24, 0.29]	430±29	(7.4±0.5)e+02	(7.2±0.5)e+02
	(0.29, 0.35]	264±23	(4.9±0.4)e+02	(4.7±0.4)e+02
	(0.35, 0.45]	131±16	221±28	215±28
	(0.45, 0.58]	12±5	25±11	24±11

Table 6.9: Total Roadset 62 yields for each target for each x_2 bin normalized to the live protons on target. The LD2 has also been corrected for contamination in the last column.

Roadset 67: All Yields Normalized to Live Protons

Target	x_2	Raw Yields	kEff-Corrected Yields	kEff+BG Corrected Yields
LH2	(0.1, 0.13]	50.2±1.8	75.7±2.7	61.8±2.2
	(0.13, 0.16]	140.5±2.9	211±5	172±4
	(0.16, 0.195]	167.5±3.2	251±5	206±4
	(0.195, 0.24]	136.6±2.9	212±5	175±4
	(0.24, 0.29]	70.9±2.1	120±4	97.8±2.9
	(0.29, 0.35]	36.6±1.5	66.4±2.9	54.1±2.3
	(0.35, 0.45]	15.2±1.0	28.1±1.9	22.6±1.5
	(0.45, 0.58]	3.3±0.5	5.8±0.8	4.8±0.7
LD2	(0.1, 0.13]	131±4	194±6	185±6
	(0.13, 0.16]	346±7	525±10	502±10
	(0.16, 0.195]	427±7	648±12	621±11
	(0.195, 0.24]	325±6	515±11	493±10
	(0.24, 0.29]	186±5	323±9	310±9
	(0.29, 0.35]	88.2±3.4	158±6	152±6
	(0.35, 0.45]	37.6±2.2	67±4	64±4
	(0.45, 0.58]	5.8±0.9	10.1±1.6	9.6±1.5
C	(0.1, 0.13]	101±5	149±8	137±7
	(0.13, 0.16]	281±9	444±15	407±13
	(0.16, 0.195]	326±10	499±15	458±14
	(0.195, 0.24]	242±8	377±13	346±12
	(0.24, 0.29]	144±6	236±11	217±10
	(0.29, 0.35]	62±4	112±8	103±7
	(0.35, 0.45]	36.1±3.2	64±6	59±5
	(0.45, 0.58]	2.8±0.9	5.6±1.9	5.0±1.7
Fe	(0.1, 0.13]	209±11	315±17	304±16
	(0.13, 0.16]	629±19	990±31	955±30
	(0.16, 0.195]	752±21	(1.24±0.04)e+03	1197±34
	(0.195, 0.24]	598±19	960±31	927±30
	(0.24, 0.29]	332±14	561±25	542±24
	(0.29, 0.35]	161±10	293±19	283±18
	(0.35, 0.45]	82±7	146±13	141±13
	(0.45, 0.58]	9.5±2.4	21±6	20±5
W	(0.1, 0.13]	256±12	404±20	392±19
	(0.13, 0.16]	721±20	1132±33	1098±32
	(0.16, 0.195]	946±24	(1.50±0.04)e+03	(1.45±0.04)e+03
	(0.195, 0.24]	742±21	1219±35	1184±34
	(0.24, 0.29]	390±15	673±27	654±26
	(0.29, 0.35]	206±11	391±22	379±21
	(0.35, 0.45]	86±7	155±14	150±13
	(0.45, 0.58]	15.9±3.1	26±5	26±5

Table 6.10: Total Roadset 67 yields for each target for each x_2 bin normalized to the live protons on target. The LD2 has also been corrected for contamination in the last column.

Roadset	x_2	Yields	Contam-Corrected Yields
57	(0.1, 0.13]	175±11	186±12
	(0.13, 0.16]	490±18	522±20
	(0.16, 0.195]	553±19	589±21
	(0.195, 0.24]	453±18	482±20
	(0.24, 0.29]	255±14	271±16
	(0.29, 0.35]	144±11	154±12
	(0.35, 0.45]	66±7	71±8
	(0.45, 0.58]	11.4±3.1	12.4±3.4
	62	(0.1, 0.13]	209±11
(0.13, 0.16]		558±18	590±19
(0.16, 0.195]		696±20	737±22
(0.195, 0.24]		549±18	580±20
(0.24, 0.29]		297±14	313±16
(0.29, 0.35]		184±12	195±13
(0.35, 0.45]		78±7	83±8
(0.45, 0.58]		9.2±2.5	9.6±2.8
67		(0.1, 0.13]	180±6
	(0.13, 0.16]	487±10	502±10
	(0.16, 0.195]	602±11	621±11
	(0.195, 0.24]	478±10	493±10
	(0.24, 0.29]	300±8.2	310±9
	(0.29, 0.35]	147±5.9	152±6
	(0.35, 0.45]	62±3.8	64±4
	(0.45, 0.58]	9.4±1.5	9.6±1.5

Table 6.11: Adjustment made to deuterium yields (per live proton) to correct for deuterium contamination. Shown are the kEff+BG corrected yields before and after the contamination correction.

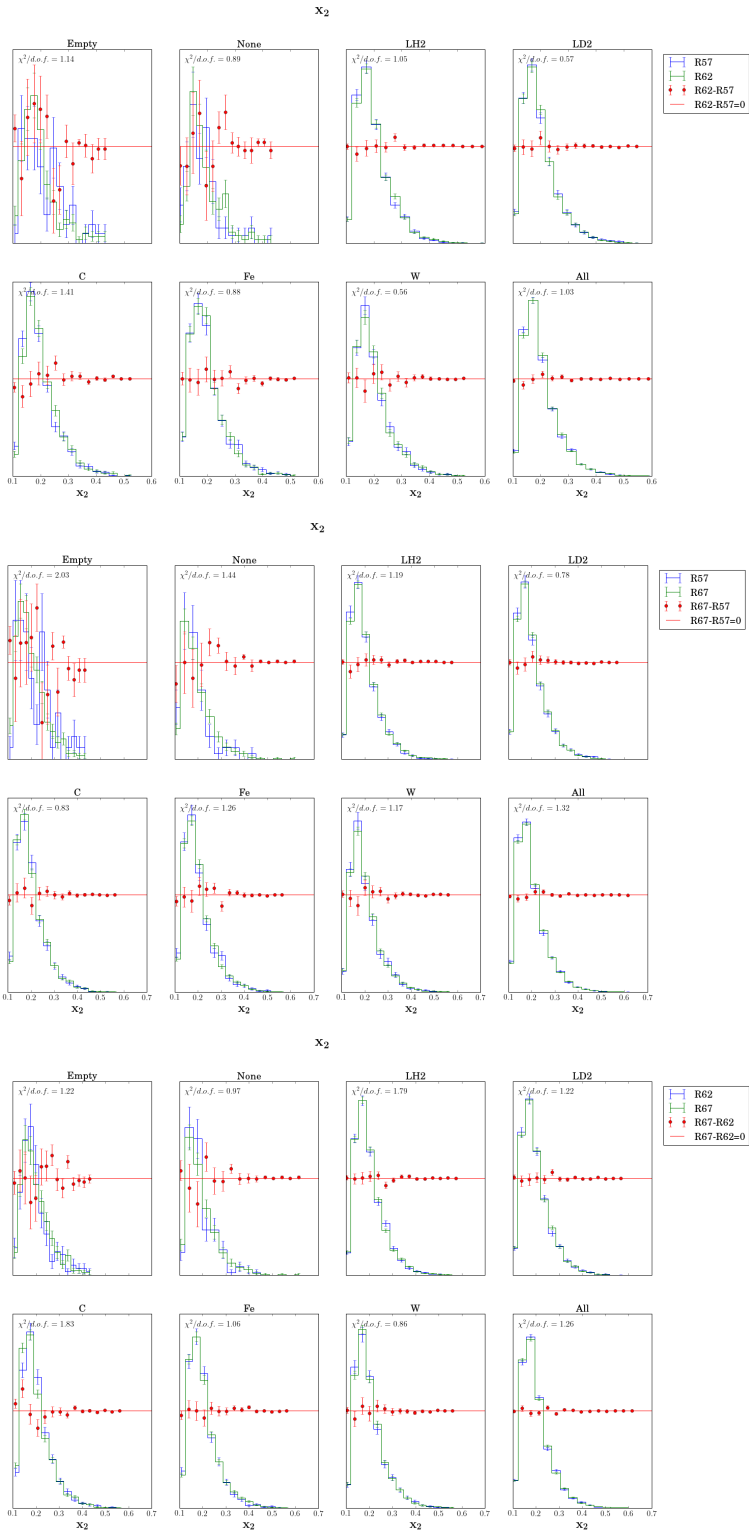


Figure 6.15: A comparison of x_2 distributions between roadsets for each target, and all targets combined. The difference between the two normalized distributions shown by the red data points, with a red line at zero. The χ^2 shown is for the comparison between the roadset differences and the line at zero.

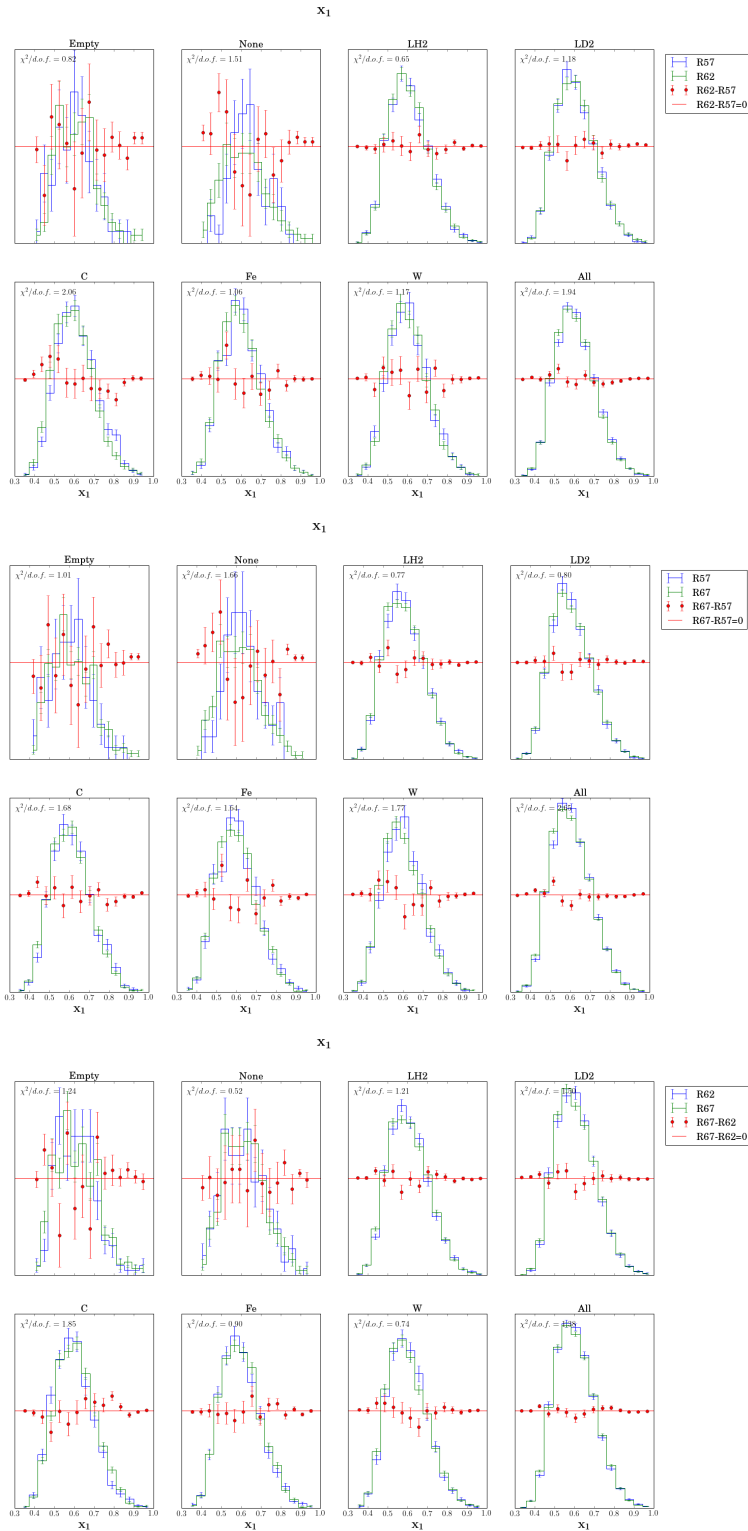


Figure 6.16: A comparison of x_1 distributions between roadsets for each target, and all targets combined. The difference between the two normalized distributions shown by the red data points, with a red line at zero. The χ^2 shown is for the comparison between the roadset differences and the line at zero.

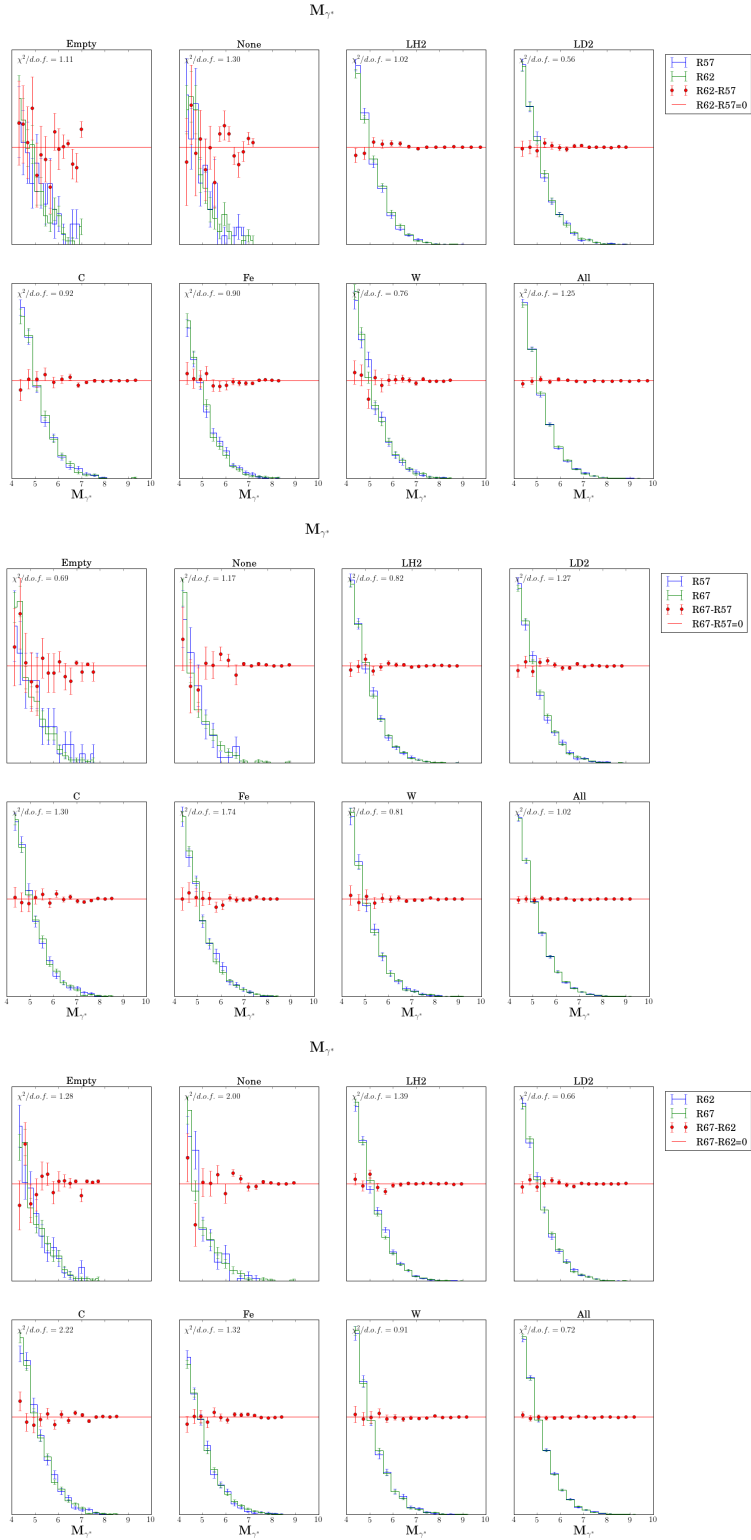


Figure 6.17: A comparison of M_{γ^*} distributions between roadsets for each target, and all targets combined. The difference between the two normalized distributions shown by the red data points, with a red line at zero. The χ^2 shown is for the comparison between the roadset differences and the line at zero.

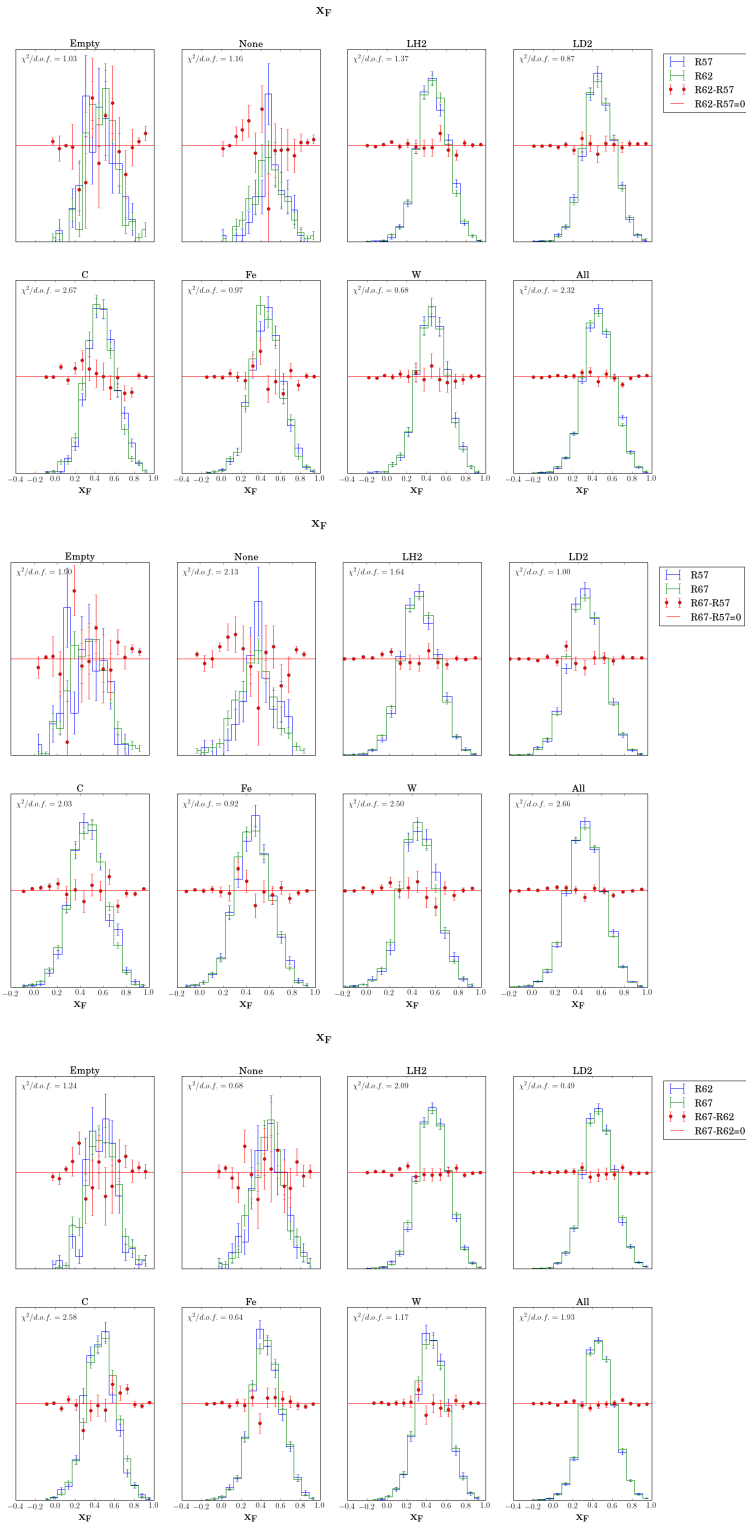


Figure 6.18: A comparison of x_F distributions between roadsets for each target, and all targets combined. The difference between the two normalized distributions shown by the red points, with a red line at zero. The χ^2 shown is for the comparison between the roadset differences and the line at zero.

Chapter 7

Results and Interpretation

Through the extensive procedures outlined in the previous chapter, a set of final results can be provided. Here, the event distributions versus M_{γ^*} , x_F , p_T , x_1 , and $\cos\theta$ are presented. Finally, the per-nucleon cross section ratios are shown for σ^C/σ^D , σ^{Fe}/σ^D , and σ^W/σ^D , along with a comparison of these ratios to previous results and some interpretation of the findings.

7.1 Drell-Yan Event Distributions

The following kinematic distributions represent the acceptance of the SeaQuest dimuon spectrometer. By measuring the momenta (p_μ) of the muon tracks at the vertex position, the quantities of invariant mass (M_{γ^*}), x_F , and p_T of the virtual photon can be calculated.

$$p_\mu = (p_x, p_y, p_z) \quad (7.1)$$

$$E_{\mu^+\mu^-} = \sqrt{|p_{\mu^+}|^2 + |p_{\mu^-}|^2} \quad (7.2)$$

$$P_{\mu^+\mu^-} = p_{\mu^+} + p_{\mu^-} \quad (7.3)$$

$$p_T = \sqrt{(p_{\mu^+}^x + p_{\mu^-}^x)^2 + (p_{\mu^+}^y + p_{\mu^-}^y)^2} \quad (7.4)$$

$$p_z = p_{\mu^+}^z + p_{\mu^-}^z \quad (7.5)$$

$$M_{\gamma^*} = \sqrt{E_{\mu^+\mu^-}^2 - |P_{\mu^+\mu^-}|^2} \quad (7.6)$$

$$x_F = \frac{p_{\mu^+}^z + p_{\mu^-}^z}{\sqrt{s}/2(1 - M_{\gamma^*}^2/s)} \approx 2(p_{\mu^+}^z + p_{\mu^-}^z)/\sqrt{s} \quad (7.7)$$

$$(7.8)$$

From these, x_1 and x_2 can be derived according to Eqs. 1.13.

The distributions of the dimuon kinematic quantities x_F , M_{γ^*} , p_T , and p_z can be found in Figure 7.1 while the quark quantities x_1 , x_2 , and $\cos\theta_\mu$ can be found in Figure 7.2. As can be seen from the distributions, the SeaQuest experiment studies muons with an invariant mass between 4.2 and 10 GeV (the range between the




Figure 7.1: Event distributions for all combined data versus dimuon kinematic variables x_F , M_{γ^*} , p_T , and p_z




Figure 7.2: Event distributions for all combined data versus quark kinematic variables x_1, x_2 , and the cosine of the polar decay angle, θ_μ

Mean Kinematics

x_2 -bin	$\langle x_1 \rangle$	$\langle x_2 \rangle$	$\langle M_{\gamma^*} \rangle$ (GeV)	$\langle x_F \rangle$	$\langle \cos \theta_\mu \rangle$	$\langle \phi_\mu \rangle$	$\langle p_z \rangle$ (GeV)	$\langle p_T \rangle$ (GeV)	$\langle p_z^+ \rangle$ (GeV)	$\langle p_z^- \rangle$ (GeV)	$\langle p_T^+ \rangle$ (GeV)	$\langle p_T^- \rangle$ (GeV)
(0.1, 0.13]	0.746	0.120	4.378	0.690	0.011	0.043	90.145	0.737	45.815	44.329	2.124	2.112
(0.13, 0.16]	0.650	0.146	4.516	0.559	-0.005	-0.030	78.487	0.762	39.415	39.079	2.224	2.186
(0.16, 0.195]	0.592	0.177	4.751	0.467	-0.011	-0.008	71.473	0.805	35.833	35.641	2.365	2.305
(0.195, 0.24]	0.567	0.215	5.143	0.404	-0.009	0.008	68.394	0.841	34.344	34.015	2.562	2.500
(0.24, 0.29]	0.557	0.262	5.626	0.348	-0.007	0.008	67.097	0.900	33.728	33.349	2.800	2.726
(0.29, 0.35]	0.543	0.314	6.079	0.280	-0.021	0.038	65.423	1.008	32.313	33.088	3.013	2.941
(0.35, 0.45]	0.535	0.386	6.713	0.191	-0.005	0.024	64.301	0.994	32.014	32.282	3.319	3.304
(0.45, 0.58]	0.512	0.485	7.331	0.040	-0.019	-0.220	61.508	1.086	31.295	30.323	3.715	3.626
Combined:	0.601	0.201	5.012	0.453	-0.002	0.000	72.553	0.865	36.444	36.109	2.472	2.441

Table 7.1: The mean dimuon kinematics in each bin of x_2 . These should be considered to be approx/imate, as they are calculated as a weighted average from the data set when corrected for rate dependence in x_2 .

ψ' and the Υ), and the acceptance measures mostly forward-moving ($x_F > 0$) muon pairs with longitudinal momentum (p_z) greater than 40 GeV in the lab frame. The acceptance of p_T of the muon pair is limited to 3 GeV. It is worth noting that, as can be seen from the $\cos\theta_\mu$ distribution, there is certainly contamination evident at the tails of the distribution, primarily in the positive $\cos\theta$ region. This is one of the key clues that SeaQuest has a problem with *combinatoric* background, or reconstructed muon pairs that arise from uncorrelated muons produced in the target or beam dump.

The beam quarks participating in the Drell-Yan process can have a fractional momentum x_1 ranging from around 0.35 to 0.9 with the distribution peaking at 0.57 and having a mean value of $\langle x_1 \rangle = 0.6$. The target antiquark participating in the interaction can possess a fractional momentum x_2 ranging from 0.1 to 0.5, with the distribution peaking at around 0.16 and having a mean value of $\langle x_2 \rangle = 0.2$. A table of the total mean kinematics and also the mean kinematics for each x_2 bin can be found in Table 7.1.

7.2 The A-Dependence of the Drell-Yan Process

The ratio for $R_{DY}^{A/D}$ are presented here, where the ratio is defined Eq. 6.1 and 6.2. Results are presented for the nuclear-to-deuterium ratios for carbon, iron, and tungsten. The results are presented in the following order:

- The ratio versus dimuon kinematics M_{γ^*} , x_F , and p_T
- The ratio versus quark kinematics x_1 , x_2 , and $\cos\theta_\mu$
- The ratio versus x_2 for the three roadsets analyzed
- The ratio versus x_2 for two bins of p_T , two bins of x_1 , and two bins of M_{γ^*} (two Q^2 intervals)
- The integrated ratio versus mass number A .

But before results are presented, a total systematic uncertainty will be estimated to consider when evaluating the results.

7.2.1 Systematic Uncertainties

The types of systematic uncertainties considered here are the following:

1. Liquid deuterium contamination
2. kEfficiency curve fit uncertainty

x_2 -bin	C/D-	C/D+	Fe/D-	Fe/D+	W/D-	W/D+
(0.1, 0.13]	-0.010	0.011	-0.018	0.019	-0.024	0.025
(0.13, 0.16]	-0.016	0.016	-0.011	0.011	-0.023	0.023
(0.16, 0.195]	-0.017	0.018	-0.025	0.026	-0.027	0.028
(0.195, 0.24]	-0.020	0.020	-0.017	0.018	-0.024	0.025
(0.24, 0.29]	-0.013	0.014	-0.017	0.018	-0.018	0.018
(0.29, 0.35]	-0.015	0.015	-0.018	0.018	-0.023	0.023
(0.35, 0.45]	-0.009	0.009	-0.022	0.023	-0.009	0.009
(0.45, 0.58]	-0.005	0.004	-0.016	0.016	-0.012	0.012

Table 7.2: Systematic error arising from the uncertainty in the kEfficiency exponential fit parameter.

3. Empty target curve fit uncertainty
4. Empty target normalization constant uncertainty
5. Remaining rate dependence uncertainty.

These factors were each independently varied from their nominal values to $\pm 1\sigma$, and the final ratio measurement was evaluated for each of the three settings. The percentage change between of the final result is what is quoted as the systematic uncertainty for that particular factor. With each uncertainty evaluated, the upper and lower bounds of asymmetric uncertainties were combined in quadrature.

In dealing with the asymmetric uncertainty in the deuterium contamination, it was found to have a uniform effect across targets and binning of $+2\%/ - 1.4\%$. The kEfficiency exponential fit uncertainty has a target dependence since the fit parameter was evaluated for different nuclear targets. These uncertainties are shown in Table 7.2, and in general are around $\pm 2\%$. The Empty target curve fit parameter uncertainty had little effect and led to $\pm 0.1 - 0.7\%$ variation in the ratios. The uncertainty in the normalization constant, B, of the empty target background subtraction also propagated a small variation of $\pm 0.1 - 0.7\%$. The remaining rate dependence was previously evaluated to be 1.07% when it was defined.

Without any approximation, the calculated asymmetric systematic uncertainties were combined in quadrature (independently for the upper and lower intervals) to yield the total estimation, as seen in Table 7.3. To a good approximation, it can be stated that the systematic uncertainty of a nuclear cross section ratio measurement at SeaQuest due to the above-described sources can, to a good approximation, be considered to be $\pm 3\%$, with it being slightly higher in the case of tungsten and slightly lower in the case of carbon.

7.2.2 M_{γ^*}, x_F , and p_T Dependence

Shown in Figure 7.3 is the nuclear cross section ratio as a function of invariant mass for the three heavy nuclei. While dimuons are accepted up to 10 GeV , there is little-to-no signal that passes all cuts above

Fractional Systematic Uncertainties

x_2 -bin	C/D-	C/D+	Fe/D-	Fe/D+	W/D-	W/D+
(0.1, 0.13]	0.021	0.025	0.028	0.030	0.032	0.035
(0.13, 0.16]	0.023	0.028	0.023	0.027	0.031	0.034
(0.16, 0.195]	0.025	0.029	0.032	0.035	0.034	0.037
(0.195, 0.24]	0.026	0.030	0.027	0.030	0.032	0.035
(0.24, 0.29]	0.022	0.026	0.027	0.030	0.028	0.031
(0.29, 0.35]	0.023	0.028	0.029	0.030	0.031	0.034
(0.35, 0.45]	0.019	0.026	0.030	0.034	0.022	0.027
(0.45, 0.58]	0.022	0.038	0.024	0.029	0.022	0.030

Table 7.3: The total estimated *fractional* systematic uncertainty for the ratio measurements versus x_2 .

7.5 GeV. Figure 7.4 shows the ratio's dependence on x_F , where an enhancement can be seen below $x_F < 0.5$ and a depletion in the ratio value can be seen above $x_F > 0.5$. Figure 7.5 shows the ratio as a function of p_T , where there is a trend of a depletion in the ratio for low $p_T < 0.5$ GeV and an excess beyond that point. Lastly, there is the ratio versus the longitudinal momentum, p_z , in Figure 7.6, where no significant trend in behavior is observed.

7.2.3 x_1, x_2 , and $\cos \theta_\mu$ Dependence

The quark variables x_1 and x_2 , which are the momentum fractions of the beam and target hadrons carried by the interacting quarks, respectively. Figure 7.7 shows the values of $R_{DY}(A/D)$ as a function of x_1 for each target. Figure 7.9 shows the values of the ratio $R_{DY}(A/D)$ as a function of x_2 , which is the main focus of this analysis. This plot is zoomed out to show an anomalously low point for the carbon target at $(x_2, R_{DY}) = (0.48, 0.39)$. The same figure zoomed in to match the rest of the plots can be found in Figure 7.10. Across the three targets, a depletion is observed at $x_2 < 0.15$. Elsewhere, there appears to be a small excess in the $0.15 < x_2 < 0.25$ for iron and tungsten. This depletion appears to be approximately an 8% depletion for all targets at $x_2 = 0.12$. A tabulation of the data from this plot can be found in Table 7.4, including estimated asymmetric systematic uncertainties.

This data for this is also broken up into the three roadsets in Figure 7.11 and into two bins of p_T in Figure 7.12. It is clear from the roadset comparison that there is considerable disagreement at higher $x_2 > 0.25$, but the tungsten data seems to agree up to $x_2 = 0.35$. When binned in high and low bins of p_T , there seems to be a systematic excess in the heavier nuclear component of the ratio for high- p_T events ($p_T > 0.6$ GeV).




Figure 7.3: The measured Drell-Yan per-nucleon cross section ratio against invariant dimuon mass. Only statistical uncertainty is shown.




Figure 7.4: The measured Drell-Yan per-nucleon cross section ratio against dimuon Feynman- x . Only statistical uncertainty is shown.




Figure 7.5: The measured Drell-Yan per-nucleon cross section ratio versus transverse momentum of the muon pair. Only statistical uncertainty is shown.




Figure 7.6: The measured Drell-Yan per-nucleon cross section ratio versus longitudinal momentum of the muon pair. Only statistical uncertainty is shown.




Figure 7.7: The measured Drell-Yan per-nucleon cross section ratio versus fractional momentum quantity, x_1 . Only statistical uncertainty is shown.




Figure 7.8: The measured Drell-Yan per-nucleon cross section ratio versus the cosine of the polar muon production angle, θ_μ . Only statistical uncertainty is shown.




Figure 7.9: The measured Drell-Yan per-nucleon cross section ratio versus fractional momentum quantity, x_2 . Only statistical uncertainty is shown. Overlaid is the data from the E-772 experiment.




Figure 7.10: The same plot, but zoomed into $R=[0.7, 1.2]$ for closer comparison with E-772 data.




Figure 7.11: The measurement of R_{DY} for the three roadsets. Their results are combined to render the values found in Figure 7.9.




Figure 7.12: The measured Drell-Yan per-nucleon cross section ratio versus fractional momentum quantity, x_2 , grouped into "high" p_T and "low" p_T . Only statistical uncertainty is shown.

$\langle x_2 \rangle$	C/D	sys.-	sys.+	Fe/D	sys.-	sys.+	W/D	sys.-	sys.+
0.12	0.91±0.04	0.019	0.022	0.93±0.05	0.026	0.028	0.90±0.04	0.029	0.032
0.15	1.038±0.030	0.024	0.029	1.011±0.029	0.024	0.027	0.964±0.026	0.030	0.033
0.18	1.010±0.027	0.025	0.029	1.032±0.027	0.034	0.036	1.035±0.025	0.035	0.038
0.22	0.962±0.030	0.025	0.029	1.049±0.031	0.028	0.031	1.048±0.029	0.034	0.037
0.26	0.98±0.04	0.022	0.026	0.97±0.04	0.026	0.029	0.94±0.04	0.026	0.029
0.31	0.89±0.05	0.021	0.025	0.98±0.06	0.029	0.030	1.07±0.06	0.034	0.037
0.39	1.07±0.09	0.020	0.028	1.11±0.10	0.033	0.037	1.01±0.08	0.022	0.027
0.49	0.36±0.13	0.008	0.014	1.24±0.30	0.030	0.036	1.08±0.25	0.024	0.032

Table 7.4: The tabulated data of $R_{DY}(A/D)$ for C/D, Fe/D, and W/D, fully with statistical uncertainty and upper and lower asymmetric systematic uncertainties.




Figure 7.13: Event distributions for all combined data versus quark kinematic variables x_1, x_2 , and the cosine of the polar decay angle, θ_μ

7.2.4 Integrated Ratio versus A

Integrating over any existing kinematic dependence, Figure 7.13 shows the ratio of total yields between σ^C/σ^D , σ^{Fe}/σ^D , and σ^W/σ^D plotted against the atomic weight A . Weights 12.0107, 55.845, and 183.84 were used for carbon, iron, and tungsten respectively. Each data point has approximately a 1.7% statistical error, with the $\sim 3\%$ systematic uncertainty now shown. The ratio of integrated carbon yields to deuterium yields appears to be consistent with unity while iron and tungsten show a very slight enhancement.

7.3 Discussion of $R_{DY}(x_2)$

As discussed in Section 1.4, a particularly interesting nuclear cross section ratio is that of DY versus x_2 . To good approximation, the following should hold:

$$R^{DY}(x_2) \approx \frac{\bar{u}^A(x_2)}{\bar{u}^d(x_2)}. \quad (7.9)$$

Looking at the results in this context, it would appear that when also considering the additional 3% systematic uncertainty, there seems to be little modification of the sea quark distribution due to nuclear effects.

Here, some topics regarding this measurement are discussed, including a criticism and alternate proposal for calculating ratios and combining separate measurements. This is followed by a comparison with the only other comparable result that exists, E-772. Finally, some of the models that make some predictions regarding the DY nuclear cross section ratio are briefly addressed and discussed.

7.3.1 Natural Log of Ratios

It has been recently discussed within the SeaQuest collaboration that perhaps the usual way of measuring, calculating, and visualizing ratios has been historically flawed – or more flawed than need be. To begin this explanation, consider the simple case of plotting a non-negative ratio of two numbers A and B. In Figure 7.14, one can see the available phase space




Figure 7.14: A depiction of the asymmetric phase space for A/B versus $\ln(A/B)$.

for such a value in the case of a direct ratio of A/B and the available phase space for the natural log of the ratio. The natural log representation shows a more desirably symmetric available phase space.

Further, consider the values of $A=200$ and $B=150$. If one wished to establish the relation of these two quantities, one might find it desirable that the plotting of A/B yields an opposite yet symmetric result with respect to B/A . Instead, what is seen is a *visually* (but not mathematically) deceiving asymmetric representation. One way shows $A/B=1.3\bar{3}$, and the other shows $B/A=0.75$. If one were to attempt to characterize an asymmetry in the quantities based upon the *magnitude of the deviation from unity*, then the impression of the magnitude of the asymmetry would change with respect to which quantity one chooses as




Figure 7.15: The combined results of the log of R_{DY} from the three roadsets.

the numerator and which one chooses for the denominator.

Another hint that a change in representation may be in order is the existence of nonsense error bars for low ratio value measurements. It is possible, though it should not be, in some cases with low statistics, that an error calculated for A/B to measure its nominal value R and uncertainty σ to have a lower bound of $R - \sigma < 0$. For a non-negative ratio, this makes no sense, and as such indicates that perhaps this should be handled differently.

The reason for the above is primarily due to uncertainties only being calculated to first order. By this, I refer to the case that the uncertainty of a function $f(x)$ can be expressed as a Taylor expansion:

$$\Delta f(x) = \sum_i^{\infty} \frac{d^i f(x)}{dx^i} (\Delta x)^i. \quad (7.10)$$

In the case of the ratio A/B , this contributes *one* term to the ratio uncertainty when considering the variation in A , but infinite terms when considering the variation in B (which are likely completely ignored). When considering the natural log of the ratio, this uncertainty is redistributed in such a way that, even if the series is only calculated to first order, it is done equally to both numerator and denominator.

Let us now look at a re-evaluation of the ratio represented in Figure 7.9. Please make special note of anomalously low value in the highest x_2 bin of σ^C/σ^D and how it appears to be $\sim 4\sigma$ away from unity. Now, we re-evaluate the ratio by calculating

$$L_{DY} = \ln(R_{DY} = \ln(Y_A/Y_D) = \ln(Y_A) - \ln(Y_D)), \quad (7.11)$$

where Y_A is the normalized yield for target A , propagating first order uncertainties along the way. Once

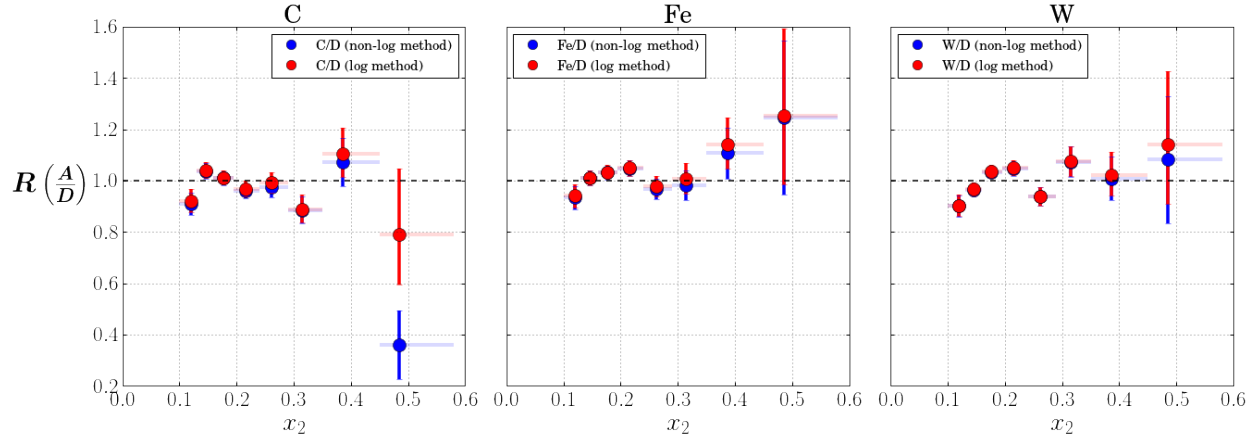


Figure 7.16: The combined results of the log of R_{DY} transformed back into a plot of R_{DY} by mapping the data points and error bounds with the exponential function.

there, if it is so desired, one can translate directly back from the natural log ratio evaluation to the bare ratio evaluation by taking the exponential of each data point and *the absolute location of the 1σ uncertainty boundaries*, not the magnitude of the uncertainties.

In Figure 7.15 shows what the combined results look like if the log ratios are calculated. It is difficult to have an intuition for the values of $\ln(R_{DY})$, but here the statistical error bars avoid the possibility of crossing a forbidden boundary. Now, what happens when all the nominal points and upper and lower statistical boundaries are mapped back into R_{DY} values with an exponential function.

Here in Figure 7.16, we see something unexpected – the anomalously low point for carbon at high x_2 that was observed in the original Figure 7.9 is now bumped up significantly from $R_{DY} = 0.36$ up to $R_{DY} = 0.8$. The tabulated data from this can be found in Table 7.5. *How could this simple transformation have such a drastic effect?*

The answer is bipartite: it lies in (1) the poor estimate in uncertainty of $R_{DY} = Y_A/Y_D$ when the ratio is taken directly, and (2) the way in which different measurements are combined. When a cross section ratio measurement is made for each of the three roadsets (R_i), the ratio values are combined (R_C) via the

	C/D-	C/D+	Fe/D-	Fe/D+	W/D-	W/D+
(0.1, 0.13]	0.878	0.967	0.897	0.987	0.862	0.944
(0.13, 0.16]	1.011	1.072	0.982	1.041	0.941	0.994
(0.16, 0.195]	0.985	1.039	1.006	1.061	1.010	1.061
(0.195, 0.24]	0.940	0.999	1.019	1.081	1.021	1.080
(0.24, 0.29]	0.952	1.033	0.938	1.018	0.903	0.977
(0.29, 0.35]	0.837	0.944	0.951	1.068	1.018	1.135
(0.35, 0.45]	1.013	1.206	1.047	1.248	0.941	1.112
(0.45, 0.58]	0.597	1.047	0.986	1.594	0.911	1.427

Table 7.5: The ratio values with asymmetric *statistical* uncertainties after combining the results using the natural log transform method.

following:

$$R_1 = Y_1^A/Y_1^D, \quad R_2 = Y_2^A/Y_2^D, \dots \quad (7.12)$$

$$R_C = \frac{\sum_i R_i/\sigma_i^2}{\sum_i 1/\sigma_i^2} \quad (7.13)$$

$$\delta R_C = \frac{1}{\sqrt{\sum_i 1/\sigma_i^2}} \quad (7.14)$$

The problem arises that when the combined measurement is weighted by the uncertainties, and the uncertainties are poorly defined, then this can really have a substantial effect on where the actual data point ends up.

By transforming to the $R_i \rightarrow L_i = \ln(R_i)$, more appropriate uncertainties can be approximated and used to combine many measurements of L_i into a combined L_C by the method described just above, and then those combined data points and absolute uncertainties are translated back to R_{DY} .

At the time of this writing, SeaQuest is actively attempting to re-evaluate the previous results of the ratio measurements made by E-866 to investigate the full extent to which this representation and uncertainty calculation may influence their results. In the E-866 σ^D/σ^H measurement, there is a *very* similar situation of three data sets being combined, resulting in the highest x -bin showing an anomalously depleted value. No model has been fully able to explain this behavior.




Figure 7.17: The measured Drell-Yan per-nucleon cross section ratio versus fractional momentum quantity, x_2 . Overlaid is the data from the E-772 experiment. Only statistical uncertainty is shown. There exists an additional $\sim 3\%$ systematic uncertainty.

7.3.2 Comparison to E-772

While there does exist some data from E-866 regarding nuclear cross section ratios for Drell-Yan, they have only reported the ratios for σ^{Fe}/σ^{Be} and σ^W/σ^{Be} 1.16. This leaves E-772's results as the only data set to compare the SeaQuest results to. The comparison of the two results can be made above in Figure 7.14.

In E-772, they observed an A-dependent depletion of the ratio at very low $x_2 < 0.1$, where shadowing would be expected (discussed in Section 1.3.6). Theoretical calculations indicated that the shadowing seen in DIS and EMC ratio studies should also be present in Drell-Yan interactions at low-x [86]. The surprising observation with the SeaQuest data is, if it persists through new data and corrections, there appears to be a depletion similar to the shadowing observed, but the behavior is observed above $x_2 > 0.1$. The behavior is not significant, however, and since each of the three measurements make use of the same denominator, an anomalous excess in deuterium yields at low-x would cause this behavior in all targets. This is a feature worth monitoring as the analysis progresses.

Regarding the rest of the x_2 range, E-772 measured ratio results consistent with unity, fitting a flat line to the data at $R_{DY}(x_2) = 0.992 \pm 0.007$ [82]. With the SeaQuest data, a similar measurement can be taken yielding $R_{DY}^C(x_2) = 0.96 + / - 0.03$, $R_{DY}^{Fe}(x_2) = 1.03 + / - 0.03$, and $R_{DY}^W(x_2) = 1.00 + / - 0.03$, which is in agreement with E-772's constant fit.

7.3.3 Models For Nuclear Modifications

The A-dependent modification of the nucleon has become a challenging and intriguing phenomenon. So many models have arisen attempting to reproduce the data, it has led some of those who study it to give it the




Figure 7.18: Virtual meson cloud model predictions for DY ratios for σ^{Fe}/σ^D [87]. Models shown are tuned for 800 GeV experiments, and may not apply exactly SeaQuest beam energies.

moniker “*Every Model is Cool*” [88]. Drell-Yan’s contributions in isolating particular sea quark components of models’ predictions can provide some discriminatory powers between models where DIS data may be satisfied by most models.

One particular model is the pion excess model. A proton is not simply three quarks, and a nucleus is not simply a convolution of nucleons. The nucleons themselves are bound by the exchange of mesons at intermediate and long range. These mesons carry a certain fraction of the momentum of the nucleons, and a modification of the meson exchange field in a nuclear medium was predicted by Berger and Coester to lead to an enhancement in the DIS and DY cross section ratios [52]. In this model, there is an enhancement in the valence distributions at high x , leading to a depletion of valence quarks at intermediate x , causing the EMC effect observed.

Seeing as pions carry a valence anti-quark, it was predicted that these intermediary pions would provide an enhancement to the DY cross section ratio for nuclear targets. In Figure 1.17, E-772’s DY nuclear cross section ratio results showed strong disagreement with this virtual meson exchange model.

It may, however, be the case that the feature is masked by another effect or the predicted scale is slightly off. In either case, if the pion excess picture is valid, the behavior at higher x_2 will still hold and become evident in the SeaQuest data (see predictions by Coester and Jung and Miller for iron in Fig. 7.18). More data at higher x_2 will be required to observe if the spike seen in the SeaQuest data for iron and tungsten persist. If so, this could lend support for the Coester view of the pion excess model as, at least in part, contributing to the nuclear modification.

A composite model regarding several nuclear effects has been assembled by Kulagin and Petti. Such

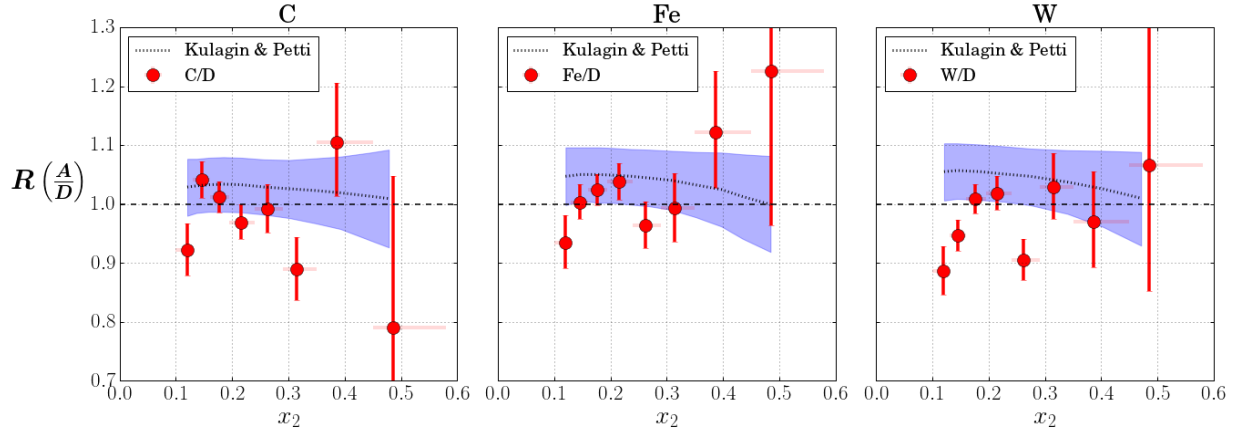


Figure 7.19: The SeaQuest results and a composite model prediction from Kulagin and Petti for DY ratios σ^C/σ^D , σ^{Fe}/σ^D , and σ^W/σ^D [89, 90].

nuclear effects are the pion excess (PI), Fermi motion (FM), off-shell effects (OS), and nuclear coherence (NS, or nuclear shadowing) [89]. They have been in communication with the SeaQuest collaboration and have provided a set of predictions [90] based upon the particular kinematics of the SeaQuest beam and dimuon acceptance. There is still some uncertainty to these predictions based on the amount of parton energy loss as the beam quark moves through the nuclear medium. This has been extracted from the E-772 and E-866 data but is still unknown for the beam energy used at the SeaQuest experiment. It is important to note that the predictions are for the *non-isoscalar corrected* data. This data and the predictions can be found in Figure 7.19.

In general, the SeaQuest data does agree with the predictions of Kulagin and Petti, but the statistics of the data over the full x_2 range are such that the wide prediction band lies within one or two standard deviations of most data points. As such, all that can be said at the moment is that the SeaQuest data does not disagree with this composite model prediction. It is worthy to note that the ratio shows a gradual enhancement with increasing A, whereas the SeaQuest data suggests a minor overall depletion from iron to tungsten. Again, this is the data before any isoscalar correction is applied.

Chapter 8

Conclusions

SeaQuest is a fixed-target, 120 GeV proton beam experiment at Fermilab that seeks to analyze the rare Drell-Yan process in pA collisions. With this beam energy and high instantaneous luminosities, SeaQuest has the potential to explore the high- x_2 kinematic range of the Drell-Yan process. This kinematic space has yet to be fully explored with Drell-Yan due to the depletion of anti-quarks at this high of x_2 and the rarity of the Drell-Yan process. In studying Drell-Yan data with its forward spectrometer, SeaQuest is able to isolate and define features of the sea quark distribution in the nucleon.

While a main focus is placed on studying deuterium to hydrogen and extracting $\bar{d}(x)/\bar{u}(x)$, the inclusion of heavy nuclear targets in the target rotation enable a study of Drell-Yan cross section ratios for heavy nuclei to deuterium. It has long been established that, in DIS experiments, quark distributions become modified in a nuclear medium. By studying this nuclear DY cross section ratio, we gain a measurement of $\sim \bar{u}_A(x)/\bar{u}_D(x)$, which lends a direct probe of the modification of the quark sea in the nuclear medium.

Presented in this work is the first dedicated attempt at extracting this ratio, $R_{DY}^{A/D}$ from existing SeaQuest data. Many challenges need be addressed in order to arrive at such a measurement, namely the handling of the rate dependence. Outlined here have been the first procedures and attempts at correcting for the rate dependence in such a way that the ratio measurements exhibit minimal rate dependence. While the correction procedure is kinematic dependent and leaves much to be desired, it is encouraging to have made progress on an analytical roadblock that has proven to be problematic. It is the opinion of the author that with further studies of the principles and methods described in this thesis with much larger quantities of data, a correction may soon be devised that utilize *only* spectrometer variables and not kinematics (i.e. efficiencies based on a single track's trigger road and that event's intensity only).

The results presented give the first peek into what is certain to be an impact measurement. Already, with very limited statistics, SeaQuest is surpassing previous measurements (E-772) in precision and range. In the regions of measurement overlap, the preliminary results seems to indicate that the signal measured by SeaQuest is largely consistent with what has been seen in the E-772 result. As the results currently stand, there appears to (still) be little nuclear A-dependence of the Drell-Yan cross section ratio, both integrated

and versus x_2 . This implies that perhaps the nuclear modification that leads to the DIS phenomenon known as the EMC effect is exclusively an effect of the valence quarks and does not originate from the quark sea.

As the rate dependence correction becomes more refined, as large swaths of data not analyzed here are included, and as the tracking and analytical cuts continue to be tuned, the author is confident that the total yield of SeaQuest will provide a truly precise window into the behavior of sea quarks at high x_2 . Additional improvements would be to precisely analyze the compositions of the liquid deuterium target used, thereby completely eliminating one of the largest sources of systematic uncertainties.

In addition to the execution of this analytical measurement, some key contributions to the operation and performance of the experiment have been outlined. This includes the testing, prototyping, manufacturing, calibration, and installation of new photomultiplier tubes bases that are more capable of handling the high instantaneous rates that at times occur at SeaQuest. Also discussed is the management and curation of the experiment's raw, processed, and tracked data on a highly-available and queryable SQL database. This platform has allowed users to perform their analysis quickly and flexibly, while allowing the collaboration to effectively share with each other and work from virtually any analysis framework.

In conclusion, this limited amount of data produced by SeaQuest has already increased the precision and constraints on the nuclear modification of sea quarks in the nucleon. With further refinement of these procedures and additional data, SeaQuest will assuredly be able to provide the most precise high- x_2 data on the A-dependence of the Drell-Yan process.

References

- [1] E. A. e. a. Hawker, “Measurement of the light antiquark flavor asymmetry in the nucleon sea,” *Phys. Rev. Lett.*, vol. 80, pp. 3715–3718, Apr 1998.
- [2] 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, Aug 1970.
- [3] W. a Heisenberg, “Über den anschaulichen Inhalt der quantentheoretischen Kinematik und Mechanik,” *Z. Phys.*, vol. 43, pp. 172–198, 1927.
- [4] B. Povh, *Particles and Nuclei: An Introduction to the Physical Concepts*. Physics and astronomy online library, Springer, 2002.
- [5] S. Bethke, “Experimental tests of asymptotic freedom,” *Prog. Part. Nucl. Phys.*, vol. 58, pp. 351–386, 2007.
- [6] J. C. Collins and D. E. Soper, “Angular distribution of dileptons in high-energy hadron collisions,” *Phys. Rev. D*, vol. 16, pp. 2219–2225, Oct 1977.
- [7] F. Halzen and A. D. Martin, *Quarks and Leptons: An Introductory Course in Modern Particle Physics*. Wiley, 1984. New York, USA (1984) 396p.
- [8] 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-ibid.25:902,1970].
- [9] S. e. a. Dasu, “Precision measurement of $r = \frac{\sigma_L}{\sigma_T}$ and F_2 in deep-inelastic electron scattering,” *Phys. Rev. Lett.*, vol. 61, pp. 1061–1064, Aug 1988.
- [10] S. e. a. Dulat, “New parton distribution functions from a global analysis of quantum chromodynamics,” *Phys. Rev.*, vol. D93, no. 3, p. 033006, 2016.
- [11] D. Griffiths, *Introduction to Elementary Particles*. Physics textbook, Wiley, 2008.
- [12] D. J. Gross and C. H. Llewellyn Smith, “High-energy neutrino - nucleon scattering, current algebra and partons,” *Nucl. Phys.*, vol. B14, pp. 337–347, 1969.
- [13] H. E. Fisk and F. Sciulli, “Charged Current Neutrino Interactions,” *Ann. Rev. Nucl. Part. Sci.*, vol. 32, pp. 499–573, 1982.
- [14] K. Gottfried, “Sum rule for high-energy electron - proton scattering,” *Phys. Rev. Lett.*, vol. 18, p. 1174, 1967.
- [15] A. S. e. a. Ito, “Measurement of the continuum of dimuons produced in high-energy proton-nucleus collisions,” *Phys. Rev. D*, vol. 23, pp. 604–633, Feb 1981.
- [16] C. N. e. a. Brown, “Dimuon Production in 800-GeV Proton Nucleus Collisions,” *Phys. Rev. Lett.*, vol. 63, pp. 2637–2640, 1989.

- [17] C. Grosso-Pilcher and M. J. Shochet, “High Mass Dilepton Production in Hadron Collisions,” *Ann. Rev. Nucl. Part. Sci.*, vol. 36, pp. 1–28, 1986.
- [18] G. E. e. a. Hogan, “Comparison of Muon Pair Production to the Quark - Anti-quark Annihilation Model,” *Phys. Rev. Lett.*, vol. 42, p. 948, 1979.
- [19] I. R. Kenyon, “The Drell-Yan Process,” *Rept. Prog. Phys.*, vol. 45, p. 1261, 1982.
- [20] T. Matsuura, S. C. van der Marck, and W. L. van Neerven, “The Calculation of the Second Order Soft and Virtual Contributions to the Drell-Yan Cross-Section,” *Nucl. Phys.*, vol. B319, p. 570, 1989.
- [21] C.-G. Duan, S.-W. Cui, Z.-Y. Yan, and G.-L. Li, “Quark energy loss and shadowing in nuclear Drell-Yan process,” *Commun. Theor. Phys.*, vol. 44, pp. 123–127, 2005.
- [22] D. e. a. Geesaman. Proposal for Drell-Yan Measurements of Nucleon and Nuclear Structure with the FNAL Main Injector, 1999.
- [23] J. e. a. Aubert, “The ratio of the nucleon structure functions $F2_n$ for iron and deuterium,” *Phys.Lett.*, vol. B123, p. 275, 1983.
- [24] A. e. a. Bodek, “Electron scattering from nuclear targets and quark distributions in nuclei,” *Phys. Rev. Lett.*, vol. 50, pp. 1431–1434, May 1983.
- [25] D. F. Geesaman, K. Saito, and A. W. Thomas, “The nuclear EMC effect,” *Ann.Rev.Nucl.Part.Sci.*, vol. 45, pp. 337–390, 1995.
- [26] P. R. Norton, “The emc effect,” *Reports on Progress in Physics*, vol. 66, no. 8, p. 1253, 2003.
- [27] R. G. e. a. Arnold, “Measurements of the a-Dependence of Deep Inelastic electron Scattering from Nuclei,” *Phys. Rev. Lett.*, vol. 52, p. 727, 1984.
- [28] J. e. a. Gomez, “Measurement of the A-dependence of deep inelastic electron scattering,” *Phys. Rev.*, vol. D49, pp. 4348–4372, 1994.
- [29] A. C. e. a. Benvenuti, “Nuclear Effects in Deep Inelastic Muon Scattering on Deuterium and Iron Targets,” *Phys. Lett.*, vol. B189, p. 483, 1987.
- [30] P. e. a. Amaudruz, “The ratio $F2(n) / F2(p)$ in deep inelastic muon scattering,” *Nucl. Phys.*, vol. B371, pp. 3–31, 1992.
- [31] K. e. a. Ackerstaff, “Nuclear effects on $R = \sigma(L) / \sigma(T)$ in deep inelastic scattering,” *Phys. Lett.*, vol. B475, pp. 386–394, 2000. [Erratum: *Phys. Lett.*B567,339(2003)].
- [32] X. Chen, J. Ruan, R. Wang, P. Zhang, and W. Zhu, “Applications of a nonlinear evolution equation II: the EMC effect,” *Int. J. Mod. Phys.*, vol. E23, no. 10, p. 1450058, 2014.
- [33] A. Bodek and J. L. Ritchie, “Fermi-motion effects in deep-inelastic lepton scattering from nuclear targets,” *Phys. Rev. D*, vol. 23, pp. 1070–1091, Mar 1981.
- [34] S. Stein, W. B. Atwood, E. D. Bloom, R. L. Cottrell, H. C. DeStaebler, C. L. Jordan, H. Piel, C. Y. Prescott, R. Siemann, and R. E. Taylor, “Electron Scattering at 4-Degrees with Energies of 4.5-GeV - 20-GeV,” *Phys. Rev.*, vol. D12, p. 1884, 1975.
- [35] J. e. a. Seely, “New measurements of the EMC effect in very light nuclei,” *Phys. Rev. Lett.*, vol. 103, p. 202301, 2009.
- [36] K. Rith, “Present Status of the EMC effect,” *Subnucl. Ser.*, vol. 51, pp. 431–449, 2015.
- [37] S. e. a. Dasu, “Measurement of kinematic and nuclear dependence of $R = \sigma-L / \sigma-t$ in deep inelastic electron scattering,” *Phys. Rev.*, vol. D49, pp. 5641–5670, 1994.

- [38] M. e. a. Arneodo, “Measurement of the proton and deuteron structure functions, $F_2(p)$ and $F_2(d)$, and of the ratio σ_L / σ_T ,” *Nucl. Phys.*, vol. B483, pp. 3–43, 1997.
- [39] V. Guzey, L. Zhu, C. E. Keppel, M. E. Christy, D. Gaskell, P. Solvignon, and A. Accardi, “Impact of nuclear dependence of $R = \sigma_L / \sigma_T$ on antishadowing in nuclear structure functions,” *Phys. Rev.*, vol. C86, p. 045201, 2012.
- [40] B. Z. Kopeliovich, J. G. Morfin, and I. Schmidt, “Nuclear Shadowing in Electro-Weak Interactions,” *Prog. Part. Nucl. Phys.*, vol. 68, pp. 314–372, 2013.
- [41] K. A. e. a. Griffioen, “Measurement of the EMC Effect in the Deuteron,” *Phys. Rev.*, vol. C92, no. 1, p. 015211, 2015.
- [42] E. Piasetzky, L. Weinstein, D. Higinbotham, J. Gomez, and O. e. a. Hen, “Short range correlations and the EMC effect,” *Nucl.Phys.*, vol. A855, pp. 245–248, 2011.
- [43] N. e. a. Fomin, “New measurements of high-momentum nucleons and short-range structures in nuclei,” *Phys. Rev. Lett.*, vol. 108, p. 092502, 2012.
- [44] K. S. e. a. Egiyan, “Measurement of two- and three-nucleon short-range correlation probabilities in nuclei,” *Phys. Rev. Lett.*, vol. 96, p. 082501, Mar 2006.
- [45] L. L. Frankfurt, M. I. Strikman, D. B. Day, and M. Sargsian, “Evidence for short range correlations from high Q^{*2} (e, e-prime) reactions,” *Phys. Rev.*, vol. C48, pp. 2451–2461, 1993.
- [46] L. B. Weinstein, E. Piasetzky, D. W. Higinbotham, J. Gomez, O. Hen, and R. Shneur, “Short range correlations and the emc effect,” *Phys. Rev. Lett.*, vol. 106, p. 052301, Feb 2011.
- [47] M. Alvioli, C. Ciofi degli Atti, and H. Morita, “Proton-neutron and proton-proton correlations in medium-weight nuclei and the role of the tensor force,” *Phys. Rev. Lett.*, vol. 100, p. 162503, Apr 2008.
- [48] K. e. a. Egiyan, “Measurement of 2- and 3-nucleon short range correlation probabilities in nuclei,” *Phys.Rev.Lett.*, vol. 96, p. 082501, 2006.
- [49] D. W. Higinbotham, J. Gomez, and E. Piasetzky, “Nuclear Scaling and the EMC Effect,” 2010.
- [50] O. Hen, L. B. Weinstein, S. Gilad, and S. A. Wood, “In Medium Nucleon Structure Functions, SRC, and the EMC effect,” 2014.
- [51] O. Hen, E. Piasetzky, and L. Weinstein, “New data strengthen the connection between Short Range Correlations and the EMC effect,” *Phys.Rev.*, vol. C85, p. 047301, 2012.
- [52] E. L. Berger, “Nuclear Effects in Massive Lepton Pair Production,” *Nucl. Phys.*, vol. B267, p. 231, 1986.
- [53] M. A. e. a. Vasilev, “Parton energy loss limits and shadowing in Drell-Yan dimuon production,” *Phys. Rev. Lett.*, vol. 83, pp. 2304–2307, 1999.
- [54] D. M. e. a. Alde, “Nuclear dependence of dimuon production at 800 gev,” *Phys. Rev. Lett.*, vol. 64, pp. 2479–2482, May 1990.
- [55] J. Carlson and R. Schiavilla, “Structure and dynamics of few nucleon systems,” *Rev. Mod. Phys.*, vol. 70, pp. 743–842, 1998.
- [56] T. Zimmerman and J. Hoff, “The design of a charge integrating, modified floating point ADC chip,” *IEEE J. Solid State Circuits*, vol. 39, pp. 895–905, 2004.
- [57] A. e. a. Alavi-Harati, “Measurements of direct CP violation, CPT symmetry, and other parameters in the neutral kaon system,” *Phys. Rev. D*, vol. 67, p. 012005, Jan 2003.
- [58] Shieldwerx, “[5http://www.shieldwerx.com/assets/swx-201\(hd\).pdf](http://www.shieldwerx.com/assets/swx-201(hd).pdf), 2015.

- [59] G. e. a. Moreno, “Dimuon production in proton-copper collisions at $\sqrt{s} = 38.8$ gev,” *Phys. Rev. D*, vol. 43, pp. 2815–2835, May 1991.
- [60] R. S. e. a. Towell, “Improved measurement of the anti-d/anti-u asymmetry in the nucleon sea,” *Phys. Rev.*, vol. D64, p. 052002, 2001.
- [61] CAEN, “v1495 general purpose vme board.” <http://www.caen.it/csite/CaenProd.jsp?idmod=484&parent=11>. Accessed: 2016-01-26.
- [62] S.-H. e. a. Shiu, “FPGA-based trigger system for the Fermilab SeaQuest experiment,” *Nucl. Instrum. Meth.*, vol. A802, pp. 82–88, 2015.
- [63] MVME, “Mvme5500 series vmebus single-board computer.” <http://www.mvme.com/manuals/MVME5500e-SPEC.pdf>. Accessed: 2016-01-26.
- [64] J. N. A. L. D. A. Support, “Coda (ceba online data acquisition).” <https://coda.jlab.org/drupal/>. Accessed: 2016-02-12.
- [65] O. S. G. Consortium, “Osg blueprint.” <http://osg-docdb.opensciencegrid.org/cgi-bin/RetrieveFile?docid=18&extension=pdf>. Accessed: 2016-02-24.
- [66] W. Schools, “Sql tutorial.” <http://www.w3schools.com/sql/>. Accessed: 27 April 2016.
- [67] Oracle, “Mysql 5.7 reference manual.” <https://dev.mysql.com/doc/refman/5.7/en/>. Accessed: 27 April 2016.
- [68] I. Mandrichenko, “Big data in high energy physics.” <http://cd-docdb.fnal.gov/cgi-bin/ShowDocument?docid=5358>. Accessed: 29 April 2016.
- [69] J. M. Andre *et al.*, “Online Data Handling and Storage at the CMS Experiment,” *J. Phys. Conf. Ser.*, vol. 664, no. 8, p. 082009, 2015.
- [70] H. Photonics, “Xp2072: High energy resolution , 10-stage, 39mm (1.5”) tube.” <http://hzcphotonics.com/products/XP2072.PDF>. Accessed: 2016-01-29.
- [71] A. Einstein, “Concerning an heuristic point of view toward the emission and transformation of light,” *Annalen Phys.*, vol. 17, pp. 132–148, 1905.
- [72] C. R. Kerns, “A High Rate Phototube Base,” *IEEE Trans. Nucl. Sci.*, vol. 24, pp. 353–355, 1977.
- [73] B. Ohare, “Mosfet structure.” https://commons.wikimedia.org/wiki/File:MOSFET_Structure.png, Mar. 2012. Accessed: 2016-02-01.
- [74] S. Hansen. private communication, 2012.
- [75] E. e. a. Ames, “Measurement of 20 hamamatsu r-5900 photomultiplier tubes for atlas tilecal module-0,” tech. rep., CERN-ATL-TILECAL-98-156, 1998.
- [76] Agilent, “Agilent 33500 series 30 mhz function / arbitrary waveform generator users guide.” <http://cp.literature.agilent.com/litweb/pdf/33520-90001.pdf>. Accessed: 2016-02-03.
- [77] LeCroy, “Lecroy 1440 high voltage power supply manual.” http://ikpe1101.ikp.kfa-juelich.de/manuals/LeCroy_1440/operation_manual/operation_manual.htm. Accessed: 2016-02-03.
- [78] P. Webb and R. Williams, “Gamma ray spectroscopy using a germanium lithium-drifted diode,” *Nucl. Instr. Methods*, vol. Vol: 22, Apr 1963.
- [79] R. M. van der Eijk, *Track reconstruction in the LHCb experiment*. PhD thesis, Amsterdam U., 2002.
- [80] 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.

- [81] S. e. a. Gorbunov, “ALICE HLT high speed tracking and vertexing,” in *Proceedings, 17th Conference on Computing Applications in Nuclear and Plasma Sciences on Real Time (RT2010)*, 2010.
- [82] M.-j. Wang, *A precision measurement of the A-dependence of dimuon production in proton - nucleus collisions at 800-GeV/c*. PhD thesis, Case Western Reserve U., 1991.
- [83] B. List, “Statistical error of efficiency determination from weighted events.” Accessed: 2016-03-16.
- [84] J. Arrington, F. Coester, R. J. Holt, and T. S. H. Lee, “Neutron Structure Functions,” *J. Phys.*, vol. G36, p. 025005, 2009.
- [85] O. Hen, D. W. Higinbotham, G. A. Miller, E. Piassetzky, and L. B. Weinstein, “The EMC Effect and High Momentum Nucleons in Nuclei,” *Int. J. Mod. Phys.*, vol. E22, p. 1330017, 2013.
- [86] S. J. Brodsky, A. Hebecker, and E. Quack, “The Drell-Yan process and factorization in impact parameter space,” *Phys. Rev.*, vol. D55, pp. 2584–2590, 1997.
- [87] H. Jung and G. A. Miller, “Pionic contributions to deep inelastic nuclear structure functions,” *Phys. Rev.*, vol. C41, pp. 659–664, 1990.
- [88] G. A. Miller, “Nuclear Wave Functions in Deep Inelastic Scattering and Drell-Yan Processes,” in *Workshop on Nuclear and Particle Physics on the Light Cone Los Alamos, New Mexico, July 18-22, 1988*, 1988.
- [89] S. A. Kulagin and R. Petti, “Nuclear Parton Distributions and the Drell-Yan Reaction,” *PoS*, vol. DIS2015, p. 040, 2015.
- [90] S. A. Kulagin and R. Petti, “Personal communication.”