A Precision Low-Energy Measurement of the Weak Mixing Angle in Møller Scattering

Thesis by

Peter A. Mastromarino

In Partial Fulfillment of the Requirements

for the Degree of

Doctor of Philosophy

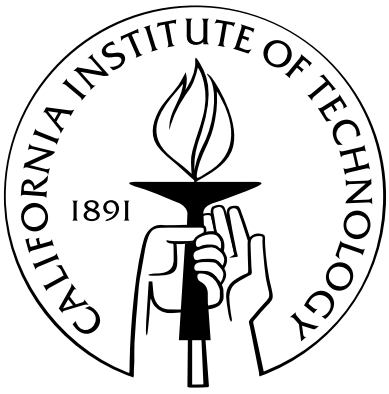

California Institute of Technology

Pasadena, California

2004

(Defended May 10, 2004)

Acknowledgements

Writing a PhD thesis is the culmination of many years of learning and labor. This thesis represents not only my own work, but also the work of countless others, some whose contributions are clear and tangible, and others who have may have contributed more indirectly. So here I'd like to give recognition to the many other members of the greater "collaboration" who have made this entire work possible.

First, of course, there is my advisor, Emlyn Hughes, who in a very real sense did actually make all of this happen through his financial support. But more than that, he has given me many interesting insights into the world around me, which helped me throughout my time as a graduate student. Emlyn has an amazing ability to penetrate right to the core of matters, immediately discerning the most important features of whatever problem lies at hand. This is as useful in science and business as it is in everyday life, and it's a skill I admire and value very much. In a way, it's basically what drew me to physics in the first place. Simply put, the goal of physics is to get to the bottom of it all, to figure out nature (or a description of nature) at its most fundamental level. In order to do this, in order to formulate an elegant, predictive and self-consistent description of any complicated system, one must always simplify, concentrating on just the essential features. Of course, the skill comes in when determining what's essential and what's not. That's what Emlyn's good at, and that's what an education in physics really teaches, I think.

Much of my time as a graduate student was spent at the Stanford Linear Accelerator Center (SLAC), and there are lots of people there I have to thank. I worked as part of a collaboration of roughly sixty physicists, and all of them deserve credit, for without them the experiment which this thesis describes would not have been possible. There are several I would like to thank specifically. I have known Krishna Kumar since I was an undergraduate, when he was my advisor for three years. Since then, I have worked with him for four of my five years as a graduate student. I am extremely thankful for his unceasing willingness

to share with me his profound knowledge of physics, and for his sound guidance during all of my years as a student. While at SLAC, I worked very closely with Michael Woods (eventually extracting from him that his middle name is Bernard). I cannot overstate how fortunate I am to have had a close mentor like him. It was a real experience to see an individual in whom passion and work fused so seamlessly. Brilliant, communicative, and possessing of a seemingly inexhaustible supply of patience, Mike was the perfect match for a young physicist at a new lab, eager to learn the ropes.

A lot of the time at SLAC I found myself involved with work at the polarized electron source, where I learned a lot about lasers and optics, photocathodes, and general hardware. During all of this time, the entire SLAC source group, but especially James Turner, Axel Brachmann, and Tom Galetto, were invaluable. In addition to being instructive (and absolutely vital to the success of the entire experiment), they were good friends. And although not a permanent member of the source group, T. Brian Humensky (the "T" stands for "Thomas") should certainly be considered in the same category. He was a fellow graduate student who had worked at the source before I arrived at SLAC, and I definitely learned a lot from him, not the least of which was, sleep is important for the human body.

There were many others who went "above and beyond" the call of duty, to whom I am indebted. Yury Kolomensky was a marvel of efficiency, skill, and the never-complain, can-do attitude that has made America (or, in his case, Russia) great. He taught me a lot, particularly about the intricacies of our analysis software. On a different note, I always enjoyed my conversations with Paul Souder, unless they were the ones that lasted from 8am to 10am after I had just finished an owl shift. His formidable intellect was augmented by a lifetime spent in physics, and so it was always a unique pleasure to be in his company. A mention of Paul deserves a mention of Ted Fieguth, for he too was a colorful character with many stories and tidbits of wisdom to share. Ted was supremely knowledgeable about SLAC, having been there for many years, and he was kind enough to take me under his wing when it was important in the beginning. He introduced me to Howard Smith and Franz-Josef Decker, both of whom were great at answering my questions about accelerator physics, even though they were terribly busy. He also provided my first introduction to Dieter Walz, who has been at SLAC from the very beginning (1962). Dieter was a veritable storehouse of information of all sorts, particularly when it came to two things: (a) beam interactions with matter and (b) running. Dieter frequently runs 100 mile ultra-marathons,

a fact that will likely never cease to amaze (and perhaps even inspire) me.

At SLAC, I captained a lot of shifts (nearly 150, to be precise). That I didn't end up captaining even more is due to the much-appreciated efforts of my fellow graduate students: G. Mark Jones (the "G" stands for "Gary," wielder of spears), Klejda Bega, David Ricci Relyea, Baris Tonguc, Imran Younus, Waled Emam, Antonin Vacheret, Mark Cooke, and Brock Tweedie. Then there was the "unstoppable force," postdoctoral scholar Carlos Gerardo Arroyo Rizo-Patron, who hails from Peru. These were my fellow shift captains, and they, more than anyone else, made my "tour of duty" at SLAC truly memorable. Thanks to this cast of characters, I have enough "gag reels" with which to entertain myself for a lifetime. I am especially indebted to Mark Jones, for many stimulating discussions (covering everything from physics to how many baboons it would take to match up against Mark McGwire with a baseball bat lashed to his forearm), lots of crazy antics, and tons of laughs. In addition, of course, there were many others at SLAC with whom I was able to forge a special bond, or who were particularly helpful, or who made my time especially enjoyable. In particular, I'd like to thank Michael Olson (MO), Owen Saxton, Zenon Szalata, Ray Arnold, Kent Paschke, Piotr Decowski, Eugene Chudakov, Ken Baird, Lisa Kaufman, David Lhuillier, and Alexander Deur, for all their help and for being all-around wonderful people. I'd also like to thank Rich Holmes for once again helping with and contributing to the dithering analysis software.

Before I went to SLAC, I had a great first year of graduate school at Caltech, largely because I was fortunate enough to wind up in a group with some really terrific people in it. Some of them I have already thanked, but others I have not. I would later run into Jaideep Singh again at SLAC (where he was kind enough to relieve some pressure on my experiment's manpower by taking some shifts), but even during my first year, I got to know what an extremely competent, funny, and kind individual he really is. Tina Pavlin was always kindhearted and a good friend, even though I may have occasionally annoyed her with certain habits of mine (whistling comes to mind). Steffen Jensen blazed the trail the other graduate students and I would later follow. Though I only knew him for one year, Steffen has left a big impression in my mind and was an incredibly fortuitous role model for a first-year graduate student like me. A special acknowledgment must go to David Pripstein, the postdoctoral scholar in our group at the time, and the person who was no doubt largely responsible for my decision to come to Caltech in the first place. Dave is one of the most

unique characters I have ever known, with enough catch-phrases and idiosyncrasies to fill a notebook, and he is also one of the most dedicated, intelligent, and genuine people I am ever likely to meet. And since I've returned to Caltech, I have been fortunate enough to get to know Guodong Wang and Wenjin Shao, who didn't join the group until about the time I left for SLAC. Both Guodong and Wenjin have been extremely gracious in accommodating the return of Mark, Klejda, and I. Finally, I'd like to thank Georgia Frueh for her tireless efforts and for always looking out for the needs of all of the graduate students.

I learned a great deal from those people whose names are listed above, but they are merely the latest in a long line of teachers and mentors I have known over the years. Of course, space constraints (which I usually pay no attention to, anyway) do not permit me to list the names of all of them. I have already thanked Krishna Kumar, my first physics professor at Princeton. I had many other wonderful professors there (Daniel Marlow, Bernhard Keimer, Edward Groth, Stewart Smith, Jeremy Goodman, and Neta Bahcall among them) who encouraged and/or inspired me to continue on the physics path. I also had many wonderful professors in subjects other than physics, math, and astronomy (the courses of Anthony Grafton and James McPherson immediately spring to mind) who, in general, fanned the flames of learning that must be kept alive within any would-be scholar. But a special debt is owed to my earlier teachers, my elementary and high school teachers (not so much middle school, which was my "difficult" phase). These were the teachers who truly had the biggest impact on me, encouraging me to write, filling me with the desire to learn as much as possible about the world around me, and giving me the confidence that is vitally important to the achievement of anything. It's an old saying that, "Education is what remains after you've forgotten all that you have been taught." Much of what I learned from these teachers, however, will *never* be forgotten. I especially want to mention the names of Mrs. Jay and Mrs. Fernelius (English), Mrs. Haley (history), Mr. Parr (calculus), Mr. Penner (economics), and Coach Kreuzer (soccer/life).

The most important thanks I have to give go to my parents, whose love and sacrifice have paved the way to where I am today. They never told me what to think, but they taught me what was important. Their examples illuminate my path, and their love urges me onwards in my explorations. I am also incredibly grateful to my wonderful sister, Carrie, and to the rest of my family (grandparents, aunts, and uncles), for always supporting and encouraging me. Lastly, I would like to thank my amazing fiancée, Catherine Wong. I

thank her for her kindness, her strength, her intelligence, her sense of humor, and her faith in me. Not only has she helped make all of my efforts possible, she has also made them worthwhile! I consider myself infinitely lucky to be spending the rest of my life with her.

Abstract

The E-158 experiment at the Stanford Linear Accelerator Center (SLAC) measures the parity-violating cross-section asymmetry in electron–electron (Møller) scattering at low Q^2. This asymmetry, whose Standard Model prediction is roughly -150 parts per billion (ppb), is directly proportional to $(1 - 4\sin^2\theta_W)$, where θ_W is the weak mixing angle. Measuring this asymmetry to within 10% provides an important test of the Standard Model at the quantum loop level and probes for new physics at the TeV scale.

The experiment employs the SLAC 50 GeV electron beam, scattering it off a liquid hydrogen target. A system of magnets and collimators is used to isolate and focus the Møller scattering events into an integrating calorimeter. The electron beam is generated at the source using a strained, gradient-doped GaAs photocathode, which produces roughly 5×10^{11} electrons/pulse (at a beam rate of 120 Hz) with ∼80% longitudinal polarization. The helicity of the beam can be rapidly switched, eliminating problems associated with slow drifts. Helicity-correlations in the beam parameters (charge, position, angle and energy) are minimized at the source and corrected for using precision beam monitoring devices.

The parity-violating cross-section asymmetry $A_{\rm PV}$ in Møller scattering is measured to be $A_{\rm PV} = -160 \pm 21\,({\rm stat}) \pm 16\,({\rm syst})$ ppb, at an average Q^2 of 0.026 GeV^2. This represents the first observation of parity violation in Møller scattering, and corresponds to the following low-energy determination of the weak mixing angle:

$$\sin^2\theta_W(Q^2 = 0.026\ {\rm GeV}^2)_{\overline{\rm MS}} = 0.2381 \pm 0.0015\ ({\rm stat}) \pm 0.0014\ ({\rm syst}).$$

This agrees with the Standard Model prediction of 0.2385 ± 0.0006. Roughly half of the experiment's total data set is represented here. This thesis provides a full description of the experimental method and analysis procedure used to obtain the above result. It also discusses the result's physical implications in terms of possible extensions to the Standard Model.

Contents

List of Figures

List of Tables

Chapter 1

Theoretical Motivation

1.1 Introduction: The Standard Model

The modern theory of elementary particle physics, built upon the idea of imposing local $\mathrm{SU}(3)_C \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ gauge invariance on a free Lagrangian, and then spontaneously breaking this underlying symmetry via the Higgs mechanism, has emerged over the past several decades as the simplest and most successful model for describing all of the fundamental particles of nature and their interactions (modulo gravity). It has thus truly earned its place as the "Standard Model" of the field. Imposing the various local gauge symmetries on the free Lagrangian automatically generates the force-mediating gauge bosons [1, 2]. In the $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ electroweak sector, these acquire mass by the Higgs mechanism, in which the self-interactions of a doublet of complex scalar fields in the original Lagrangian produce nonzero vacuum expectation values [3]. Perturbation theory can be applied around a physical groundstate; this, however, spoils (or "breaks") the symmetry of the original Lagrangian.[1] A convenient gauge transformation can be performed whereby only a single scalar Higgs field out of the original complex doublet remains. Diagonalizing the mass matrices in the transformed Lagrangian, one finds that three of the four physically observable vector bosons in the electroweak theory, the $W^{\pm}$ and Z, have acquired mass terms (or, equivalently, longitudinal polarization states) [4]. The fourth vector boson, the photon, remains massless, a reflection of the unbroken $\mathrm{U}(1)_Q$ symmetry.

The masses $m_{W^{\pm}}$ and m_Z are not completely independent parameters. The Standard

[1] Spontaneous symmetry-breaking also generates the fermion masses. In a sense, the starting Lagrangian is not really "free," since it does include Yukawa couplings between the fermions and the Higgs doublet. When a nonzero vacuum expectation value is introduced by the Higgs scalar, mass terms directly proportional to this value, but including the original Yukawa coupling strengths as multiplicative factors, result for all of the different fermions.

Model makes firm predictions for how they should relate to one another. First, the masses of the W^+ and the W^- are equal, so that $m_{W^\pm} = m_W$. Second, m_Z and m_W are related by the following equation:

$$\cos\theta_W = \frac{m_W}{m_Z} \tag{1.1}$$

Here θ_W is the so-called weak mixing angle, which describes the degree of mixing between the gauge bosons of the underlying $\mathrm{SU}(2)_L$ and $\mathrm{U}(1)_Y$ symmetries. More specifically, after the spontaneous symmetry-breaking procedure described above, the following mass eigenstates for the electroweak vector gauge fields are found [4]:

$$W_\mu^\pm = \frac{1}{\sqrt{2}}\left(A_\mu^1 \mp iA_\mu^2\right) \tag{1.2}$$

$$Z_\mu = -\sin\theta_W B_\mu + \cos\theta_W A_\mu^3 \tag{1.3}$$

$$A_\mu = \cos\theta_W B_\mu + \sin\theta_W A_\mu^3 \tag{1.4}$$

Here $(A_\mu^1,\, A_\mu^2,\, A_\mu^3)$ is the original isotriplet of vector bosons needed for local $\mathrm{SU}(2)_L$ gauge invariance, and B_μ is the original isosinglet vector boson needed for local $\mathrm{U}(1)_Y$ gauge invariance. Writing the interaction Lagrangian in terms of these eigenstates, the combinations $W_\mu^\pm$ are recognized as the mediators of the charged weak current, while the combinations Z_μ and A_μ are recognized as the mediators of the neutral weak and electromagnetic currents, respectively. The $\mathrm{SU}(2)_L{\times}\mathrm{U}(1)_Y$ electroweak theory is often referred to as the Glashow-Weinberg-Salam (GWS) theory, after the three physicists who made significant contributions towards its complete formulation [4, 5, 6].

At this point in any discussion of the Standard Model, it is customary to reflect on its many impressive successes, beginning with its prediction of the neutral weak current.[2] Unification of the charged weak and electromagnetic interactions into a single $\mathrm{SU}(2){\times}\mathrm{U}(1)$ symmetry group required the existence of a fourth gauge boson. At the time the model was being developed in the 1960's, however, no experimental evidence existed for such a particle, presumably because it was predicted to be very heavy, with a mass of at least 80 GeV. Confirmation of its existence did not come until 1973, when the Gargamelle bubble chamber at CERN finally gave the first evidence for a $\nu_\mu e \to \nu_\mu e$ scattering event, the unequivocal sign of a neutral weak current interaction [8].

[2] The prediction of the neutral weak current actually goes back to 1958 [7]. However, it was not until the GWS theory that its role was actually appreciated.

In addition to correctly predicting a fundamentally new interaction, the theory also predicted a very specific helicity structure, which could be experimentally tested. The GWS theory incorporated left-handed quarks and leptons into isospin doublets, and right-handed quarks and leptons into isospin singlets. Alternative theories also existed, an example being the so-called "hybrid" model, wherein right-handed leptons were assigned to isospin doublets instead of singlets [9]. Continued neutrino observations at CERN had proven unable to rule out such alternatives. Thus, in 1978, a fixed target experiment at SLAC, involving inclusive electron–deuteron deep inelastic scattering, measured the first-ever parity-violating asymmetry in a neutral current interaction [10, 11]. In general, different theories — such as the GWS $\mathrm{SU}(2)_L{\times}\mathrm{U}(1)_Y$ theory and the $\mathrm{SU}(2)_L{\times}\mathrm{SU}(2)_R{\times}\mathrm{U}(1)_Y$ "hybrid" theory — gave different predictions for the parity-violating asymmetry in this process [12]. The value obtained by the experiment agreed precisely with that of the GWS theory, further establishing the theory's increasing reputation as the "standard model" of particle physics.

The final major piece of experimental evidence in support of the GWS theory came in the early 1980's, after the completion of CERN's proton–antiproton collider ring. Used in conjunction with the results of a variety of weak interaction experiments, including additional neutrino scattering experiments at CERN and muon decay measurements, the data from the SLAC parity-violating electron scattering experiment could provide a value for $\sin^2\theta_W$ that was good to within a couple percent.[3] The GWS theory made the following prediction for how $\sin^2\theta_W$ should relate to the coupling constant G_F of Fermi's original beta decay theory [4, 13]:

$$G_F = \frac{\pi\alpha}{\sqrt{2}m_W^2\sin^2\theta_W} \tag{1.5}$$

Since G_F was already known fairly accurately from earlier nuclear experiments, a measurement of $\sin^2\theta_W$ could be turned into a prediction for m_W and, using Equation (1.1), m_Z:

$$m_W = 82 \pm 2 \text{ GeV} \qquad m_Z = 92 \pm 2 \text{ GeV} \tag{1.6}$$

In 1983, a group at CERN announced the discovery of the W and Z vector gauge bosons at 81 ± 5 GeV [14] and 95 ± 3 GeV [15], respectively. This was in stunning agreement with

[3] Specifically, the value obtained from this early electroweak data was $\sin^2\theta_W = 0.207 \pm 0.005$. However, this result is quoted in the on-shell renormalization scheme, which is different from the modified minimal subtraction scheme used throughout the rest of this thesis.

the predicted values obtained from the combined analysis of the world's electroweak data and further confirmation of the GWS theory.

This example perfectly typifies the fundamental success of the Standard Model, which contains only a finite number of free parameters (see below), but yet which must accommodate the findings of numerous different experiments. In general, each experiment provides sensitivity to a different combination of the theory's free parameters. Thus, combining results from multiple experiments tests the internal consistency of the model. Occasionally it may even result in the predictions of particle masses before the particles themselves are actually observed, as in the case of the W and Z. This in fact happened once again with the top quark. By the mid 1990's, the precision of the world's electroweak data had improved to the point that the mass of the as-yet-undiscovered top quark could be predicted at $m_t = 177 \pm 20$ GeV [16]. In 1995 the elusive quark was discovered at the Tevatron, precisely where the electroweak data had predicted it to be [17]. Today its mass can be placed at $m_t = 174 \pm 5$ GeV [18]. Currently, a similar story could be unfolding with the Higgs scalar, whose mass is proportional to the vacuum expectation value of its self-interaction, another free parameter of the Standard Model (in one formulation, at least). A global fit to all electroweak precision data yields the prediction $45 \leq m_H \leq 191$ GeV (90% C.L.) for the Higgs mass [18]. Meanwhile, direct searches at LEP provide a 95% C.L. lower bound of $m_H \gtrsim 114$ GeV [19]. Run II of the Tevatron is currently extending this limit [20]. Ultimately, the Large Hadron Collider at CERN should be able to explore a mass range of 100 GeV to 1 TeV [21]. Should the Higgs be discovered somewhere in this range of energies (and self-consistency within the model suggests that it should be), and should the discovery agree with the prediction obtained from the electroweak fit, it would be yet another triumph for the Standard Model.

Despite the unflagging and at times spectacular success of the Standard Model, it is likely not destined to be a "final" theory. The reason is that it is too *ad hoc* and leaves too many unanswered questions [22]. For instance, why the $\mathrm{SU}(3)_C$, $\mathrm{SU}(2)_L$ and $\mathrm{U}(1)_Y$ gauge symmetries, and why is the $\mathrm{SU}(2)_L$ symmetry only left-handed? Why are there three generations of fermions? Why is charge quantized? Why is the $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ symmetry spontanesouly broken by the vacuum down to $\mathrm{U}(1)_Q$, providing masses for all of the fermions as well as the weak gauge bosons? The puzzles surrounding the origin of mass naturally lead to questions concerning the many free parameters of the Standard Model, of which there

are 19. First, there are the three gauge couplings g_3, g_2, and g_1 of the underlying $\mathrm{SU}(3)_C$, $\mathrm{SU}(2)_L$ and $\mathrm{U}(1)_Y$ gauge symmetries, respectively. Then there are the nine charged fermion masses (or, equivalently, their couplings to the Higgs scalar). Also, there are the three angles and one phase of the Cabibbo-Kobayashi-Maskawa quark mixing matrix. Then there is the Higgs mass and its vacuum expectation value. Finally, there is the QCD parameter θ, a measure of the degree of CP violation in the strong interaction.[4] It should be noted that whereas the Standard Model has been formulated as a theory with massless neutrinos, recently evidence has accumulated for neutrino mass. This can easily be accommodated in the model, though at the cost of at least seven new parameters, three masses and four mixing angles. Because there are so many free parameters that cannot be calculated from first principles, and because it leaves so many fundamental questions unanswered, ultimately the Standard Model leaves one unsatisfied.

Perhaps the most glaring problem with the Standard Model is that it is not truly unified, as one might hope a final theory would be. There are three separate gauge couplings, each of whose values have been precisely measured by experiment and found to be completely different. In so-called grand unified theories (GUT's), the three gauge couplings are assumed to unify at some very high energy scale $m_X \sim 10^{15}$ GeV, resulting in a single gauge coupling of the larger symmetry group, in which the Standard Model is embedded [23]. Below m_X, the theory undergoes spontaneous symmetry breaking, ultimately resulting in the $\mathrm{SU}(3)_C$, $\mathrm{SU}(2)_L$ and $\mathrm{U}(1)_Y$ symmetries of the Standard Model. Besides unification, GUT's also attempt to provide natural explanations for many of the questions posed above by employing general symmetry arguments. The simplest such theory is the $SU(5) \rightarrow \mathrm{SU}(3)\times\mathrm{SU}(2)\times\mathrm{U}(1)$ model proposed by Georgi and Glashow in 1974 [24]. However, this theory, while aesthetically pleasing due to its simplicity, suffers from three major defects. The first is that it predicts the proton to decay with a characteristic lifetime much shorter than is experimentally allowed. The second is that the three gauge couplings, when extrapolated to the supposed unification scale m_X using the renormalization group equations, actually fail to meet at a single point, indicating that unification is not attained. This is illustrated in Figure 1.1(a), which shows the Standard Model prediction for the evolution of the gauge couplings, assuming no new physics between $\mu \sim m_t$ and m_X. The third is

[4] That θ is very nearly zero, even though in general one would not expect it to be, is one of the more severe "fine-tuning" problems in the Standard Model.

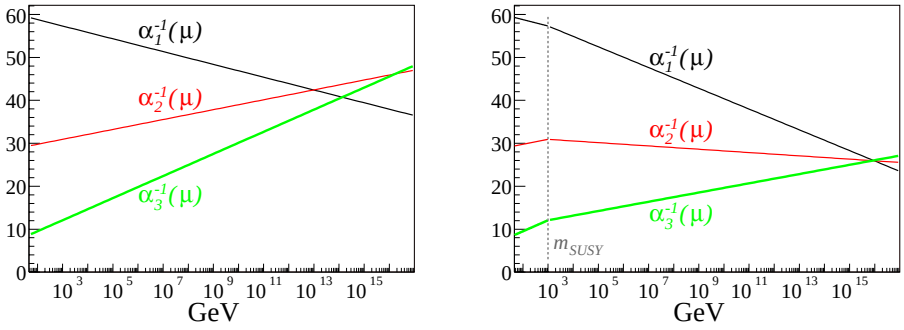

Figure 1.1: Gauge coupling evolutions in the Standard Model and in the minimal supersymmetric extension to the Standard Model. Note that $\alpha_i = g_i^2/4\pi$. Based purely on the number of particles contained in the model, the renormalization group equations can be used to extrapolate $\alpha_i(\mu)$ to any energy scale μ given its value at any other energy. Figure (a) shows how α_i^{-1} evolve assuming the Standard Model is complete (i.e., no new physics above the electroweak breaking scale). In Figure (b), minimal supersymmetry has been assumed above ~1 TeV. Adapted from Reference [22].

that it introduces a severe gauge hierarchy problem. That is, why is the scale of the unified symmetry breaking at $\sim 10^{15}$ GeV so vastly different from that of the electroweak symmetry breaking at ~250 GeV? This would imply extremely large and seemingly coincidental cancellations between the Higgs' bare mass and its radiative corrections [25].

Refusing to give up on the beautiful idea of ultimate unification, some have sought to rectify the situation by appealing to even higher symmetries in nature. This has given rise to GUT's involving more complicated symmetry breaking patterns, such as SO(10) and the superstring-inspired E_6. Most notably, however, it has long been observed that by adding supersymmetry to the simple SU(5) theory, in which spin-1/2 fermions have spin-0 boson partners and spin-1 or spin-0 bosons have spin-1/2 fermion partners, grand unification can be restored [26]. This is shown in Figure 1.1(b). Above the spontaneous supersymmetry breaking scale m_{SUSY}, new particles appear and modify the renormalization group equations, changing the evolution of the gauge couplings. This has three intriguing consequences. The first is that the unification scale m_X is raised to $\sim 10^{16}$ GeV, which is even closer to the Planck mass $m_P \approx 10^{19}$ GeV, perhaps hinting at a role played by gravity. The second is that the proton decay lifetime is raised to $\sim 10^{35}$ years, safely beyond the experimental limits. Finally, supersymmetry can provide an "explanation" for the gauge hierarchy problem, as a Higgs boson mass term would violate chiral supersymmetry transformations. The spontaneous breakdown of supersymmetry then allows the Higgs particle to acquire a mass $m_H \approx m_{SUSY}$. It should be noted that models with $m_{SUSY} \approx 1$ TeV, as assumed in

Figure 1.1(b), are fully consistent with precision electroweak measurements, though somewhat higher values for m_{SUSY} are also allowed. If m_{SUSY} gets much higher than 10 TeV, however, not only does the consistency with the precision electroweak data start to suffer, but the hierarchy problem is reintroduced.

There are thus theoretical reasons to expect new physics at the TeV scale. The physics program should therefore proceed as it has for decades, with the various colliders exploring the high-energy frontier and precision electroweak experiments testing the internal consistency of the model. As has happened in the past, it may even be that a high-precision but low-energy experiment will be able to reveal signs of new physics before it is eventually discovered at a high-energy facility. Accordingly, the rest of this chapter motivates a precision measurement of the parity-violating cross-section asymmetry in electron–electron scattering, describing the specific types of new physics signatures it might expect to observe. In general, such a measurement is very sensitive to the helicity structure of new physics models, and thus provides information complementary to that obtained from current high-energy colliders. Chapter 2 details the experimental apparatus and methodology. Chapter 3 gives a full description of the analysis procedure, with the results being presented in Chapter 4. Finally, Chapter 5 discusses the findings in terms of the implications they have for various possible extensions to the Standard Model.

1.2 Parity Violation in Møller Scattering

The tree level diagrams representing the processes governing electron–electron (Møller) scattering are given in Figure 1.2. Since Møller scattering is a neutral current interaction, the only forces through which the process can occur are the electromagnetic, mediated by the massless photon, and the neutral weak force, mediated by the massive Z. The helicity-dependent cross section can therefore be written as:

$$\sigma_{\pm} = |\mathcal{M}_{\gamma} + \mathcal{M}_{Z\pm}|^2 \tag{1.7}$$

Here $\mathcal{M}_{\gamma}$ and $\mathcal{M}_{Z}$ are the matrix amplitudes for the electromagnetic and neutral weak processes, respectively, and the $\pm$ denotes the helicity of the beam electron (positive for right-helicity, negative for left-helicity). The target electron is unpolarized, so that any

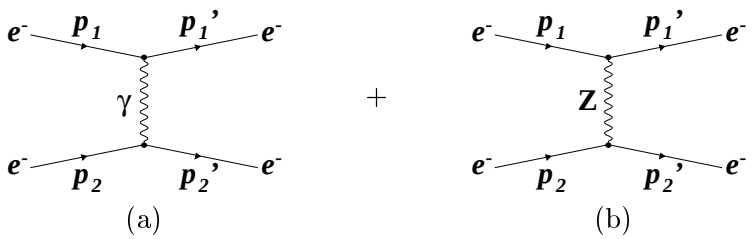

Figure 1.2: Neutral current amplitudes contributing to the parity-violating cross section asymmetry $A_{\rm PV}$ in Møller scattering at tree level.

helicity-dependence in the cross section can only come from parity-violating terms in the neutral currents, which themselves arise from the simultaneous existence of both vector and axial-vector current components. However, the electromagnetic force conserves parity, so that its matrix amplitude does not depend on the helicity of the beam electron. In contrast, the weak force is known to violate parity, so that its matrix amplitude *does* depend on the helicity of the beam electron. At low energies ($Q^2 \ll m_Z^2$), the electromagnetic amplitude, which goes as α/Q^2, completely dominates over the neutral weak amplitude, which contains a $1/m_Z^2$ suppression factor because of the massive propagator. Therefore, the ratio of the Z amplitude to the electromagnetic amplitude should be roughly Q^2/m_Z^2, which is on the order of a few parts per million (ppm).

Going further, one can use this extreme disparity in the relative strengths of the two amplitudes at low Q^2 to estimate the size of the accompanying parity-violating asymmetry $A_{\rm PV}$ in Møller scattering. Defining the asymmetry as $A_{\rm PV} \equiv (\sigma_+ - \sigma_-)/(\sigma_+ + \sigma_-)$, one finds that the asymmetry is due to the interference between the electromagnetic and weak neutral currents:

$$\begin{aligned} A_{\rm PV} &= \frac{|\mathcal{M}_\gamma + \mathcal{M}_{Z+}|^2 - |\mathcal{M}_\gamma + \mathcal{M}_{Z-}|^2}{|\mathcal{M}_\gamma + \mathcal{M}_{Z+}|^2 + |\mathcal{M}_\gamma + \mathcal{M}_{Z-}|^2} \\ &\approx \frac{2{\rm Re}\left\{\mathcal{M}_\gamma \left(\mathcal{M}_{Z+} - \mathcal{M}_{Z-}\right)^*\right\}}{|\mathcal{M}_\gamma|^2} \end{aligned} \tag{1.8}$$

Here the fact that $\mathcal{M}_Z \ll \mathcal{M}_\gamma$, appropriate for $Q^2 \ll m_Z^2$, has been used. This leads to an order-of-magnitude estimate for $A_{\rm PV}$ similar to that obtained for the relative sizes of the amplitudes themselves, namely something on the order of a few ppm. A more precise

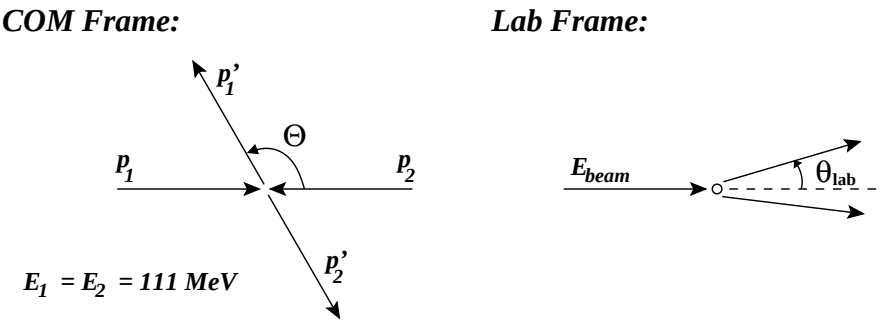

Figure 1.3: Scattering kinematics in center-of-mass frame and laboratory frame. For an incident beam energy of $E_{beam} = 48$ GeV, the energy of each electron in the center-of-mass frame is approximately 111 MeV.

calculation yields the following result [27]:

$$A_{\rm PV} = -\frac{2\sqrt{2}G_F Q^2}{\pi\alpha}\frac{1+\cos\Theta}{(3+\cos^2\Theta)^2}(1-4\sin^2\theta_W) \tag{1.9}$$

where $G_F \approx 1.166 \times 10^{-5}$ GeV^{-2} is the Fermi coupling constant, determined from the muon lifetime; $\alpha \approx 1/137.036$ is the fine structure constant, determined by Thomson scattering at $Q^2 = 0$; $Q^2 = (p-p')^2$ is the square of the transfer four-momentum; Θ is the scattering angle in the center-of-momentum frame; and $\sin^2\theta_W \approx 0.2311$ is the square of the sine of the weak mixing angle, more of which will be said below. The small size of the asymmetry can now be seen to be due primarily to the tiny factor of $G_F Q^2$ in the numerator, but also in part to the $1 - 4\sin^2\theta_W$ suppression factor. Finally, the following relation will often prove useful:

$$Q^2 = \frac{s}{2}\cdot(1-\cos\Theta) = s\cdot\sin^2(\Theta/2) \tag{1.10}$$

where $s = 2m_e^2 + 2m_e E_{beam}$ is the square of the total energy in the center-of-momentum frame, with m_e the electron mass and E_{beam} the beam energy.

The kinematics for the scattering process are illustrated in Figure 1.3. For a given incident luminosity, it is appropriate to consider which set of kinematic parameters (namely, E_{beam} and Θ) minimizes the the relative statistical uncertainty $\delta(A_{\rm PV})/A_{\rm PV}$ achievable by the experiment. The (spin-averaged) differential cross section for Møller scattering is given by the following formula [28]:

$$\frac{d\sigma}{d\Omega} = \frac{\alpha^2}{2m_e E_{beam}}\frac{(3+\cos^2\Theta)^2}{\sin^4\Theta} \tag{1.11}$$

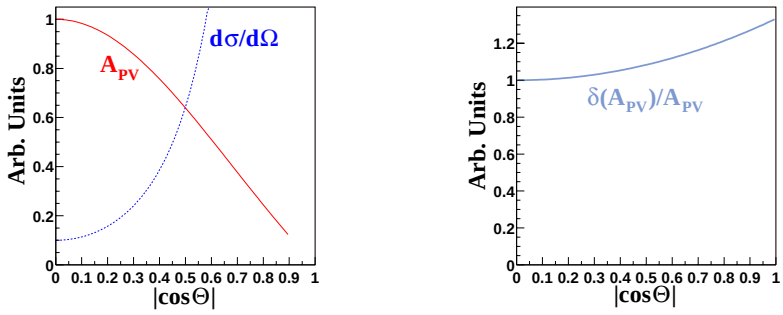

(a) (b)

Figure 1.4: Figure (a) shows the differential cross section $d\sigma/d\Omega$ and parity-violating asymmetry $A_{\rm PV}$ (both in arbitrary units) as a function of $|\cos\Theta|$. Figure (b) shows the statistical uncertainty that can be achieved for a given luminosity, proportional to $\left(A_{\rm PV}\sqrt{d\sigma/d\Omega}\right)^{-1}$, as a function of $|\cos\Theta|$. Small values of $|\cos\Theta|$ are seen to be optimal.

Combining this result with that of Equation (1.9), and substituting the variable Θ for the variable Q^2 using Equation (1.10), the relative statistical uncertainty may be expressed in terms of E_{beam} and Θ as:

$$\begin{aligned} \frac{\delta(A_{\rm PV})}{A_{\rm PV}} &\propto \frac{1}{\sqrt{d\sigma/d\Omega}} \cdot \frac{1}{A_{\rm PV}} \\ &\propto \frac{1}{\sqrt{E_{beam}}}(3+\cos^2\Theta) \end{aligned} \tag{1.12}$$

From this equation, one can see that maximizing E corresponds to minimizing the relative statistical uncertainty. Likewise, the function $3+\cos^2\Theta$, plotted in Figure 1.4(b), shows a slow variation with respect to Θ, with 90° scattering in the center-of-momentum frame ($|\cos\Theta| = 0$) being optimal.

The experiment uses the electron beam at the Stanford Linear Accelerator Center (SLAC), which can generate electrons of energy up to 48 GeV. The experiment is designed to integrate scattered events in the kinematic range $\Theta = 90$ to 119 degrees, corresponding to $E' = 13$ to 24 GeV, $\theta_{lab} = 4.4$ to 7.5 mrad. This range of energies is wide enough to allow for as much signal acceptance as possible, while just narrow enough to avoid double-counting problems (since Møller electrons come in pairs with $E'_1 + E'_2 = E_{beam}$). It also avoids the significant low-energy, wide-angle backgrounds. It should be noted that this kinematic range corresponds to $Q^2 = 0.023$ to 0.036 GeV^2. The relationships between the various kinematic parameters quoted here are illustrated in Figure 1.5.

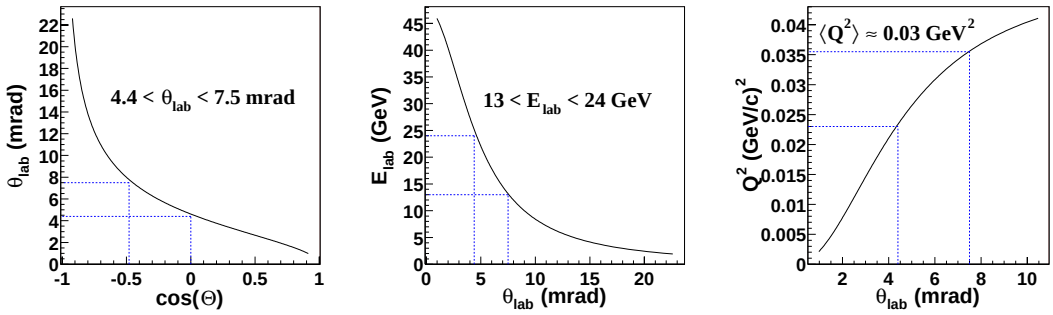

(a) (b) (c)

Figure 1.5: Transformations between various kinematics parameters. Figure (a) plots θ_{lab} vs. $\cos\Theta$. Figure (b) plots E_{lab} (the energy of the scattered electron) vs. θ_{lab}. Figure (c) plots Q^2 vs. θ_{lab}. These plots are useful for getting a feel for the kinematics. In general, wide angles in the lab frame correspond to backwards scattering in the center-of-mass frame, lower scattering energies, and higher momentum transfers.

1.2.1 Radiative Corrections and New Physics Sensitivity

Properly averaged over the experimental kinematics, the value of the tree-level asymmetry is approximately $A_{\rm PV} \approx -270$ parts per billion (ppb). The asymmetry gets significantly modified, however, by one-loop electroweak radiative corrections [29]. When discussing radiative corrections, it is necessary to pick a set of renormalized parameters consistent with a particular renormalization scheme. Here the modified minimal subtraction ($\overline{\rm MS}$) scheme will be employed. In general, the definitions given earlier for α and G_F still suffice, absorbing, for instance, many of the one-loop corrections to muon decay [30]. However, for some of the less important one-loop effects (specifically, box diagrams involving two massive bosons), $\alpha(m_Z) = 1/127.9$, defined to be the Z-pole value of the fine structure constant in the $\overline{\rm MS}$ renormalization scheme, is a necessary substitute. Finally, the following definition for the weak mixing angle will be used exclusively [18]:

$$\sin^2\theta_W \equiv \sin^2\theta_W(m_Z)_{\overline{\rm MS}} = 0.23113 \pm 0.00015 \tag{1.13}$$

This is the sine-squared of the Z-pole value of the weak mixing angle in the $\overline{\rm MS}$ renormalization scheme, related to the "effective" definition of the weak mixing angle used at e^+e^- colliders at CERN and SLAC by a simple translation, $\sin^2\theta_W(m_Z)_{\overline{\rm MS}} = \sin^2\theta_W^{\rm eff} - 0.0003$ [31].

The diagrams representing the most significant one-loop electroweak radiative corrections are shown in Figure 1.6. These include γ-Z mixing as well as the W loop contribution to the anapole moment. In the γ-Z mixing diagrams, mixing occurs through vacuum po-

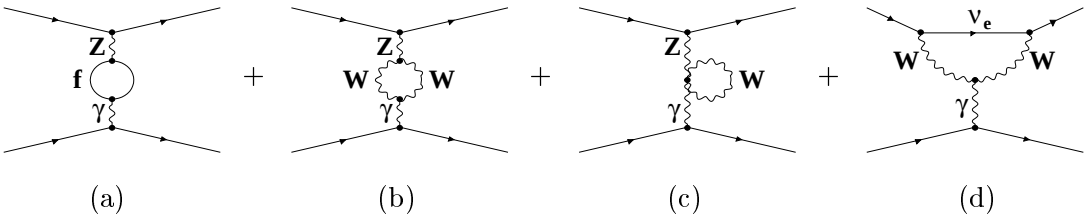

Figure 1.6: Most signficant one-loop electroweak radiative corrections to the Møller scattering asymmetry. Figures (a)–(c) show the vacuum polarization loops contributing to the γ-Z mixing diagrams. Figure (d) shows the W-loop contribution to the anapole moment.

larization loops involving either W bosons or fermions. In the fermionic loops, all leptons and quarks with mass less than m_Z contribute.[5] When evaluating the effects of these diagrams on the low-Q^2 experimental asymmetry $A_{\rm PV}$, one may use $Q^2 \approx 0$ with negligible error [29]. The general contributions for non-zero Q^2 will be discussed later. The low-Q^2 hadronic contributions, however, cannot be handled perturbatively, since at low energies the strong coupling constant diverges. Rather, their contributions must be extracted from $e^+e^- \to$ hadrons data via a dispersion relation [33, 34].

The overall effect of one-loop radiative corrections for $Q^2 \approx 0$ is to reduce the size of the asymmetry by roughly 40%, from -270 to -150 ppb [29]. Because the Standard Model is a renormalizeable quantum field theory, potential "new physics" at higher energies can introduce one-loop effects that modify $A_{\rm PV}$ in a manner that is entirely analogous to the way in which the standard electroweak radiative corrections do. For instance, "oblique" corrections introduce their own vacuum polarization loops in the boson propagators, similar to the γ-Z mixing diagrams shown in Figure 1.6. These modify the self-energies of the photon, W, and Z, as well as the γ-Z mixing (parameterized by the vacuum polarization functions $\Pi_{\gamma\gamma}$, Π_{WW}, Π_{ZZ}, and $\Pi_{\gamma Z}$, respectively). Such corrections, caused, for example, by a new generation of heavy fermions ($m_f \gg m_Z$) found in many supersymmetric models and possessing highly suppressed couplings to the light fermions e, μ, τ, u, d, s, and c, can be described by three new parameters S, T, and U [35]. Data from LEP at CERN and the SLC at SLAC heavily constrain these new parameters [36]. However, if the new heavy fermions instead have masses not much larger than the electroweak scale ($m_f \sim \mathcal{O}(100)$ GeV), then

[5] According to the Marciano–Rosner convention adopted by many of the theory references [31, 32], the top quark's contribution is absorbed into the definition of $\sin^2\theta_W$.

three additional parameters V, W, and X need to be introduced [37]. The X parameter, in particular, is of interest, for it parameterizes the contributions of electroweak scale oblique corrections to the off-shell value of the weak mixing angle, i.e.:

$$\frac{\sin^2\theta_W(0)}{\sin^2\theta_W(m_Z)} \simeq 1.032 - 0.033X \tag{1.14}$$

The current constraint on X is not nearly as tight as those on S and T, which are provided by the Z-pole measurements at LEP and SLC, but roughly equal to the present constraint on U set by the W mass measurement [36]. An off-shell measurement of the weak mixing angle at the ± 0.001 level could actually improve the current constraint on X by a factor of two or three.

So far only new phenomena that affect low-energy observables via radiative corrections (and then only through "oblique" corrections) have been considered. New physics that would contribute to the Møller scattering process through non-radiative means can also be probed by a measurement of $A_{\rm PV}$. For instance, many extensions to the Standard Model, including many grand unified theories and models involving supersymmetry and/or extra dimensions, propose the existence of new massive neutral gauge bosons, collectively referred to as Z' bosons [18, 38, 39]. In order to explain the gauge hierarchy problem satisfactorily, some models require such bosons to exist at the TeV scale. The current limits on the S and T parameters discussed above require the mixing between the regular Z boson and any new Z' bosons to be extremely small. Nevertheless, a measurement of $A_{\rm PV}$ provides sensitivity to the possible existence of Z' bosons, provided there are parity-violating terms in their interactions [34]. The limit on $m_{Z'}$ that could be placed by a $\sim$10% measurement of $A_{\rm PV}$ varies depending on the model, but for typical models (such as SO(10) and E_6) energies of 600 to 900 GeV are explored. This is comparable to (if slightly less than) the discovery reach equivalent of one or two years of Run II data at the Tevatron [40]. However, a low-energy Møller asymmetry measurement complements the direct Z' searches of the Tevatron and other colliders in the following important ways. First, should a deviation from the Standard Model prediction for $A_{\rm PV}$ be observed, the sign of the relative shift would indicate whether the supposed Z' (if indeed a Z' is responsible for the shift) couples more strongly to right- or left-handed electrons. Second, the size of the shift would depend strongly upon the details of the particular model being considered, which could be very useful in discriminating between

competing theories, should a collider yield a future discovery.

Finally, lepton compositeness has long been postulated as a possible extension to the Standard Model [41]. In such models, quarks and leptons are comprised of constituents held together by forces whose interactions become important at the scale of the binding energies, parameterized by the quantity Λ_{ff} (where f denotes the fermion type). Well below the compositeness scale Λ_{ee}, the general four-electron contact interaction takes the following form, assuming that both helicity and flavor are conserved [41]:

$$\mathcal{L}_{ee} = \frac{4\pi}{2\Lambda_{ee}^2}\left[\eta_{LL}(\overline{\psi}_{_L}\gamma_\mu\psi_{_L})^2 + \eta_{RR}(\overline{\psi}_{_R}\gamma_\mu\psi_{_R})^2 + 2\eta_{LR}(\overline{\psi}_{_L}\gamma_\mu\psi_{_L})(\overline{\psi}_{_R}\gamma^\mu\psi_{_R})\right] \tag{1.15}$$

Provided that there are parity-violating terms (i.e., $\eta_{RR} \neq \eta_{LL}$), a measurement of $A_{\rm PV}$ possesses great sensitivity to the possibility of electron substructure, probing Λ_{ee} at a level approaching 10 TeV. This is comparable to the current limits from e^+e^- colliders [42]. However, even when in the future the limits are pushed to the $10 - 20$ TeV range and above, a measurement of $A_{\rm PV}$ could provide important insight into the precise nature of the new interactions, should they exist.

1.2.2 Weak Mixing Angle

In order to facilitate comparisons to various theoretical predictions as well as to other experimental results, it is convenient to cast the measurement of the parity-violating Møller asymmetry into a measurement of the weak mixing angle. This is accomplished through a straightforward rearrangement of Equation (1.9), discussed further in Chapter 5. Because of the electroweak radiative corrections discussed in the last section, however, the weak mixing angle effectively becomes a function of energy. The following substitution for $\sin^2\theta_W$ is therefore made in the tree-level expression for $A_{\rm PV}$ [43]:

$$\sin^2\theta_W \;\longrightarrow\; \sin^2\theta_W(Q) = [1 + \Delta r_\ell(Q) + \Delta r_q(Q) + \Delta r_b(Q)]\sin^2\theta_W(m_Z)_{\overline{\rm MS}} \tag{1.16}$$

where Δr_ℓ, Δr_q, and Δr_b are the radiative contributions coming from the lepton, quark, and W boson loops, respectively, in the diagrams depicted in Figure 1.6, namely the γ-Z mixing diagrams and the W loop contribution to the anapole moment. These are the *only* radiative corrections described by the functions Δr_ℓ, Δr_q, and Δr_b.

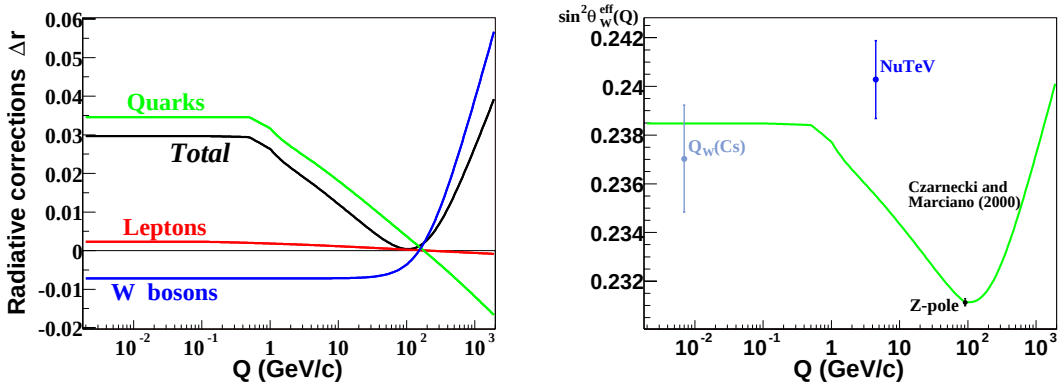

Figure 1.7: Electroweak radiative corrections to Møller scattering plotted as a function of energy. Figure (a) shows the contributions of lepton, quark, and W boson loops separately, but only for the diagrams contained in Figure 1.6. The quark loop contribution can seen to be capped at low-Q^2. Figure (b) shows the effective running of $\sin^2\theta_W$, along with some of its measurements.

The functions Δr_ℓ, Δr_q, and Δr_b are plotted versus Q in Figure 1.7(a). It is important to recognize that these functions do not parameterize all radiative corrections modifying $\sin^2\theta_W$ (and, hence, $A_{\rm PV}$) at arbitrary energies. In particular, at large Q^2, appropriate for future e^+e^- or e^-e^- colliders, the effects of WW box diagrams must be taken into account [43]. Nevertheless, Figure 1.7(a) gives a good general impression of the relative contributions of each higher-order process over a wide range of energies. It is particularly relevant for low energies, where the total radiative corrections are dominated by the diagrams described by the functions being plotted. One can readily see the "cap" placed on the low-energy quark contributions, necessary since low-energy QCD is nonperturbative. At the experimental kinematics, $Q \approx 0.2$ GeV, $\sin^2\theta_W$ has risen by roughly 3% from its Z-pole value, which corresponds to a decrease in $A_{\rm PV}$ of roughly 37% from its tree-level value. Other radiative corrections, not parameterized by the functions Δr_ℓ, Δr_q, and Δr_b, reduce $A_{\rm PV}$ by a further 3%, resulting in a total reduction of 40% as reported earlier.

In Figure 1.7(b), Equation (1.16) is plotted as a function of Q. Various measurements of the weak mixing angle are also shown on the plot. These include the $Q = m_Z$ measurements at LEP and SLC [42], the $Q \approx 4$ GeV measurement by the NuTeV collaboration [44, 45], and the $Q \approx 0$ measurement by the cesium atomic parity violation experiment [46]. These measurements all suffer from one or more drawbacks, however. In the case of the Z-pole measurements, that drawback is limited sensitivity to new physics which does not couple strongly to the Z. For on-shell Bhabha scattering, note that the matrix amplitude $\mathcal{M}_Z$ of

the Z-mediated process is large and purely imaginary, whereas the matrix amplitude $\mathcal{M}_X$ of a new physical process involving an undiscovered massive particle with $m_X \gg m_Z$ (to which the Z couples only weakly) would be purely real and relatively small. Hence, the total cross section would go as:

$$\begin{aligned} \left.\frac{d\sigma}{d\Omega}\right|_{Q=m_Z} &\propto \left|\mathcal{M}_Z + \mathcal{M}_X\right|^2 \sim \left|\mathcal{M}_Z\right|^2 + \left|\mathcal{M}_X\right|^2 \\ &\sim \left|\mathcal{M}_Z\right|^2 \left(1 + \frac{\left|\mathcal{M}_X\right|^2}{\left|\mathcal{M}_Z\right|^2}\right) \end{aligned} \tag{1.17}$$

No interference between the two processes occurs, and, since $\left|\mathcal{M}_X\right|^2 \ll \left|\mathcal{M}_Z\right|^2$ for $Q^2 \approx m_Z^2$, the total cross section is basically the same as if the new interaction was not contributing at all. Hence, the sensitivity of the Z-pole measurements to things like lepton compositeness, exotic Higgs particles, and Z' bosons that do not couple strongly to the Z is highly suppressed. Low-energy measurements offer much greater sensitivity.

One such low-energy measurement is provided by the NuTeV experiment at FNAL, which measures the ratios of neutral current to charged current inclusive cross sections in the deep inelastic scattering of neutrino and antineutrino beams from quarks in heavy nuclei. Using the so-called Paschos-Wolfenstein relation [47], these ratios can be used to extract a value for the weak mixing angle:

$$\frac{\sigma(\nu_\mu N \to \nu_\mu X) - \sigma(\overline{\nu}_\mu N \to \overline{\nu}_\mu X)}{\sigma(\nu_\mu N \to \mu^- X) - \sigma(\overline{\nu}_\mu N \to \mu^+ X)} = \frac{R^\nu - r \cdot R^{\overline{\nu}}}{1 - r} = \frac{1}{2} - \sin^2\theta_W \tag{1.18}$$

where

$$R^\nu = \frac{\sigma(\nu_\mu N \to \nu_\mu X)}{\sigma(\nu_\mu N \to \mu^- X)} \qquad R^{\overline{\nu}} = \frac{\sigma(\overline{\nu}_\mu N \to \overline{\nu}_\mu X)}{\sigma(\overline{\nu}_\mu N \to \mu^+ X)} \qquad r = \frac{\sigma(\overline{\nu}_\mu N \to \mu^+ X)}{\sigma(\nu_\mu N \to \mu^- X)} \approx \frac{1}{2} \tag{1.19}$$

While this measurement does a better job of providing sensitivity to many classes of potential new physics (such as new "oblique" radiative corrections or Z' bosons, in addition to quark compositeness), it suffers from theoretical uncertainties surrounding the complicated hadronic interactions that must be included in its analysis. Use of the Paschos-Wolfenstein relation in Equation (1.18) suppresses sensitivity to the considerable theoretical uncertainties associated with charm quark production from scattering off low-momentum sea quarks, which in the past has severely limited extraction of process-independent observables (such

as the weak mixing angle) from the results of neutrino-nucleon scattering experiments. However, there are still large uncertainties associated with the scattering process's electroweak and pure QED radiative corrections, calculated by Reference [48]. It has been suggested that effects primarily related to the calculation's treatment of initial and final state radiation from quarks and the use of experimental energy cuts can significantly modify the interpreted value of the weak mixing angle, perhaps by as much as the total quoted experimental uncertainty [49]. In addition, an independent study of the QCD corrections relevant for the NuTeV experiment identified several important sources of uncertainty that remain accounted for in the experiment's published result [50]. For instance, this study found that a small (1%) violation of isospin symmetry in the valence parton distribution functions could by itself reduce the observed discrepancy by a factor of two. These and other related issues may well be responsible for at least part of the $3\,\sigma$ deviation of the NuTeV experiment's result from the Standard Model prediction, without needing to invoke a "new physics" explanation.

Proceeding further to even lower energies, the atomic parity violation group based at the University of Colorado have reported on a measurement of the weak charge Q_W of the nucleus of ^{133}Cs, which can be related to $\sin^2\theta_W$ in the context of the Standard Model [46]. The measurement relies on the observation of the relative rates of the $6S_{F=3} \rightarrow 7S_{F=4}$ and $6S_{F=4} \rightarrow 7S_{F=3}$ parity-violating hyperfine transitions in a spin-polarized atomic cesium beam. The error bar on the measurement, however, is relatively large and is completely dominated by systematics and theoretical uncertainties, chiefly those associated with the proper computation of the atomic wavefunction [51]. In order to appreciate the magnitude of these uncertainties, it is instructive to consider a sampling (not comprehensive, but fairly representative of the literature) of the published results interpreting Q_W from the measurements of Reference [46]. Such a sampling is pictured in Figure 1.8. First, in 1999, Bennett and Wieman updated the two-year-old analysis of Reference [46] using newly available atomic structure data, resulting in a 2.5σ deviation of Q_W from the Standard Model prediction [52]. In 2000, Derevianko included the Breit interaction in the wavefunction calculations, while also using a new nuclear charge distribution, bringing the Q_W result to within 1.0σ of the Standard Model prediction [53]. In 2001, Derevianko further refined his treatment of the Breit interaction, employing relativistic many-body perturbation theory and further reducing the Q_W discrepancy to 0.6σ [54]. In 2002, however, Dzuba, Flambaum,

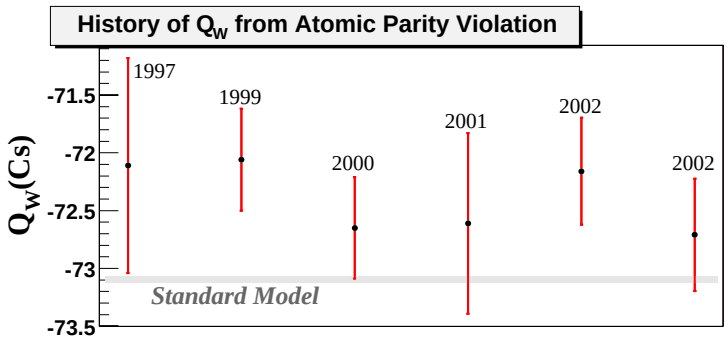

Figure 1.8: Time-history of the theoretical corrections to the $Q_W(^{133}\mathrm{Cs})$ measurement of the Boulder atomic parity violation group. Each data point corresponds to an updated analysis. The error bars in each case are dominated by systematics and theoretical uncertainties. Significant deviations from the Standard Model prediction have been reported over the years, but the mean value has not yet settled, due to persistent theoretical difficulties.

and Ginges improved their methods for handling Coulomb screening effects, as well as improved the overall numerical precision of their calculations. The result was a 2.0σ deviation from the Standard Model prediction, restoring the former discrepancy [55]. Around the same time, Sushkov showed [56] that vacuum polarization effects in the electronic orbitals were capable of inducing corrections as large as those found earlier by Derevianko, which were due to the Breit interaction. Thus, in late 2002, Kuchiev and Flambaum reported on a calculation that included these QED radiative corrections and brought Q_W once again into agreement with the Standard Model prediction, this time to within 0.8σ [57]. Theoretical work continues, however, with Derevianko suggesting that a fully self-consistent many-body treatment of vacuum polarization effects has yet to be performed and could change the value of Q_W significantly [58].

The point of this discussion is not to reflect poorly on the results of the atomic parity violation group, which after all has accomplished a significant feat, namely the first-ever measurement of a nuclear anapole moment, whose existence was originally predicted by Zel'dovich in 1958 [59]. Rather, the point is merely to emphasize that interpreting the group's measurements of the parity-violating $6S \to 7S$ transition rates in cesium in terms of fundamental physics parameters, such as Q_W or $\sin^2\theta_W$, remains a major theoretical challenge. The original work quoted $Q_W = -72.11 \pm 0.27 \pm 0.89$ [46], compared to the Standard Model prediction of $Q_W = -73.10 \pm 0.03$ [18]. Since then, the theoretical uncertainty has been reduced and the central value has shifted to $Q_W = -72.71 \pm 0.29 \pm 0.39$ [57]. From the discussion above, however, it is quite clear that the theoretical uncertainty has mostly

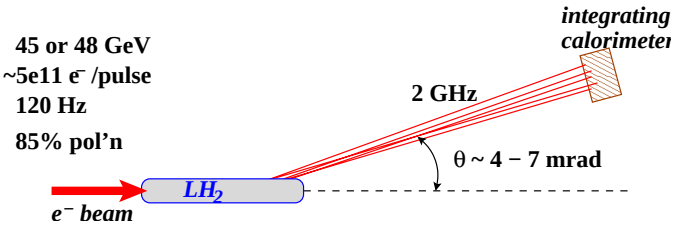

Figure 1.9: Conceptual design for the experiment. An intense, high-energy, polarized electron beam impinges on a liquid hydrogen target, and the $e^-e^- \longrightarrow e^-e^-$ (Møller) scatters are focused into an integrating calorimeter for detection. By rapidly flipping the helicity of the incident electron beam, the parity-violating scattering asymmetry $A_{\rm PV}$ can be measured.

been underestimated.

The above discussion also serves to highlight the principle motivation for the measurement of the parity-violating asymmetry $A_{\rm PV}$ in Møller scattering. The interaction is purely leptonic, and therefore the electroweak radiative corrections are conceptually simple and have relatively small uncertainties. There are few hadronic effects to consider and no atomic wavefunctions to compute. Interpreting the Møller asymmetry $A_{\rm PV}$ in terms of the fundamental Standard Model parameter $\sin^2\theta_W$ is a relatively straightforward procedure requiring comparatively few corrections and assumptions. Measuring this asymmetry is therefore a very attractive means for obtaining a precision low-energy measurement of the weak mixing angle, thereby testing the Standard Model at the quantum loop level and continuing the search for new physics.

1.3 Experimental Goals and Requirements

The experimental goal is to measure the parity-violating Møller asymmetry $A_{\rm PV}$, predicted to be approximately -150 ppb (before accounting for initial and final-state radiation effects), to a precision of roughly 10%. This will allow for a nearly ± 0.001 determination of $\sin^2\theta_W$. The basic experimental design concept is shown in Figure 1.9. A beam of high-energy, longitudinally polarized electrons impinges on a liquid hydrogen target. Some of the electrons pass through unimpeded and proceed to the beam dump, while others scatter off the target electrons or protons. Electron–proton (ep) scatters will have a different energy–angle correlation than the electron–electron (Møller) scatters. A system of magnets and collimators can therefore be used as a spectrometer, purifying the signal (i.e., reducing the backgrounds discussed below) for integration in an electromagnetic calorimeter. The

polarized electron source at SLAC is able to deliver as many as 5×10^{11} electrons in a single beam pulse. At 120 Hz, and for a liquid hydrogen target, this corresponds to a very high luminosity $\mathcal{L} \approx 4 \times 10^{38}$ cm^{-2}s^{-1}. At the scattering angles being considered, which have already been chosen so as to maximize the experiment's statistical power, the total cross section is approximately 10 μbarn, making the detector rate roughly 2 GHz.

The helicity of the electron beam can be manipulated between right and left states at 120 Hz, forming pulse-pairs of opposite helicity. The scattering rates for the two different beam helicities can then be compared by constructing the following pulse-pair asymmetry:

$$A_{\mathrm{PV}} = \frac{\sigma_R - \sigma_L}{\sigma_R + \sigma_L} \tag{1.20}$$

where σ_R and σ_L are the scattering rates for the right and left helicity states, respectively. The scattering rates are obtained by dividing the detector signal by the incident beam intensity, as is described in much more detail in Chapter 3. Since each pulse is expected to deliver approximately 20 million electrons into the detector, the statistical resolution of the detector should be as low as $\left(\sqrt{2 \times 10^7} \cdot \sqrt{2}\right)^{-1} \approx 160$ ppm per pair. This implies that an uncertainty on the asymmetry of roughly 15 ppb, corresponding to a $\sim$10% measurement of A_{PV}, should be achievable after just over 100 million pairs. Of course, systematic uncertainties will also contribute to the total uncertainty $\delta(A_{\mathrm{PV}})$.

As will be seen in Chapters 4 and 5, the primary systematic uncertainties contributing to $\delta(A_{\mathrm{PV}})$ arise from beam effects, physics backgrounds, and asymmetry normalizations. In addition, theoretical uncertainties associated with the calculation of radiative corrections (both electroweak and purely electromagnetic) contribute to $\delta(\sin^2\theta_W)$, but these are relatively small and will not be discussed until Chapter 5. Asymmetry normalizations include scale factors such as the beam polarization P_{beam} and the detector linearity ϵ. Both modify the measured asymmetry in the following way: $A_{meas} = A_{real} \cdot P_{beam} \cdot \epsilon$. That is, the measured asymmetry will at most be equal to, and in general smaller than, the real physical asymmetry. Since these quantities enter so directly into the interpretation of A_{real} from A_{meas}, it is important to keep their relative uncertainties as small as possible. The beam polarization uncertainty should therefore be kept below 5%, and the detector should be linear to within 1%.

Because the scattering rate is dependent on the incident beam energy and on the scat-

tering angle, as can be seen from the cross-section formula in Equation (1.11), helicity-correlated differences in either the beam energy or in the beam's trajectory will in general produce false asymmetries via purely geometric effects. For example, a right-helicity beam hitting the target at a position that is 10 nm different from that of the corresponding left-helicity beam will have to scatter by a slightly different angle in order to make it into the detector's acceptance, resulting in a false asymmetry of, say, 10 ppb. The same logic applies to all of the beam parameters E, x, y, x', and y', where x' and y' are the incident horizontal and vertical beam angles, respectively. Hence, the right-left (or "helicity-correlated") differences in all of the beam parameters must be carefully monitored throughout the experiment to correct for beam-related false asymmetries.

The asymmetry correction procedure will be discussed in more detail in Chapter 3. For now it suffices merely to motivate the goals for how small the various right-left beam parameter differences should be kept over the course of the experiment. The correlation coefficients parameterizing how sensitive the asymmetry is to each of the various beam parameters are roughly of the following size:

$$\begin{aligned} \frac{\partial A}{\partial E} &\sim 10 \text{ ppb/keV} \\ \frac{\partial A}{\partial x},\ \frac{\partial A}{\partial y} &\sim 1 \text{ ppb/nm} \\ \frac{\partial A}{\partial x'},\ \frac{\partial A}{\partial y'} &\sim 50 \text{ ppb/nrad} \end{aligned} \tag{1.21}$$

If each correction is assumed to be accurate to roughly 10%, then right-left differences of the sizes given below will generate a total contribution of roughly ± 3 ppb to the systematic uncertainty $\delta(A_{\mathrm{PV}})$ [60]:

$$\begin{aligned} \Delta E &< 2 \text{ keV} \\ \Delta x,\ \Delta y &< 10 \text{ nm} \\ \Delta x',\ \Delta y' &< 0.25 \text{ nrad} \end{aligned} \tag{1.22}$$

The charge asymmetry must be handled differently. Since the detector signal is charge-normalized, a large charge asymmetry should in theory have no real effect on the observed asymmetry. However, as will become evident in Section 2.2, a large charge asymmetry hurts the experiment in a number of ways. First and foremost, since charge couples to all beam parameters at some level, a large charge asymmetry can produce helicity correlations in

other beam parameters. In order to ensure that other beam parameters (most importantly energy) do not develop large helicity correlations, one should start by nulling the charge asymmetry. Second, a charge asymmetry can produce a residual false asymmetry due to detector nonlinearities. This can be assuaged by performing an "extra regression" of the measured asymmetry against the charge asymmetry. This detail will be left until Chapter 3 to be discussed further. For now, one can assume a detector nonlinearity of 1% in order to set a limit on the maximum allowable charge asymmetry:

$$\begin{aligned} \sqrt{(3 \text{ ppb})^2 + (0.01 \cdot A_Q)^2} &< 4 \text{ ppb} \\ \Longrightarrow A_Q &\lesssim 200 \text{ ppb} \end{aligned} \tag{1.23}$$

The same limit for A_Q can be obtained by noting that beam loading[6] is estimated to be as large as 10%, meaning that a 200 ppb charge asymmetry will tend to produce a 20 ppb energy difference. For a beam energy of ∼50 GeV, this corresponds to a right-left energy difference of $\Delta E \approx 2$ keV, equal to the limit given in Equation (1.22). The beam asymmetry limits given in Equations (1.22) and (1.23) have been calculated so as to keep the total systematic uncertainty contribution from the asymmetry correction procedure at the ±4 ppb level, which is a reasonable goal.

The per-pair detector resolution σ_{det} is defined to be the width of the detector asymmetry distribution. The final statistical uncertainty on $A_{\rm PV}$ will be governed by this resolution, via the ordinary statistical averaging formula: $\delta(A_{\rm PV})_{stat} = \sigma_{det}/\sqrt{N}$, where N is the total number of integrated pulse-pairs. The detector resolution receives contributions from statistical fluctuations, electronics noise, and beam jitter. Beam jitter refers to the random pulse-to-pulse beam fluctuations that naturally arise in the accelerator. While such random fluctuations may end up averaging to zero when integrated over the course of the entire experiment, they still contribute to the asymmetry width. The detector resolution σ_{det} can therefore be written as:

$$\sigma_{det}^2 = \frac{1}{N_e} + \sigma_{elec}^2 + \sigma_{beam}^2 \tag{1.24}$$

where N_e is the number of electrons reaching the detector per pair ($N_e \approx 40$ million), σ_{elec} is the random electronic noise, and σ_{beam} is the noise due to beam jitter. Carrying out the

[6]Beam loading is a measure of the degree to which the beam energy is coupled to the beam intensity. Its origin is discussed in Section 2.3.

asymmetry correction procedure on a pair-by-pair basis essentially removes σ_{beam} from this equation, drastically reducing the detector resolution. This is of crucial importance, since this produces a smaller statistical uncertainty.

In order to properly account for beam helicity correlations, the beam parameters must be carefully monitored on a pulse-by-pulse basis. All first-order beam parameters E, x, y, x', and y' are monitored, as are some second-order beam parameters. In general, precision in beam monitoring is more important than accuracy. That is, calibration constants such as offsets and scale factors, both of which determine accuracy, cancel one another when composing asymmetry correction quantities such as $(x_R - x_L) \cdot \partial A/\partial x$. They help in interpreting beam measurements in terms of physical units, but do not contribute significantly to the $A_{\rm PV}$ measurement. On the other hand, high precision is essential, since any beam correction is only as good as the resolution of the device making the beam measurement. That is, systematic uncertainties will be produced that go as $(\sigma_\xi/\sqrt{N}) \cdot \partial A/\partial \xi$ where σ_ξ is the per-pair resolution of a device monitoring beam parameter ξ, and N is the total number of pairs. Given the approximate values for $\partial A/\partial \xi$ listed in Equation (1.21), and assuming that these systematic uncertainties should be kept at or below the ppb level after 100 million pairs, the following per-pair device resolutions are required:

$$\begin{gathered} \sigma_E < 1 \text{ MeV} \qquad \sigma_x, \sigma_y < 10 \ \mu\text{m} \\ \sigma_{x'}, \sigma_{y'} < 0.2 \ \mu\text{rad} \end{gathered} \tag{1.25}$$

Because of the "extra regression" against charge, the resolution of the charge monitoring devices is not terribly important; per-pair resolutions of 50 ppm or lower will suffice. The numbers provided in this section are meant merely to serve as approximate guidelines, not as strict limits. For instance, if 200 million pairs are integrated instead of 100 million, the resolution requirements for all beam monitoring devices can be relaxed by $\sqrt{2}$ without having an impact on the final systematic uncertainty.

1.3.1 Physics Backgrounds

Besides Møller scatters, the detector will also have to contend with a number of different background fluxes [61]. The primary backgrounds are electron–proton (or ep) scatters, photons, and pions. Some of these backgrounds can be subdivided further according to the

specific process by which they are created. Each background will be discussed briefly below, along with the special requirements they may impose upon the experiment. More details can be found in the next chapter.

The ep background represents the largest background, both in terms of absolute numbers and in terms of how it affects the asymmetry measurement. The ep scatters can be divided into elastic and inelastic scatters. The spectrometer is designed to take advantage of the differences in the energy–angle correlations for the Møller and elastic ep (or Mott) distributions, greatly reducing the Mott background flux. However, initial and final state radiation is expected to alter the kinematics of the Mott scatters, thereby increasing the number that make it into the Møller detector's acceptance region. The asymmetry carried by the Mott scatters is roughly of the same order as the Møller asymmetry itself. The inelastic ep scatters, on the other hand, are harder to separate from the Møller scatters due to the lack of a well-defined energy–angle correlation. Furthermore, the asymmetry carried by the inelastic ep flux can be very large (a few ppm), particularly in the Δ^+ resonance region. However, the cross section for inelastic ep scattering is much smaller than for Mott scattering, so the bulk of the ep background in the Møller detector comes from Mott scatters [61]. Regardless, an accurate measurement of the Møller asymmetry will require a correction to be made for the ep background. The details of this correction are given in Section 4.8.3.

The photonic background consists mainly of bremsstrahlung radiation emitted from the target, as well as photons emitted when charged particles hit objects like collimators and the beampipe. As the beam passes through dipole magnets in the spectrometer, synchrotron radiation is also produced. All of these photons, which can be classified as either "hard" (high-energy) or "soft" (low-energy), will add noise and possibly false asymmetries if they are allowed to reach the detector. They are therefore blocked as much as possible by special photon and synchrotron radiation collimators. Neutral hadrons produced by charged particles interacting within the calorimeter can also present a small background. Corrections must be made for the various photonic backgrounds, as well as the neutral hadron background, by carefully measuring the detector's response under a variety of different experimental conditions, as is described in Section 4.8.2.

Finally, pions are produced in the target through real and virtual photoproduction, as well as through deep inelastic scattering [61]. The total cross section for real photo-

Run	*Start date*	*Stop date*	*Pairs*
Run I	12 March 2002	28 May 2002	100 million
Run II	10 October 2002	13 November 2002	86 million
Run III	8 July 2003	8 September 2003	155 million

Table 1.1: Summary of data collection periods. For each period, the number of pairs listed is the approximate number of pairs remaining after all analysis cuts have been applied (analysis cuts are discussed further in Section 4.3). This thesis focuses exclusively on the Runs I and II data collection periods.

production is large ($\sim$100 μbarn), but the asymmetry is of the same order as the Møller asymmetry [61]. On the other hand, the total cross section for virtual photoproduction and pion production through deep inelastic scattering is considerably smaller ($\sim$2 μbarn), but the asymmetry is much larger (at the level of a few ppm) [61]. The spectrometer acceptance for the pions is very small, such that the total pion background is expected to contribute to the primary detector's signal at the sub-1% level. A simultaneous measurement of the pion flux and asymmetry is made with a separate, dedicated pion detector. The details of this correction are given in Section 4.8.4.

1.3.2 Summary of Physics Runs

The total data collection for the SLAC E-158 experiment has occurred in three periods, referred to as Runs I, II, and III. These are summarized in Table 1.1. This thesis concentrates exclusively on Runs I and II, which together account for roughly one-half of the total data set. In general, the experimental conditions were nearly identical for all three runs, with one exception being the installation of several new collimators between Runs I and II. This will be discussed further in the next chapter. Occasionally Run III will be referred to, particularly in the discussion of beam-related systematics. In order to avoid confusion, the data collection period being discussed will always be made explicit.

Chapter 2

Experimental Method

2.1 Overview of SLAC Experiment E-158

Figure 2.1 gives a general overview of the setup for the SLAC E-158 experiment. A discussion of the experimental design naturally divides into seven basic parts: the polarized electron source, the linac, beam monitoring devices, the target, the spectrometer, the detectors, and the data acquisition system. This chapter describes each of these subsystems in turn.

The longitudinally polarized electron beam is generated at the source and accelerated up to an energy of nearly 50 GeV in the two-mile-long linac. When the beam has reached an energy of 1.2 GeV, it enters an area known as the Accelerator System Setup for Experimental Testing (ASSET), which houses beam monitoring devices used as part of feedback loops intended to control helicity correlations in the beam. At the end of the linac, the beam enters the Beam Switch Yard (BSY), at which point it can be steered in one of three

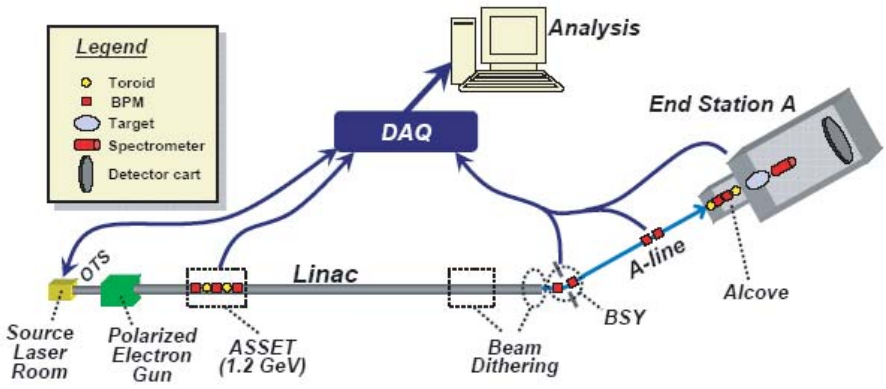

Figure 2.1: Experimental overview. For purposes of clarity, the wire array and two toroids in the alcove are omitted from the picture.

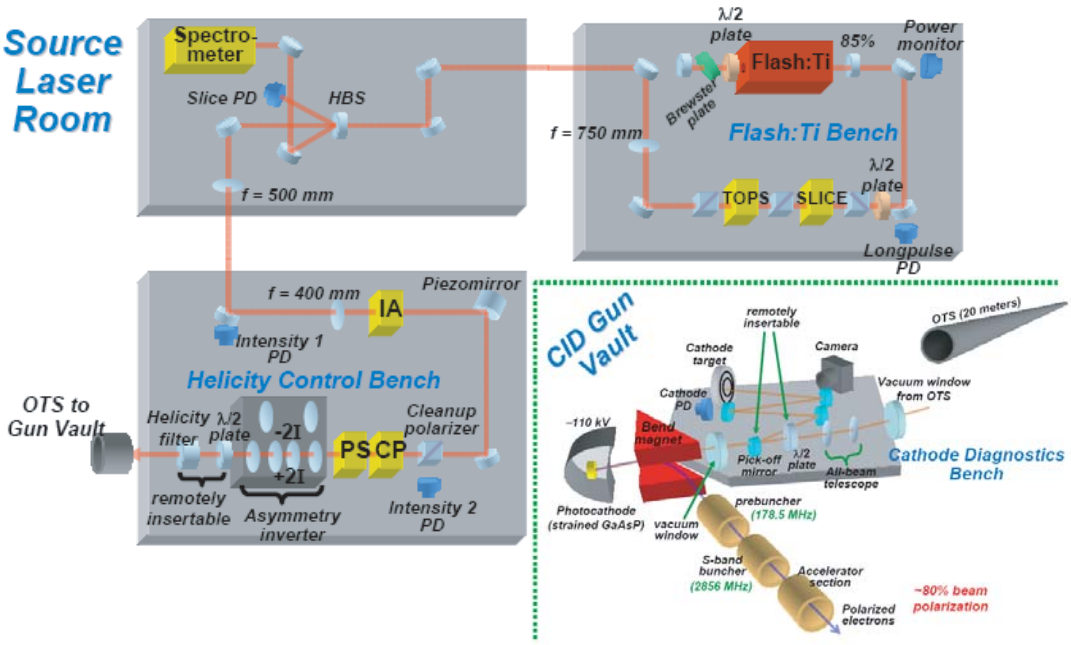

Figure 2.2: Overview of polarized electron source. The mirror box (which brings the laser down to the level of the photocathode immediately after the OTS) is not shown.

directions. For this experiment, it is steered into the A-line, which eventually brings the beam to the main experimental area, known as End Station A (ESA). Along the way, the beam's current, position, energy, and spot size are monitored by various devices. The target, spectrometer, and detectors are all housed in ESA. The spectrometer selects from the many assorted particles that scatter from the target just those of physical interest, and focuses them onto the main integrating calorimeter. Signals from all devices are both monitored on-line and stored for subsequent off-line analysis by the data acquisition system.

2.2 Polarized Electron Source

The Polarized Electron Source (usually abbreviated as the PES, but often referred to simply as "the source") consists of two rooms, the Polarized Light Source (PLS) room and the gun vault, separated by a 20 m long tube called the Optical Transport System (OTS). An overview of the entire source can be found in Figure 2.2. The tunable pulsed laser system generates a 300 ns pulse of left or right circularly polarized light, which first travels through optics designed to minimize all helicity correlations, and then is brought into the gun vault. Here the laser light impinges on the dc-biased photocathode to produce the polarized electrons which travel into the injector and then into the linac. The following

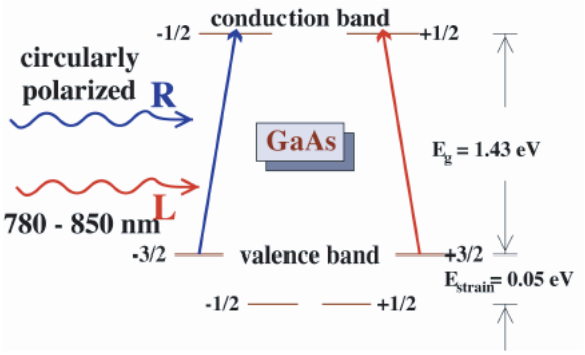

Figure 2.3: Source operating principle. When circularly polarized light strikes the GaAs crystal (the laser wavelength tuned to the GaAs bandgap energy), selection rules exist that only favor certain transitions ($\Delta m_j = \pm 1$), thanks to the semiconductor's direct–bandgap nature. A mechanical strain lifts the degeneracy of the $P_{3/2}$ valence band, making some transitions energetically disallowed.

sections describe the entire source in a fair amount of detail. First, a description of the physical setup will be given, starting with a discussion of the final stage of the source, photoemission from the photocathode, and then proceeding with a description of the rest of the source optics system. This will lead into a discussion of the major sources of beam systematics and the methods used to control and correct for them.

2.2.1 Photocathode

The operating principle behind the source, illustrated in Figure 2.3, is basically the same as it was for SLAC E-122 back in the late 1970's [62, 63, 64]. A gallium arsenide (GaAs) photocathode sits in vacuum in the gun vault, and circularly polarized laser light, its wavelength (typically 780–850 nm) tuned to the bandgap energy of the GaAs crystal, enters the gun vault from the PLS via the OTS, where it impinges on the photocathode. If the band gap energy of the GaAs is close to its work function, then photoemission can occur. If the band gap energy is less than the work function (a characteristic known as negative electron affinity, or NEA), then photoemission is likely [65]. Because GaAs is a direct band gap semiconductor, the electron is left with zero momentum ($\vec{k} = 0$) after photoemission. Consequently, the angular momentum of the photon goes directly into the spin of the photoemitted electron. Selection rules therefore exist that govern the transitions between the $P_{3/2}$ valence band and the $S_{1/2}$ conduction band when circularly polarized light is used. For example, for right circularly polarized light, only $\Delta m = +1$ transitions are allowed. The extent to which this feature can be used to produce high polarization beams will be

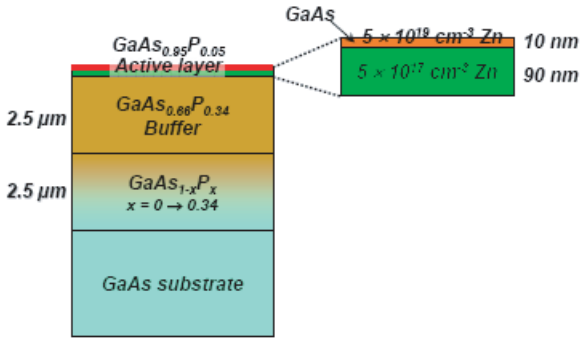

Figure 2.4: Photocathode structure for Runs I and II. The high doping level (5×10^{19} cm^{-3} Zn) in the 10 nm GaAs surface overcomes the surface charge limit effect. The relatively low doping level (5×10^{17} cm^{-3} Zn) in most of the active layer is to ensure minimal polarization degradation.

discussed below.

The photocathode used for E-158 Runs I and II is illustrated in Figure 2.4 [66, 67]. The photocathode is grown by a process known as Metal Organic Chemical Vapor Deposition (MOCVD) on a GaAs substrate. In MOCVD, other Group V elements (for example phosphorous) can be substituted for arsenic during the growth process, and these will take the place of arsenic in the lattice (since they are chemically similar), thereby altering the spacing of the lattice (since they are not of the same size). Immediately atop the GaAs substrate is a 2.5 μm thick layer of a $GaAs_{1-x}P_x$ binary mixture, where x varies linearly from 0 at the bottom to 0.34 at the top. The 2.5 μm thick buffer layer immediately below the active layer, from which the photoelectrons are drawn (and whose 100 nm thickness corresponds roughly to the penetration depth of the laser photons), consists of $GaAs_{0.66}P_{0.34}$. The amount of phosphorous added during the growth process then suddenly changes going from the so-called buffer layer to the active layer, causing the lattice constants of the two respective layers to be different. This sudden lattice mismatch between the active layer (which contains only 5% P) and the buffer layer (which contains 34% P) produces a small mechanical strain in the system which lifts the degeneracy of the $P_{3/2}$ states of the valence band electrons.

The lifting of the degeneracy of the $P_{3/2}$ valence band states via the strain mechanism allows for beam polarizations of greater than 50% to be achieved [68]. For example, if the degeneracy was not lifted, then either the $-3/2 \rightarrow -1/2$ or the $-1/2 \rightarrow +1/2$ transition could occur during photoemission initiated by a right-circularly-polarized photon. The electron beam would still be longitudinally polarized, but its maximal polarization would then

be only 50%. The lattice mismatch is thus essential for a precision parity violation experiment like E-158 because of its need for high beam polarization, but it has the unfortunate side effect of introducing a sensitivity in the response of the photocathode to residual linear polarization in the incoming laser light [68, 69]. In other words, the photocathode will emit more electrons for light that is linearly polarized along some preferred axis. Effectively, the photocathode's quantum efficiency gains an analyzing power of a few percent, which plays an important role in producing unwanted beam asymmetries.

Negative electron affinity is induced by depositing cesium and nitrogen trifluoride (NF_3) on the photocathode's heat-cleaned surface [70]. This treatment (part of the photocathode's activation procedure) increases the quantum efficiency (QE) of the photocathode by several orders of magnitude. This enhancement, however, is very sensitive to the photocathode's complex surface chemistry and typically deteriorates over the course of several days as residual gases interact with the NEA surface. To minimize such interactions, the photocathode is kept in an ultra-high vacuum environment. Cesiations occur roughly every three days to restore the QE, typically to $\sim$0.4%.

Doping the active layer with zinc overcomes the phenomenon known as surface charge limit [71, 72]. Too much doping, however, degrades polarization. In order to maximize polarization, previous experiments tended to use relatively low doping levels. As these photocathodes were hit with more and more laser power, they became saturated, emitting fewer and fewer electrons per microjoule of incident beam energy. The solution was found by doping most of the active layer (90 nm out of 100 nm) with relatively low levels, and then increasing the doping level by a factor of 100 (to 5×10^{19} cm^{-3}) for the topmost 10 nm

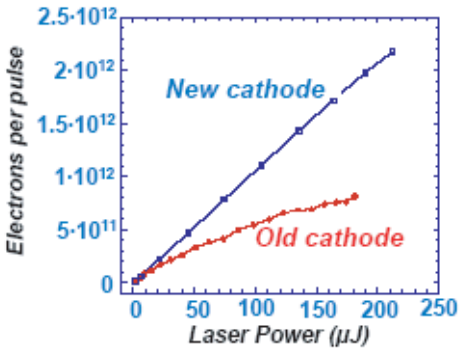

Figure 2.5: Comparison of surface charge limit for old vs. new photocathodes. Photocathode output is plotted versus laser power for the old and new photocathodes. Because of the highly doped surface layer of the new photocathode, its output remains linear when the laser power is increased.

portion of the active layer. As can be seen in Figure 2.5, the new photocathode does not exhibit any signs of charge limit.

2.2.2 Polarized Light Source

The functions of the PLS can be broken into two basic, separate parts: 1) production of a stable pulse of laser light tuned to the band gap energy of the GaAs photocathode, and 2) ensuring that it is as perfectly circularly polarized as possible at the photocathode, and contains minimal helicity correlations in intensity, position, and angle. The first two optics benches in the PLS room (the Flash:Ti and diagnostic benches) are largely concerned with the former function, whereas the last bench (the helicity control bench) is exclusively concerned with the latter. Each part will now be described.

Flash:Ti Bench

The Flash:Ti laser system is documented in much more detail elsewhere [73, 74, 75]. Inside the laser cavity, designed at SLAC and built by Big Sky Laser Technologies, two flashlamps are used to pump a rod-shaped Ti:Sapphire crystal. The flashlamps can fire at any rate up to 120 Hz, with their firing rate of course determining the maximum beam rate (see Section 2.3). The one-meter-long cavity is formed by an 85%-reflectivity planar output coupler mirror and a 99.9%-reflectivity end mirror with a 42 m concave curvature, chosen to minimize the effects of thermal lensing , thereby optimizing laser stability. A quartz quarter-waveplate acts as both a Brewster plate, selecting out horizontal polarization, and a birefringent tuner, which by rotation about its normal axis selects out a wavelength appropriate for the GaAs band gap energy (typically in the range of 780 nm to 850 nm) with a bandwidth of $\sim$0.7 nm (FWHM). A half-wave plate was at one time required to compensate for the arbitrary orientation of the Ti:Sapphire laser rod, thereby maximizing transmission through the Brewster plate. However, between Runs I and II modifications to the laser head assembly procedure were made that allowed for control over the crystallographic orientation of the laser rod, thereby eliminating the need for the half-wave plate (whose removal was seen to further improve laser stability).

When laser light exits the cavity, it is horizontally polarized and has a temporal profile, shown in Figure 2.6, largely determined by the pulse shape of the voltage signal driving the flashlamps. It is then necessary to cut out a "slice" from this $\sim$15 μs long pulse, which is

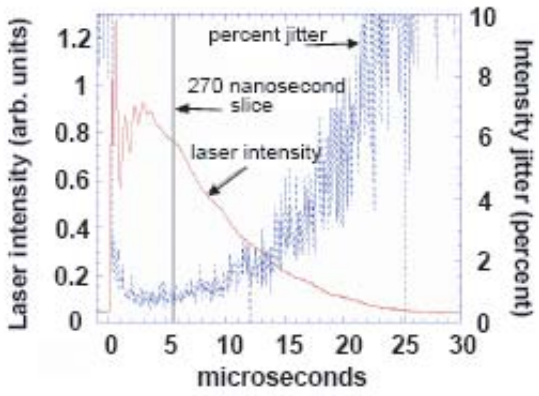

Figure 2.6: Laser temporal profile, before pulse shaping. On the time axis, $t = 0$ corresponds to when the flash lamp fires. After a settling time of a few microseconds, the relative energy jitter in the laser steadily grows as a function of time (as the intensity drops). The timing of the SLICE high voltage signal corresponds to a period in which the relative energy jitter is low ($< 1\%$).

done with the slice-selection optics. The slice-selection optics consist of a Pockels cell sandwiched between two crossed polarizers. The first polarizer ensures horizontal polarization, whereas the second naturally extinguishes all light. A Pockels cell is an electro-optical device consisting of a crystal whose birefringence is proportional to the voltage applied across its face. It can therefore act as a variable waveplate whose phase retardation (or advancement) can be controlled electrically. For slice-selection, the so-called "SLICE" Pockels cell receives a high voltage pulse of duration equal to the desired length of the electron beam pulse and of amplitude equal to the half-wave plate voltage of the Pockels cell (about 2800 V). While the high voltage pulse lasts, the orientation of the laser beam's linear polarization rotates by 90°, enabling it to pass through the second polarizer.

Using the slice-selection optics, the duration, intensity, and energy jitter of the laser pulse can be controlled. For instance, lengthening or shortening the SLICE high voltage signal lengthens or shortens the pulse. Typical sliced pulse lengths are 50–370 ns. Tweaking the amplitude of the high voltage signal (away from the half-wave voltage) rotates the polarization by less than or more than 90°, allowing less light to pass through the second polarizer and lowering the intensity of the pulse. In fact, the amplitude of the SLICE high voltage signal is used as part of a linac feedback intended to maintain constant intensity in the electron beam, compensating for the slow decrease in the quantum efficiency of the photocathode during its 3-day cesiation cycle. Lastly, changing the timing of the high voltage signal relative to the flash lamp timing determines the energy jitter of the pulse, as the relative jitter primarily grows over time. This can be seen in Figure 2.6.

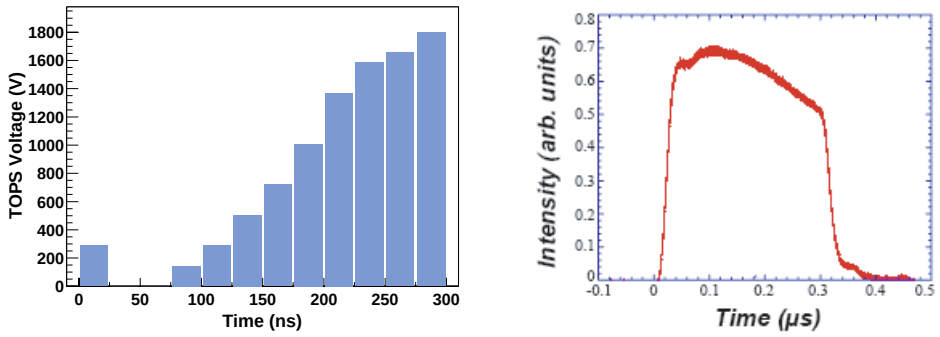

(a) (b)

Figure 2.7: Example of pulse shaping waveform. Figure (a) shows the voltage supplied to the TOPS Pockels cell versus time. Figure (b) shows the resulting time profile of the beam intensity. Shaping the charge time profile in this way helps alleviate the effects of beam loading.

After the laser slice has been formed, it needs to be shaped in order to compensate for beam loading effects in the linac (described more in Section 2.3). The pulse-shaping optics consist of another Pockels cell, called the "TOPS" Pockels cell, sandwiched between two co-aligned polarizers. The dc power supply for the TOPS cell is connected to a function generator allowing for modulation of the output high voltage signal according to an arbitrary waveform generated in 25 ns steps. In practice, a waveform such as that shown in Figure 2.7(a) is used, so that, as one moves along in time, the orientation of the laser light's polarization gets rotated by small amounts, meaning that less light passes through the downstream polarizer. This configuration allows the controller to give the intensity of the laser pulse a negative time gradient, as shown in Figure 2.7(b), reducing the overall effects of beam loading in the electron beam.

Diagnostics Bench

Immediately downstream of the Flash:Ti bench, the diagnostics bench contains optics connected to monitoring various aspects of the sliced laser pulse. Here and elsewhere, photodiodes are located behind the broadband NIR-coated high-reflectivity mirrors used to reflect the laser beam along its designated path. These photodiodes monitor leakage light and provide measurements of intensity, wavelength, and spot size. On the Flash:Ti bench, a photodiode (sometimes called the long-pulse photodiode) located behind the mirror immediately upstream of the slice-selection optics monitors the total laser output. Also on the Flash:Ti bench, a CCD camera installed behind the mirror immediately downstream of the pulse-shaping optics monitors the spatial profile of the sliced laser pulse. On the diagnostics

bench, a holographic beam sampler diverts two one-percent samples of the laser beam at small angles to the primary beam path. One of these samples goes to a photodiode, while the other can either be focused onto a scanning monochromator for wavelength measurements or (if a remotely insertable mirror is used) can be imaged onto another CCD camera for spatial profile monitoring.

Helicity Control Bench

By the time the laser beam reaches the helicity control bench, it is vertically linearly polarized and sliced into a short pulse of 50–370 ns. As discussed in Section 2.2.1, circularly polarized light is needed in order to produce longitudinally polarized electrons. All of the optical components located on the helicity control bench therefore have the dual and intimately-related responsibilities of ensuring near-perfect circular polarization at the photocathode while also minimizing helicity correlations in all laser beam parameters (besides polarization, obviously). Furthermore, they must allow for rapid helicity flips, for which purpose electro-optical devices such as Pockels cells are ideally suited.

Immediately downstream of the intensity asymmetry (IA) Pockels cell and the piezomirror (a description of whose functions will be deferred until Section 2.2.6), the beam encounters the polarization optics, consisting of a vertically aligned polarizer followed by two Pockels cells. The so-called cleanup polarizer serves two purposes. First, it determines the orientation of the linear polarization for all downstream optics, most importantly the CP and PS Pockels cells, whose fast optical axes can be accurately set relative to this orientation, rather than that of some other optical element further upstream. Second, it is used as the point of combination for the YLF:Ti beam that is used to generate electrons for the PEP rings as part of the BaBar experiment, which collected data concurrent with this experiment during Run I. Both the Flash:Ti and the YLF:Ti beams share a common path to the photocathode after they are combined at the cleanup polarizer. The first Pockels cell in the polarization optics is called the circular polarization, or just CP, cell. It is supplied with a high voltage pulse of amplitude $\sim$2800 V, corresponding roughly to quarter-wave voltage, the sign of which determines the helicity of the pulse, either right or left. The fast axis of the CP cell is oriented at 45° with respect to the vertical, so that, for vertical incoming polarization, circular polarization of either helicity can be produced. The second Pockels cell, referred to as the phase shift, or PS, cell, with a vertical fast axis, nominally operates

at zero voltage. Small (on the order of 10 V) helicity-correlated deviations from zero are used to correct for phase shifts in the downstream optics and maintain near perfect circular polarization at the photocathode. The exact way in which the polarization Pockels cells' voltages are set, and the importance of the choice of voltages, will be more fully explained in Section 2.2.4, once the various sources of systematics have been described.

Two photodiodes on the helicity control bench monitor the total energy in the pulse. One of these photodiodes detects leakage light from the cleanup polarizer, and so is sensitive both to the Flash:Ti pulse as well as to the YLF:Ti pulse. The last optical element on the helicity control bench, the insertable helicity filter is used when determining the optimal operating voltages for the CP and PS cells. The helicity filter itself is a linear polarizer affixed to a quarter-wave plate in such a way that the whole arrangement transmits right helicity light while extinguishing left helicity light. In practice, the CP and PS voltages are swept one at a time over an appropriate range while the helicity filter is inserted into the beam path and a downstream photodiode monitors transmitted intensity. Fitting the data to parabolas yields the voltages to use in order to maximize or minimize transmission, as the case may be, thereby ensuring optimal circular polarization. An extinction ratio R is defined by:

$$R = \frac{T_{\text{right}}}{T_{\text{left}}}, \tag{2.1}$$

where T_{right} (T_{left}) is the transmission for right (left) helicity light. An extinction ratio of $R > 1000:1$, for instance, implies circular polarization better than 99.8%, which in turn implies a linearly polarized component of less than 0.2%. We require that R satisfies this criterion. If it does not, then realignment of the CP and PS cells is necessary. This procedure does not finalize the CP and PS cells' voltages, however. First, when the electron beam becomes available (after initial commissioning and setup), a toroid early on in the linac is used to measure the charge asymmetry as the CP and PS voltages are tweaked about their nominal values. This procedure, which will be explained in Section 2.2.5, compensates for various effects, such as phase shifts in the optics downstream of the CP and PS cell, and minimizes the asymmetry in the linear polarization at the photocathode. Second, when the electron beam has been successfully transported all the way down the linac machine and into End Station A, the Møller polarimeter is used to check that the Pockels cell voltages are indeed those that result in the highest polarization at the target. Typical results of

these types of scans will be given in Section 4.2.

Located just after the PS cell, the asymmetry inverter (sometimes abbreviated as the AI) comprises two sets of lenses, a lens doublet and a lens quadruplet, which act as telescopes of magnification $M = \pm 2.25$. In the doublet configuration, which possesses the magnification factor $M = -2.25$, the trajectory and spatial profile of the beam undergo a reflection about the origin. This is useful in cancelling helicity-correlated position differences in the electron beam. The two sets of lenses sit on an optics stage which can be slowly translated by hand from one configuration to the other. Typically the experiment ran for a week in one configuration and then toggled the AI into the other configuration, hoping in this way to achieve some level of cancellation of position-related systematics. An insertable half-wave plate is mounted just downstream of the asymmetry inverter. Along with another half-wave plate located on the cathode diagnostics bench, this half-wave plate can be used to provide cancellation of certain classes of beam asymmetries.

Optical Transport System

For practical reasons involving space, safety, and maintenance concerns, the source laser room is located 20 meters away from the gun vault. It is the job of the optical transport system (OTS) to transport the beam from the laser room to the optics bench located just before the entrance to the vacuum containing the photocathode. The OTS consists of a 20-m-long pipe containing a single 5-m imaging lens and a mirror box that brings the laser beam down to the height of the optics bench (the cathode diagnostics bench, described next) located just in front of the entrance to the photocathode gun. The imaging lens compensates for beam divergence. The mirrors in the mirror box are designed to preserve the high degree of circular polarization. There are two pairs of helicity-compensating mirrors, with the mirrors within each pair acting to interchange s and p–polarization upon reflection. In this way, differences in phase shifts and reflectances for s and p–polarizations are cancelled out.

The OTS was once held under vacuum, with window plates installed on either end, to eliminate air currents that could potentially steer the beam (there is a $\sim 10^{\circ}$C temperature gradient between the laser source room and the gun vault). However, for this experiment the OTS was brought up to atmospheric pressure in order to eliminate the stress-induced birefringence in the window plates, which could otherwise produce helicity correlations in the electron beam (as will be discussed in Section 2.2.4).

Cathode Diagnostics Bench

The first item on the cathode diagnostics bench, the imaging telescope, isn't a diagnostic at all, but is used to determine the beam spot size on the photocathode. The photocathode is ∼20 mm in diameter and, unlike previous photocathodes, does not need to be fully illuminated in order to meet the experiment's beam current requirements. This flexibility allows one to reduce the spot size by almost half while simultaneously moving the object point to within a few centimeters of the CP cell. This is a significant reduction of the ∼20 m lever arm actually separating the CP cell and photocathode, and reduces the sensitivity to Pockels cell lensing effects. After the imaging telescope, another remotely insertable half-wave plate is mounted. The position (in or out of the beam) of this half-wave plate was toggled about every two days throughout the entire physics run. This was very important for reducing the experimental sensitivity to certain classes of beam-related systematics.

Opposite the window leading out of the OTS and immediately after the remotely insertable pick-off mirror (which will be described in a moment), the beam encounters the vacuum window leading to the beam pipe connected to the photocathode gun. Inside the gun, the photocathode is held in ultra-high vacuum at a dc bias voltage of −120 kV [70]. The photoemitted, polarized electrons are thus drawn away from the photocathode back the way the light came in (the gun is thus said to operate in "reflection mode"), until a dipole magnet bends them into the injector for bunching and subsequent acceleration. The HV power supply recharges the photocathode in between every pulse as part of an RC-circuit.

The remaining optical elements on the cathode diagnostics bench come into play only when a 50% pick-off mirror is inserted into the beam path just upstream of the vacuum window. A paper target is monitored by a video camera. Multiple mirrors reflect the beam so that the path length from the pick-off mirror to the cathode target is exactly the same as to the real photocathode. The video camera thus gives an image of the beam spot exactly as it appears on the actual photocathode. This is an extremely useful feature, especially for mapping the quantum efficiency over the photocathode's surface. A diagnostic HeNe laser (having a very small ∼1 mm spot) is used with the pick-off mirror inserted, and the current drawn from the photocathode is measured as the beam is swept back and forth both horizontally and vertically. The apparatus is also useful for determining the object point of the imaging telescope. A wire mesh screen is placed in the beam path near the polarization

Wavelength:	805 nm (750 – 850 nm)
Bandwidth:	0.7 nm FWHM
Repetition rate:	120 Hz (max.)
Cavity output:	45 mJ (5.4 W at 120 Hz)
Pulse length:	250 ns (50 – 370 ns)
Pulse intensity:	60 μJ in 250 ns (typ.) 600 μJ in 370 ns (max.)
Intensity jitter:	0.5% RMS
Position jitter at photocathode:	< 70 μm RMS
Circular polarization:	$> 99.8\%$
Photocathode type:	gradient-doped, strained GaAs
Electron beam polarization:	$\sim 85\%$
Electrons per pulse:	5×10^{11} (typ.) 2×10^{12} (max.)
Energy spread (e^- beam):	$\sim$0.15%

Table 2.1: Summary of the source performance. Unless otherwise noted, all items refer to laser performance. The laser position jitter at the photocathode is given for a beam $\sim$1 cm in diameter. Electron beam performance is further discussed and summarized in Chapter 4.

Pockels cells on the helicity control bench, and its z-position is varied as the resulting image on the cathode target is observed. By finding the position for which the wire mesh screen's shadow appears sharpest, the object point is determined to be within a few centimeters of the CP cell, as noted above. Behind one of the mirrors on the cathode diagnostics bench, a photodiode monitors leakage light and can be used to give measurements of transmission and of the actual laser power reaching the photocathode.

2.2.3 Source Performance

A summary of the overall source performance appears in Table 2.1. This table mainly summarizes the performance of the source laser system, but it also sheds light on three key aspects of the electron beam — polarization, charge, and energy spread (related to beam loading) — crucial to the success of the experiment and intimately related to the performance of both the laser and photocathode. The feasibility of this experiment depends on the ability to produce a high-charge, high-polarization, and high-energy electron beam, making SLAC an ideal site for its implementation.

Serious demands were placed on the source to provide a highly stable electron beam of high current and high polarization. These demands were indeed met, thanks to the new

gradient-doped photocathode structure, which SLAC developed, as well as to some of the novel techniques employed in the source optics setup. For example, the TOPS pulse-shaping is essential for minimizing energy spread due to beam loading, which in turn is essential for the delivery of a stable, usable ∼50 GeV electron beam.

One quantity that had an especially big impact on numerous aspects of the experiment was the laser intensity jitter. Historically pulse-to-pulse laser intensity jitter was at the level of $1-2\%$ or even higher. Much effort was made to reduce the jitter, and thanks to many small improvements in the Flash:Ti cavity, the laser power supply electronics, and the photocathode,[1] the intensity jitter was gradually reduced to its present value of no more than 0.5%. Minimizing laser intensity jitter, by lowering the intensity jitter in the resulting electron beam, helps the experiment in many ways. It lowers energy jitter in the electron beam, improves beam delivery as well as beam quality, and plays a part in reducing systematic effects.

Table 2.1 contains information about the beam's polarization, charge, and jitter, but it makes no mention of any helicity correlations present in either its intensity, energy, position, or spot size. Minimizing such helicity correlations, in addition to producing a high-current beam that is both stable and of high polarization, is the primary concern of the source. The next section describes how unwanted helicity correlations can arise in the electron beam. The ways in which these potential sources of systematics are minimized will then be discussed.

2.2.4 Sources of Systematics

Helicity correlations (wanted and otherwise) are introduced for the first time on the helicity control bench, starting with the polarization optics. For the purposes here, it will be helpful to use the Jones matrix formalism, whereby a 100% polarized beam is represented in vector notation in terms of its electric field vector [76]:

$$\left|\vec{E}\right\rangle = \begin{pmatrix} E_x(t) \\ E_y(t) \end{pmatrix} = E_0\, e^{i(kz-\omega t)} \begin{pmatrix} \cos\theta \\ e^{i\phi}\sin\theta \end{pmatrix} \tag{2.2}$$

[1] Improving the photocathode doesn't itself reduce laser intensity jitter, but it allows for modifications to the laser system that *can* reduce its intensity jitter. For example, the new photocathode requires a wavelength of 805 nm for peak polarization. This wavelength is much closer to the gain maximum of the Ti:Sapphire laser crystal, which significantly enhances its performance.

The varying amplitude term out front can be dropped since it is common to both components. Using this notation, then, the laser light after the cleanup polarizer can be represented as follows:

$$\left|\vec{E}\right\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \tag{2.3}$$

This defines the vertical axis. Next come the CP and PS Pockels cells. In general, the Jones matrix for a Pockels cell pulsed at voltage V can be written as:

$$\begin{aligned} &\begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & \mathrm{e}^{i\pi V/V_{hw}} \end{pmatrix} \begin{pmatrix} \cos(-\theta) & -\sin(-\theta) \\ \sin(-\theta) & \cos(-\theta) \end{pmatrix} \\ &\qquad = R(\theta) \begin{pmatrix} 1 & 0 \\ 0 & \mathrm{e}^{i\delta} \end{pmatrix} R(-\theta) \end{aligned} \tag{2.4}$$

where $R(\theta)$ is a rotation matrix used to account for the orientation of the Pockels cell's fast optical axis with respect to the incoming polarization vector, and δ is the phase retardation introduced between the Pockels cell's fast and slow axes. It follows that V_{hw} is the *half-wave voltage* of the Pockels cell, or the voltage for which the Pockels cell acts as a half-wave plate. The CP cell has its fast axis at 45°, while the PS cell has a vertical fast axis. Using this formalism, then, the electric field vector of the laser light after the PS cell takes on the following form:

$$\left|\vec{E}\right\rangle = \begin{pmatrix} \sin(\delta_{CP}/2) \\ \mathrm{e}^{i(\pi/2+\delta_{PS})}\cos(\delta_{CP}/2) \end{pmatrix} \tag{2.5}$$

where δ_{CP} and δ_{PS} are the phase retardations for the CP and PS Pockels cells, respectively. Under nominal conditions, the CP cell behaves like a quarter-wave plate ($V_{CP} = \pm V_{hw}/2$), whereas the PS cell is not pulsed with any voltage. This gives $\delta_{CP} = \pm\pi/2$ and $\delta_{PS} = 0$, which from Equation (2.5) can be seen to produce perfect circular polarization of either helicity. In order to generate perfect circular polarization at the photocathode, however, we must be able to compensate for residual birefringence both in the Pockels cells themselves as well as in various optical elements downstream of the polarization optics. Residual birefringence throughout the transport system is a major concern and, coupled with any analyzing power anywhere in the system (such as that of the strained photocathode), is capable of producing large helicity-correlations in the electron beam.

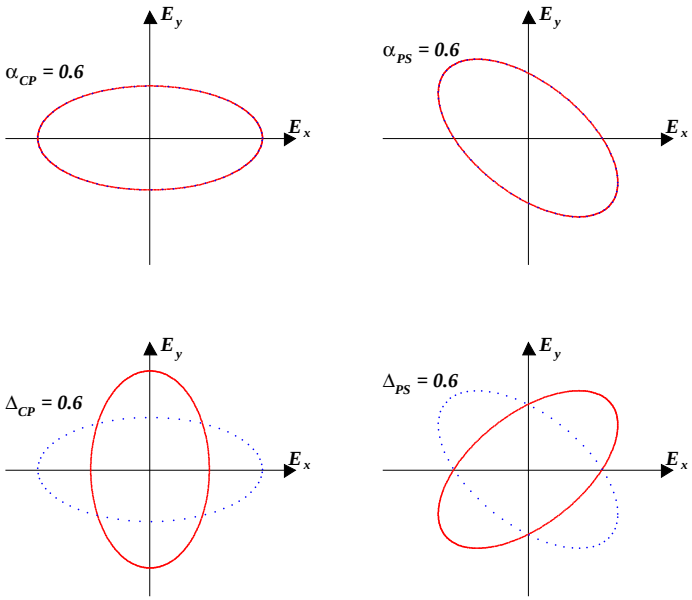

Figure 2.8: Effects of non-zero phase shifts $\alpha_{\rm CP}$, $\alpha_{\rm PS}$, Δ_{CP}, and Δ_{PS}. Polarization ellipses are shown for the laser light at the photocathode. The solid red line denotes the left-helicity ellipse, while the dotted blue line denotes the right-helicity ellipse. Only nonzero Δ parameters are capable of producing linear polarization asymmetries at the photocathode.

In order to compensate for the imperfections produced by residual birefringence, deviations from the nominal operating voltages are necessary. Four new parameters ($\alpha_{\rm CP}$, Δ_{CP}, $\alpha_{\rm PS}$, and Δ_{PS}) are needed to completely span the choice of voltages. The right and left phase retardations for the CP and PS cells can thus be written in the following manner:

$$\begin{aligned} \delta_{CP} &= \pm(\pi/2 + \alpha_{\rm CP}) + \Delta_{CP} \\ \delta_{PS} &= \pm\alpha_{\rm PS} + \Delta_{PS} \end{aligned} \tag{2.6}$$

From this equation, one can see that $\alpha_{\rm CP}$ and $\alpha_{\rm PS}$ correspond to *symmetric* deviations (they affect the voltages of both helicity states by equal but opposite amounts, keeping them *symmetric* about the zero-axis), while Δ_{CP} and Δ_{PS} correspond to *asymmetric* deviations (they ruin the symmetry about the zero-axis). Consequently the α and Δ parameters are often referred to as symmetric and asymmetric phase shifts, respectively.

Another, more physical way in which the phase shifts can be seen to be symmetric and asymmetric is in the effects they have on the polarization of the laser light. Any deviations from the nominal operating setpoints ($\delta_{CP} = \pm\pi/2$, $\delta_{PS} = 0$) distort the perfect circular polarization, turning it into elliptical, but symmetric and asymmetric phase shifts differ in

how they do this. The effects of non-zero α and Δ are shown in Figure 2.8. For $\alpha \neq 0$ and $\Delta = 0$, the major and minor axes of the polarization ellipse are the same for right and left helicity light. In other words, the linear polarization component has the same direction for right helicity light as it does for left helicity light. By contrast, $\alpha = 0$ and $\Delta \neq 0$ produces elliptically polarized light with the major and minor axes of the ellipse being switched for right and left helicity light. In other words, the linear polarization components for right and left helicity light are orthogonal.

This turns out to be a very important distinction. On the one hand, for $\Delta \neq 0$, one has a linear polarization asymmetry at the photocathode, while for $\alpha \neq 0$ one does not. Consider what would happen if an element was introduced into the system that possessed a higher transmission efficiency for linearly polarized light along some preferred axis. To a very small level, at least, every optical element that is not perfectly normal to the beam will possess a finite analyzing power. The dominant analyzing power in our system, however, is that of the strained photocathode, which possesses a QE analyzing power on the order of a few percent. Qualitatively, this will produce an intensity asymmetry equal to the level of the linear polarization asymmetry (itself proportional to Δ) multiplied by the analyzing power in the system. For instance, a linear polarization asymmetry of 1% coupled with an analyzing power in the photocathode of 1% will naturally produce a charge asymmetry in the electron beam on the order of $0.01 \times 0.01 \approx 100$ ppm.

Of course, this treatment just gives an order of magnitude estimate for the effects one may expect to observe; the actual size of the effects will be larger or smaller depending on the exact size of the linear polarization asymmetry, the size of the analyzing power, and the orientation of the analyzer's preferred axis with respect to that of the linear polarization asymmetry. A derivation of the general form of the so-called PITA (Polarization Induced Transport Asymmetry) effect can be found in several sources [75, 77], with the final result being:

$$A_Q \equiv \frac{q_R - q_L}{q_R + q_L} = m_{\mathrm{CP}} \cdot \Delta_{CP} + m_{\mathrm{PS}} \cdot \Delta_{PS} \tag{2.7}$$

where A_Q is assumed to be the charge asymmetry in the resulting electron beam, q_R and q_L are individual charge measurements for right and left helicities, respectively, and m_{CP} and m_{PS} are referred to as the PITA slopes for the CP and PS Pockels cells, respectively.

To summarize, looking back at Figure 2.8, one realizes that by manipulating α_{CP}, α_{PS},

Δ_{CP}, and Δ_{PS}, it is possible to generate light with arbitrary elliptical polarization at the photocathode. The symmetric phase shifts $\alpha_{\rm CP}$ and $\alpha_{\rm PS}$ produce linear polarization components that are symmetric with respect to helicity. For $\alpha_{\rm CP} \neq 0$ the linear polarization is along the x or y axis, while for $\alpha_{\rm PS} \neq 0$ it is along the u or v axis (the u–v axis system is the x–y axis system rotated by 45°). By contrast, the asymmetric phase shifts Δ_{CP} and Δ_{PS} produce linear polarization components that are asymmetric with respect to helicity. For $\Delta_{CP} \neq 0$ the linear polarizations are along the x and y axes (e.g., x-axis for right helicity, y-axis for left helicity). For $\Delta_{PS} \neq 0$ the linear polarizations are along the u and v axes (e.g., u-axis for right helicity, v-axis for left helicity).[2]

In addition to phase shifts that are controlled experimentally, imperfections in optical components can give rise to their own contributions to Δ_{CP} and Δ_{PS}. These contributions are called residual birefringences, and they will be the first source of systematics specifically discussed below. Residual birefringences couple only to charge asymmetries and, because of beam loading effects described in Section 2.3, energy asymmetries. Other effects will be seen to produce helicity correlations in all of the various beam parameters. After describing each effect and the types of beam asymmetries it is capable of producing, the next section will turn towards the different ways in which these effects may be minimized.

Residual Birefringence

Up until this point, it has been implicitly assumed that there are no birefringences in the system other than those possessed by the CP and PS Pockels cells, and that these birefringences can be totally ascribed to the Pockels effect. In reality, not only do the Pockels cells contain natural birefringence (which will not be proportional to the $\vec{E}$ field applied to the cell), but so do many other optical elements in the source optics system. For example, a typical glass window will produce a phase shift on the order of 5 nm/cm, which, for perfectly circularly polarized light at 800 nm, can produce charge asymmetries in the electron beam on the order of 1000 ppm, given typical photocathode QE anisotropies of $\sim 5\%$. These natural birefringences are referred to as residual birefringence.

The residual birefringence of the Pockels cells themselves can be accounted for by sep-

[2]Readers familiar with the work of George Stokes will recognize these sensitivities as being to Stokes parameter 1 (S_1) and Stokes parameter 2 (S_2), respectively.

arating Δ_{CP} and Δ_{PS} in Equation (2.7) into two variables each, like so:

$$\begin{aligned} \Delta_{CP} &\Longrightarrow \Delta_{CP} + \Delta^0_{CP} \\ \Delta_{PS} &\Longrightarrow \Delta_{PS} + \Delta^0_{PS} \end{aligned} \tag{2.8}$$

Here the Δ^0 terms are the residual birefringences resulting from imperfections in the Pockels cells, and Δ_{CP} and Δ_{PS} are the phase shifts actually controlled by the experimenter. Substituting Equation (2.8) into Equation (2.7) clearly shows, given some finite amount of residual birefringence, voltage offsets will be necessary from the nominal operating points of $V_{CP} = \pm V_{QW}$ and $V_{PS} = 0$.

In general, phase shifts in any optical component downstream of the polarization optics can be accounted for by making the following substitutions in Equation (2.7):

$$\begin{aligned} \Delta_{CP} &\Longrightarrow \Delta_{CP} + \Delta^0_{CP} + \Delta^1_1 + \cdots + \Delta^n_1 \\ \Delta_{PS} &\Longrightarrow \Delta_{PS} + \Delta^0_{PS} + \Delta^1_2 + \cdots + \Delta^n_2 \end{aligned} \tag{2.9}$$

where the i superscript refers to the i^{th} optical element being considered, and the subscripts 1 and 2 denote sensitivity to the Stokes parameters S_1 and S_2, respectively (see the discussion connected with the footnote on page 43).

Birefringence Gradients

Having introduced the variables Δ^i, it is natural to ask what would happen if these new variables possessed spatial variation. In this case, Equation (2.7) would gain spatial variation, and the electron beam would thus have a charge asymmetry that varied with position. A spatially varying charge asymmetry is immediately recognized as a helicity-correlated position difference. An example scenario is illustrated in Figure 2.9, which shows a birefringence with a linear position gradient and the helicity-correlated position difference in the electron beam that it would cause, assuming a Gaussian laser beam profile. In general, higher order moments in the spatial profile of the birefringence functions Δ^i produce helicity-correlated differences in the higher order moments of the electron beam profile. The next higher order moment is spot size. The effects that helicity-correlated spot size differences can have on the physics asymmetry measurement will be discussed in Section 2.5.1.

Because of its natural birefringence, an improperly aligned Pockels cell is capable of

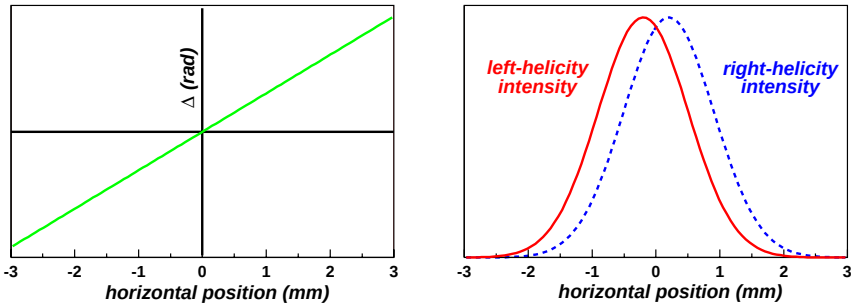

Figure 2.9: Effects of birefringence gradients. A linear gradient in a birefringence Δ is shown on the left. On the right, the effect such a birefringence gradient might have on right- and left-helicity beams is shown. The birefringence gradient will produce a linear polarization asymmetry that varies with position. The photocathode's analyzing power will then transform this into a spatially varying charge asymmetry, which is the same as a helicity-correlated position difference.

generating helicity-correlated position differences by a means very similar to that just described. If a laser beam passes through a Pockels cell at an angle, different rays will travel different lengths through the cell, and will therefore experience more or less retardation. This then becomes exactly equivalent to the situation pictured in Figure 2.9. Likewise, a diverging or converging beam, even if it is perfectly aligned with the Pockels cell, can produce a similar effect, as different rays will again experience varying amounts of retardation depending on their location within the beam.

Photocathode Gradients

Gradients in the analyzing power of the photocathode can also produce spatial variations in the electron beam charge asymmetry given in Equation (2.7), by effectively making the PITA slopes m_{CP} and m_{PS} functions of position. This effect has been observed in special tests in the Gun Test Lab (GTL) facility (which is a scale mock-up of the entire polarized electron source, along with a reproduction of the first few meters of the injector), but the observations indicate that the effects are of only moderate importance, being able to produce helicity-correlated position differences on the order of 100 nm [75]. It should be noted, however, that such effects will be proportional to Δ_{CP} and Δ_{PS}.

Another photocathode gradient to consider is one not in the QE anisotropy, but in the QE itself. Given any helicity-correlated steering effects in the laser system (such as those to be discussed immediately below), so that right and left helicity beams hit the cathode at different spots, a QE gradient could cause unwanted helicity correlations in the electron

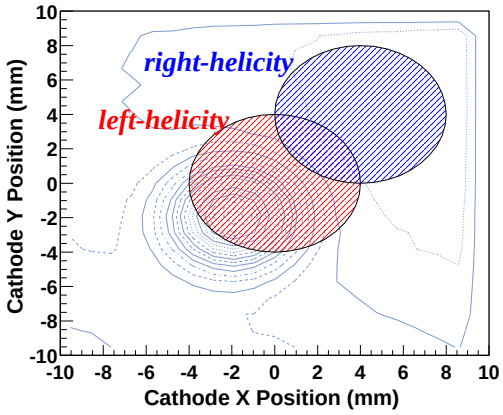

Figure 2.10: Effects of photocathode quantum efficiency (QE) gradients. The contour lines represent a supposed map of the photocathode's QE. The left-helicity laser beam is shown hitting a spot on the photocathode with large QE gradients, while the right-helicity beam is hitting a spot with small QE gradients. Such a scenario will generate a charge asymmetry in the emitted electron beam.

beam, most noticeably in charge and position,[3] but to a lesser extent in spotsize as well. All such effects will tend to be "averaged out" according to how much of the cathode the laser spot illuminates, which for the current setup is roughly half. Nonetheless, some effects likely attributable to this scenario, which is sketched in Figure 2.10, have been observed, and are referred to briefly in Section 4.5.

Pockels Cell Lensing

It is known that Pockels cells can, to some extent, act as imperfect lenses, focusing and steering the light that passes through them. What's more, these steering and lensing effects can be different depending on the voltage at which the Pockels cells are operated (in general they become more apparent at higher voltage, and depend on the sign of the voltage). This can have dire consequences, since the CP cell, in particular, continually flips between ± 2600 V. Since this voltage determines the helicity of the electron beam, the CP cell is therefore capable of introducing helicity-correlated steering effects into the source system.

Exhaustive studies of Pockels cell lensing effects were never carried out at SLAC, mainly because they had been studied in detail in previous parity violation experiments [78] employing polarization optics systems on which the E-158 optics system has been modeled. These studies suggested a simple and straightforward means of suppressing such lensing

[3] The position of the laser beam on the photocathode determines the position of the resulting electron beam, so that helicity-correlated laser beam position differences inevitably show up in the electron beam. However, a photocathode QE gradient would "muddy" the neat correspondence.

effects — image the Pockels cell onto the photocathode. Observations in the E-158 optics system seem to indicate that, uncorrected for, Pockels cell lensing effects can result in helicity-correlated position differences comparable to, if not greater than, those produced by birefringence gradients. Imaging the CP cell onto the photocathode has reduced the electron beam's sensitivity to lensing effects by a factor of 100 or more.

Electronic Cross Talk

Electronic cross talk refers to any unwanted communication of the helicity state of the electron beam to the data acquisition (DAQ) system. It is of course necessary to transmit the helicity state to the DAQ at least once, so that it can be recorded for data analysis. However, special measures (described in the next section) can be taken to minimize the chances of unwanted transmission of the helicity state.

Electronic cross talk can stem from the CP cell, whose perfectly helicity-correlated voltage signal of ± 2600 V can act like an antenna, broadcasting the helicity for outside "listeners" (other devices' signals) to pick up. They can also arise from ground loops, which surreptitiously connect the signals of other devices to the helicity controller. However it arises, electronic cross talk can result in a mismeasurement of the detector asymmetry and/or helicity-correlated beam differences. Because of the measures taken both to minimize the chances of it ever occurring as well as to provide cancellation of whatever systematic effects it may cause, electron cross talk is not expected to have a major impact on the experiment.

2.2.5 Passive Systematics Minimization

The previous sections discussed how helicity-correlated systematics can creep into parameters (such as intensity, position, angle, spot size, and in principle even more) of the electron beam, summarizing the formal discussions provided by Ref. [75]. As was seen in Section 1.3, such helicity-correlations cause false asymmetries (i.e., asymmetries caused purely by geometrical effects, and *not* new physical processes) to be measured by the detector, thus obscuring the true underlying parity-violating asmmetry A_{PV}. By setting a limit on the systematic uncertainty of A_{beam} at 4 ppb, where A_{beam} was defined to be the total false asymmetry caused by helicity-correlations in the beam parameters, the same section moti-

Systematic Source	*Charge*	*Energy*	*Position*	*Spot size*
Residual birefringence	1000 ppm	10 MeV	—	—
Birefringence gradients	—	—	1 μm	1 μm
Beam misalignment, divergence	—	—	1 μm	1 μm
Pockels cell lensing	—	—	1 μm	100 nm
Cathode gradients	10 ppm	100 keV	100 nm	100 nm
Electronic cross talk	1 ppm	10 keV	1 nm	1 nm
TOTAL	$\sim$1000 ppm	$\sim$10 MeV	$\sim$2 μm	$\sim$1 μm

Table 2.2: Summary of the potential sources of beam systematics. The estimates for the sizes of all effects are only approximate (i.e., order of magnitude), and are intended to be representative of the sizes of the beam asymmetries one would observe if no minimization measures (such as those discussed in Section 2.2.5) were taken.

vated the following upper limits for the sizes of beam systematics:

$$\begin{aligned} A_Q &< 200 \text{ ppb} & \Delta x, \Delta y &< 10 \text{ nm} \\ \Delta E &< 2 \text{ keV} & \Delta x', \Delta y' &< 0.25 \text{ nrad} \end{aligned} \tag{2.10}$$

These limits represent extraordinary levels of parity between left and right helicity beams. For instance, since the beam energy is nearly 50 GeV, a 2 keV limit on the right–left energy difference corresponds to requiring that the energy of the right beam be the same as that of the left beam to within 40 parts out of a billion.

Meeting the limits put forth in Equation (2.10) requires special attention to all potential sources of systematics discussed in Section 2.2.4. All of the things one does with the source optics setup before ever shining light on the photocathode and generating an electron beam are known as passive systematics minimization, as opposed to active systematics minimization, which is discussed in the next section.

If nothing was done to optimize the source optics system, one could expect to observe charge asymmetries in the electron beam on the order of 1000 ppm and right–left position differences[4] of roughly 2 μm, as indicated in Table 2.2. Because beam loading effects couple energy jitter to intensity jitter at roughly the 5% level, energy differences of $5 - 10$ MeV would also be observed. Beam helicity correlations of this size are clearly unacceptable, given the limits of Equation (2.10), and must be reduced.

The Pockels cells used for the polarization optics are selected for the uniformity of their

[4]For the purposes of this discussion, the "right–left" in "right–left ___ difference" (where "___" can refer to any beam parameter) will be dropped when its absence is judged not to cause confusion.

residual birefringence. Multiple pairs of Pockels cells were studied, and those which generated the smallest helicity-correlated position differences in the laser light they transmitted were picked for further commissioning. Because residual birefringence gradients are typically smooth and gradual, minimizing the spot size of the laser beam as it passes through optics suspected of carrying them is generally considered a good idea, since this reduces the sensitivity to their effects (by sampling a smaller area over which it may vary). However, this conflicts with the requirement that the laser be well collimated when passing through birefringent optics, since reducing the beam's diameter tends to increase its divergence.

An improperly aligned Pockels cell or an uncollimated laser beam can mimic gradients in residual birefringence. Therefore, when installing the Pockels cells, great care is taken to ensure that they are perfectly aligned with the beam. The exact procedure by which this is accomplished is documented in the SLAC source manual. In addition, three lenses are placed throughout the system (these have their focal lengths clearly labeled in Figure 2.2) in order to control the beam's spot size and divergence at the polarization optics. A safe compromise between the contradictory goals of reducing the beam's spot size *and* divergence is found by bringing the beam through a gentle focus at the CP and PS cells, with the optimal beam width at the cells being $\sim$1 mm.

Properly setting the voltages of the CP and PS cells provides the most important passive means of minimizing systematics. Before laser light is ever brought to the photocathode, the voltages of the CP and PS cell are scanned in order to find the best set that maximizes the so-called extinction ratio R, defined in Equation (2.1). This procedure effectively finds the best set of voltages that compensate for the asymmetric phase shifts Δ^0_{CP} and Δ^0_{PS} caused by the Pockels cells' residual birefringence (see Equation (2.8) and related discussion). This maximizes the beam's circular polarization just after the PS cell. However, perfect circular polarization is wanted at the photocathode. In essence, all of the various Δ^i variables in Equation (2.9) must be found.

The practical way of accomplishing this is to measure the PITA slopes m_{CP} and m_{PS}. Once the PITA slopes are known, they can be used to adjust Δ_{CP} and Δ_{PS} in order to null whatever charge asymmetry is measured in the electron beam. The PITA slopes are measured once the laser has been brought to the photocathode and electrons are being accelerated at least part-way down the linac. A toroid (a beam measurement device sensitive to charge, described in Section 2.4.1) is used to monitor the electron beam charge

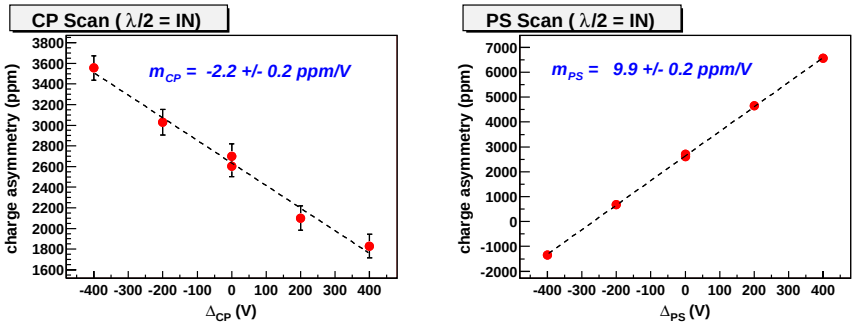

Figure 2.11: Example PITA scans. These data were taken with the half-wave plate ($\lambda/2$) inserted in the path of the laser beam. The charge asymmetry is plotted as either Δ_{CP} or Δ_{PS} is varied. The PITA slopes m_{CP} and m_{PS} are used to compensate for phase shifts downstream of the polarization optics, resulting in near-perfect circular polarization at the photocathode.

asymmetry while Δ_{CP} and Δ_{PS} are varied. An example so-called PITA scan is shown in Figure 2.11. From data like these, m_{CP} and m_{PS} are extracted. Finally, Δ_{CP} and Δ_{PS} are adjusted in order to null the charge asymmetry measured at what *were* the nominal operative voltages. This compensates for residual birefringence in optics downstream of the PS cell, most notably that of the cathode vacuum window at the end of the cathode diagnostics bench (which has unavoidable stress-induced birefringence).

When correcting for measured charge asymmetries using the PITA slopes, one in theory has freedom to decide which lever, either Δ_{CP} or Δ_{PS}, to use. There exists a line in voltage space, illustrated in Figure 2.12, upon which the charge asymmetry is zero. The slope of this line is $s = -m_{\mathrm{CP}}/m_{\mathrm{PS}}$. If a non-zero charge asymmetry A_0 is measured at some nominal operating point (which can be taken as the origin, and so denoted by the $_0$ subscript), it means that this operating point is not on this line. The most efficient change to Δ_{CP} and Δ_{PS} will be that which brings one back to the $A_Q = 0$ line via the shortest route possible in voltage space:

$$\begin{aligned} &\min \left| \left(\Delta_{CP},\ -\frac{A_0}{m_{\mathrm{PS}}} + s \cdot \Delta_{CP} \right) \right| \\ &\Longrightarrow \begin{cases} \Delta_{CP} = \dfrac{-A_0 \cdot m_{\mathrm{CP}}}{m_{\mathrm{CP}}^2 + m_{\mathrm{PS}}^2} \\ \Delta_{PS} = \dfrac{-A_0 \cdot m_{\mathrm{PS}}}{m_{\mathrm{CP}}^2 + m_{\mathrm{PS}}^2} \end{cases} \end{aligned} \tag{2.11}$$

Here the fact $m_{\mathrm{CP}} \cdot \Delta_{CP} + m_{\mathrm{PS}} \cdot \Delta_{PS} = -A_0$ has been used.

Finally, the means by which the photocathode is imaged onto the CP cell has already been described. This greatly reduces the beam's sensitivity to Pockels cell lensing effects.

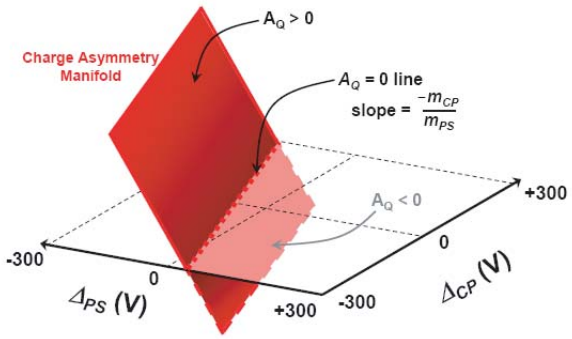

Figure 2.12: Charge asymmetry vs. Δ_{CP} and Δ_{PS}. There exists a line in voltage space where $A_Q = 0$. If a nonzero charge asymmetry is measured, changing Δ_{CP} and Δ_{PS} according to Equation (2.11) will move the charge asymmetry along a manifold back to the $A_Q = 0$ line.

Helicity Selection and Timeslots

One of the key requirements of the electron source is that it be able to rapidly switch between helicity states in order to minimize sensitivity to slow drifts. If not done properly, however, rapid helicity switching can add unwanted noise to the measurements or even generate false asymmetries via electronic cross talk. Measures must be taken in order to ensure that neither of these possibilities result.

The helicity of each pulse is controlled by the Polarization Monitor (PMON) system. PMON uses a 33-bit shift register algorithm [79] to generate two pseudo-random helicity states, either 00, 01, 10, or 11. It then briefly halts using the pseudo-random bit shift algorithm and automatically generates the complements of these helicity states. It then recommences with the bit shift algorithm to generate two more pseudo-random helicity states, and so on and so forth. It therefore operates in modified pair-wise pseudo-random mode:

$$h_1\, h_2\, \overline{h_1}\, \overline{h_2} \quad h_3\, h_4\, \overline{h_3}\, \overline{h_4} \quad h_5\, h_6\, \overline{h_5}\, \overline{h_6} \quad \ldots \tag{2.12}$$

Here the subscripts are used to denote the so-called *pulse pairs*, which consist of two pulses with opposite helicities. Since the maximum repetition rate for the source is 120 Hz, it is clear from this equation that the maximum rate at which pairs may be generated is 30 Hz. This shields out the considerable 60 Hz line noise that would otherwise add a large amount of random noise into the experiment, ruining the statistical precision.[5] If 0 is used

[5]This line noise, however, would not bias the measurement in any way, since it would have no knowledge of helicity and would therefore tend to cancel out when taking the average.

to denote a *left* helicity state, and 1 is used for a *right* state, and these are denoted by L and R, respectively, then a possible helicity sequence would be:

$$\overset{\text{pair}}{\overbrace{L\,R\,R\,L}}\,R\,R\,L\,L\,R\,L\,L\,R\,\ldots \qquad (2.13)$$

As can be seen, the pairs are interleaved, and are said to inhabit two "timeslots," each timeslot effectively a completely separate experiment, since, in general, it sees a consistently different phase of the 60 Hz line noise. Timeslots are an unavoidable but, as it turns out, a very useful feature of the experiment. Much more about them will be discussed in Sections 3.4 and 4.5.

In order to guard against the risks of electronic cross talk, all helicity bits are delayed by one pulse and RF modulated prior to broadcast. In addition, ground loops are carefully avoided in all data acquisition systems by using fiber optics (instead of copper cables) to connect the various VME crates, which house the electronics used by the data acquisition system. These measures greatly reduce the risks of false asymmetries arising from electronic pick-up. Any effects that do show up can be cancelled by the insertion of either of the half-wave plates.

2.2.6 Helicity-Correlated Feedback Loops

Passive systematics minimization can usually bring beam asymmetries down to the level of roughly 100 ppm for charge and a few 100 nm for position differences. This is an improvement of roughly one order of magnitude from where the systematics started, but a huge amount of further reduction is still needed in order to reach the goals put forth in Equation (2.10). This further reduction is achieved through the use of active feedbacks.

The source uses three active feedback loops on helicity-correlated measurements: an intensity asymmetry feedback, a right–left position difference feedback, and a source phase feedback. Each feedback works along the same principle: integrate measurements from a particular device for a given number of pairs to average out statistical fluctuations, then induce an asymmetry using a particular control device to compensate for whatever was

measured. The procedure can be written out as:

$$\begin{aligned} A^1_{ind} &= \ 0 \\ &\vdots \\ A^i_{ind} &= \ A^{i-1}_{ind} - g \cdot A^{i-1}_{obs} \end{aligned} \tag{2.14}$$

where

$$A^i_{obs} = A^i_{ind} + A^i_{syst} + A^i_{stat} \tag{2.15}$$

Here g is the gain of the feedback ($g > 0$), A_{ind} is the asymmetry induced using the control device, A_{obs} is the beam asymmetry observed in the measurement device, A_{syst} is some underlying beam asymmetry caused by any of the sources of systematics discussed in Section 2.2.4, and A_{stat} is the random fluctuation measured due to beam jitter. If the feedback loop integrates N pairs at each time interval i, the statistical fluctuations A_{stat} will decrease as $\sigma/\sqrt{N}$, where σ is the pulse-to-pulse beam jitter for whatever beam parameter is being considered.

Following the prescription above cancels out the beam systematic A_{syst}, so that the measured beam asymmetry A_{obs} tends towards zero. What's more, the measured beam asymmetry can tend towards zero at a faster rate than normal statistical scaling would seem to allow. This is because the statistical error bars A^i_{stat} are all arranged to cancel one another out, with the exception of the very last one, when taking the average of all observed beam asymmetries A^i_{obs}. Therefore only one error bar enters into the error propagation formula.[6] In this way, the error bar on the final averaged asymmetry will scale as $1/N$, and not as $1/\sqrt{N}$, where N is the number of time intervals.

The intensity asymmetry (IA) feedback uses the IA Pockels cell on the helicity control bench to apply a helicity-correlated phase shift to the laser beam, rotating its polarization in a helicity-correlated manner. The cleanup polarizer then transforms this helicity-correlated phase shift into an intensity asymmetry. In this way, the IA cell and cleanup polarizer combination is capable of producing charge asymmetries in the electron beam as large as a few percent, limited by the IA cell's HV power supply (currently able to produce voltages up to 750 V).

[6] The correlation between the random error bars A^i_{stat}, leading towards super-statistical $1/N$ scaling, is only perfect for unitary gains ($g = 1$). Non-unitary gains will slow the rate of scaling.

The position difference (POS) feedback uses a mirror mounted on a piezostack on the helicity control bench as its control device. This piezostack is connected to three pairs of voltages that can manipulate the mirror's orientation on a pulse-to-pulse basis. The so-called piezomirror can therefore induce helicity-correlated position differences on the laser beam. The piezomirror can only generate small tilts, though, roughly ± 50 μrad (the voltages required are only ± 10 V). The point-to-point imaging of the photocathode onto the CP cell leaves an effective lever arm of approximately 50 cm for the piezomirror, giving it a dynamic range of ~ 50 μm at the photocathode.

Both the IA feedback and the POS feedback use measurement devices situated in an early region of the linac where the beam has only been accelerated to 1.2 GeV. Feeding back on the beam early, before it has been accelerated to very high energies, will hopefully suppress helicity correlations in beam parameters over which the source has no direct control. For instance, beam loading couples intensity to energy. Nulling the intensity asymmetry in the electron beam early in the machine should then also null the energy asymmetry, even though no direct control over the energy is being exercised.

For much the same reason, a phase feedback is employed, which continually adjusts Δ_{CP} and Δ_{PS} according to Equation (2.11) in order to compensate for whatever charge asymmetry correction the IA loop has had to make. This keeps the average correction made by the IA loop small, by maintaining near–zero residual linear polarization asymmetries at the photocathode. In effect, whereas the IA feedback merely addresses the visible result of a problem (a charge asymmetry in the electron beam, for instance), the phase feedback goes after the root of the problem itself (a linear polarization asymmetry at the photocathode, caused by residual birefringence). Photocathode gradients offer another example of a class of systematic that benefits from the phase feedback's habit of adjusting the CP and PS cell voltages to null linear polarization asymmetries at the photocathode.

For each feedback loop, the number of pairs N integrated at each time interval is determined by the dynamic range of the particular control device used and by the beam jitter. For instance, the intensity asymmetry loop feeds back on a toroid that typically measures pulse-to-pulse intensity jitter of 0.5%, which in turn produces jitter in the asymmetry measurements of $0.5\%/\sqrt{2} = 0.35\%$. Since this is far below the dynamic range of the IA Pockels cell, which can easily produce charge asymmetries (with the cleanup polarizer's help, of course) as large as 3%, N need not be very large. Consequently, $N = 800$ pairs for

the IA loop, corresponding to typical corrections on the order of 100 ppm.

On the other hand, the POS loop only has a dynamic range of roughly ± 1 μm at the 1.2 GeV region, due to adiabatic damping effects. [7] The pulse-to-pulse beam position jitter is typically on the order of 10 μm. In order to avoid driving the piezomirror too hard, voltages on the order of only 1 V are desired. Thus, the number of pairs the POS loop should integrate at each time interval is $N = (10 \times 10)^2 = 10$k pairs. The exact value of N was periodically adjusted during data taking to compensate for changes in jitter and betatron phase adjustments that the linac operators occasionally made, but N always stayed within the range of 10k$-$20k pairs.

The phase feedback has a very long time constant. At every step, it averages the previous 30 IA loop corrections, meaning that it effectively averages $N = 30 \times 800 = 24$ kpairs. The phase feedback is designed to compensate for slow drifts in the system, such as slow birefringence drifts. In theory it is not strictly needed, yet it provides a valuable degree of automation and flexibility.

The gain of each feedback is determined by the stability of its calibration. For instance, the IA feedback can be calibrated very precisely simply by measuring the half-wave voltage V_{hw} of the IA Pockels cell, which does not fluctuate over time. Consequently, the gain of the IA feedback is unity. The POS loop calibration, however, is extremely sensitive to small changes in a number of system parameters, including the "beam tune" of the injector and early linac and slow drifts in the phasing of the BPM's.[8] For this reason, the POS feedback is typically run with a gain of 0.25. In fact, because random phase drifts in the linac were observed to be leading towards natural cancellations of position systematics in the beam by the time it had reached the target, the POS feedback was actually found to be unneccessary. As keeping it properly calibrated required careful attention, and if not calibrated it could actually harm the experiment by producing gross systematics, the POS loop was turned off for large portions of the data.

[7] Because of adiabatic damping, beam emittances scale inversely proportional to the beam energy, and transverse beam sizes and divergences scale inversely to the square root of the beam energy [80].

[8] The "beam tune" refers to the betatron phase advance from the source to the 1.2 GeV region where the POS loop's measurement devices are located. BPM phases are discussed in Section 2.4.2, and, if not carefully monitored and controlled, can drift due to day–night temperature fluctuations, among other things.

Systematic Source		*Passive Means*	*Active Means*
Residual birefringence		Null A_Q with Δ_{CP}, Δ_{PS}; beam alignment at CP, PS	IA, phase
Birefringence gradients		Select CP, PS for uniformity; minimize beam spot at CP, PS	POS
Beam misalignment, divergence		Beam alignment, collimation at CP, PS	POS
Pockels cell lensing		Image photocathode onto CP	POS
Cathode gradients		Keep linear polarization asymmetries small	phase
Electronic cross talk		Delay helicity bits; rf modulate helicity signals	—
TOTAL	*Charge asym:*	100 ppm	100 ppb
	Pos. diffs:	100 nm	5 nm

Table 2.3: Summary of the passive and active means of beam systematics minimization. The "active means" column lists which feedback (IA, POS, or phase) specifically addresses which potential systematic source. The "total" row lists the approximate levels of beam asymmetries that are to be expected given either passive or passive + active minimization techniques.

2.2.7 Systematics Minimization Summary

Section 2.2.4 described the various ways in which helicity correlations can arise in the electron beam parameters. Charge asymmetries mainly arise when residual birefringences induce linear polarization asymmetries at the photocathode. The photocathode's analyzing power then transforms this linear polarization asymmetry into a charge asymmetry. The dominant source of position differences in the electron beam arise mainly from birefringence gradients and poor beam alignment and/or large divergence through birefringent optics, in particular the CP and PS Pockels cells. Gradients in the photocathode's analyzing power can also play a significant role. As described in the Section 2.2.5, many precautions in the initial setup of the source optics must be taken in order to reduce sensitivity to these potential sources of systematics. Even after these precautions are taken, however, further suppression is needed from active feedback loops. The descriptions contained in the previous sections, regarding the minimization and control of helicity-correlated systematics in the electron beam, are summarized in Table 2.3.

2.2.8 Systematics Reversals

Devices that reverse the experiment's sensitivity to certain classes of systematics, while leaving its sensitivity to others unchanged, are often referred to as systematics reversals. There are two basic types: ones that change the sign of the physics asymmetry A_{PV}, and ones that don't. In addition, each reversal may flip the sign of various systematic contributions A_{sys}. Reversals that change the sign of the physics asymmetry can be termed *physics asymmetry reversals*, while those that do not can be termed *false asymmetry reversals*. Their respective effects can be tabulated as follows:

	Measured asymmetry	
	Physics reversal	*False reversal*
Reversal state 1:	$A_{PV} + A^1_{sys} + A^2_{sys}$	$A_{PV} + A^1_{sys} + A^2_{sys}$
Reversal state 2:	$-A_{PV} - A^1_{sys} + A^2_{sys}$	$A_{PV} + A^1_{sys} - A^2_{sys}$
$\frac{1}{2} \times$ *TOTAL:*	A^2_{sys}	$A_{PV} + A^1_{sys}$
$\frac{1}{2} \times$ *DIFFERENCE:*	$A_{PV} + A^1_{sys}$	A^2_{sys}

(2.16)

where A^1_{sys} is a systematic that flips sign like a physics asymmetry, and A^2_{sys} is one that is insensitive to the actual beam helicity, and so does not flip sign. The above equations show that by taking roughly equal portions of data in both states of a reversal device, one can cancel out at least some classes of systematics. It is useful to have as many systematics reversals as possible, since different systematics sources will in general respond differently to different kinds of reversals. The experiment implemented three types of systematics reversals: a half-wave plate, the asymmetry inverter (AI), and an energy flip.

The half-wave plates, along with the asymmetry inverter, have already been described in Section 2.2.2. Inserting a half-wave plate into the path of the laser beam between the polarization optics and the photocathode simply reverses the relationship between the helicity bit (0 or 1), and by extension the voltages applied to the polarization Pockels cells, and the actual orientation of the beam's elliptical polarization at the photocathode. This, in turn, can cancel out effects like those called by Pockels cell lensing and electronic cross talk. It can also provide imperfect cancellation of the position and spot size helicity correlations that result from birefringent gradients. There are actually two half-wave plates downstream of the CP and PS cells. The one furthest downstream, located on the cath-

ode diagnostics bench just before the vacuum window entrance to the photocathode gun, provides the greatest level of cancellation for the greatest number of classes of systematics, and is therefore the one used.

The asymmetry inverter, located on the helicity control bench, can essentially reverse the trajectory of all the light rays contained in the laser beam. It does *not* affect the laser beam's polarization in any way, and so does nothing to the actual physical helicity of the electron beam. Therefore it is a false asymmetry reversal. The asymmetry inverter can provide cancellation for helicity-correlated position differences in the electron beam, caused by effects like birefringence gradients, Pockels cell lensing, and beam divergence through birefringent optics.

The electron beam must bend after the end of the linac in order to be directed into End Station A (where the main experimental apparatus, including target and detector, is located). While the beam is bending, the electrons' spins precess about the dipole magnetic field. At 45 GeV, the spins precess by 14π, leaving their longitudinal polarization at the target the same as it was in the linac. At 48 GeV, however, the spins precess by 15π, reversing their longitudinal polarization at the target. This is an especially powerful type of reversal, since it does not alter the source setup in any way. Consequently, all source–related systematics should receive at least partial cancellation. In the nomenclature of Equation (2.16), the energy flip is a physics asymmetry reversal wherein the greatest number of systematics sources as possible are of the A_{sys}^2 type. Averaging data taken at both energy states, multiplying the 48 GeV data by -1 to take into account the minus sign that is picked up due to the spin precession, should therefore provide a large number of systematics cancellations. Alternatively, if the same value for the physics asymmetry is measured for 45 GeV data as for 48 GeV data (correcting for the $\sim$6 ppb/GeV asymmetry change with energy), then an upper limit can be placed on the systematic contributions, at least at the level of statistics. The results of all of the systematics reversals will be discussed in Section 4.4.

2.3 Linac

Electrons are drawn from the photocathode with longitudinal polarizations typically $\sim$80%. The beam must then be made ready for acceleration in the linac. The first step is called

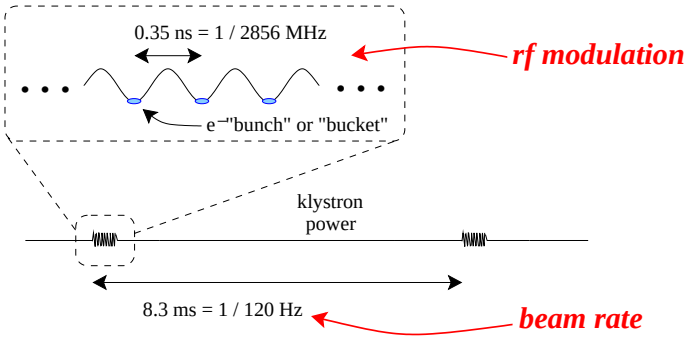

Figure 2.13: Linac rf time–structure. Each beam pulse consists of hundreds of electron bunches, separated by distances determined by the rf structure of the klystron power.

electron capture, which happens in the bunchers and prebunchers. After the bunchers, the beam has achieved roughly 250 keV of (kinetic) energy and is appropriately synchronized with the frequency of the rf power provided by the machine klystrons (2856 MHz, in the S-band). In other words, the electrons ride along in the troughs of the rf waves. The beam is said to be "bunched," or "rf modulated," with the individual bunches of electrons often referred to as *buckets*. The bunched beam is then ready for injection into the linac. In the injector, the beam gains 40 MeV of energy. By the end of the linac, it will have been accelerated up to a final energy near 50 GeV. Capture, injection, and acceleration do not significantly degrade the beam polarization.

The relevant timescales for linac machine operation are illustrated in Figure 2.13. A *pulse* refers to the $50-300$ ns of time in which the electrons are mainly distributed. Usually there are $1-6\times10^{11}$ electrons per pulse. It is sometimes useful to think in terms of physical length as opposed to time. By the end of the linac, for instance, a 300 ns pulse will be over 270 feet long, weaving and winding its way through apertures and beam pipes. Again, the rf modulation means that it will actually be composed of over 850 individual bunches of electrons, each separated by roughly 4 inches.

The *beam rate* is defined to be the inverse of the time in between pulses (from the start of one pulse to the start of the next). This rate is determined by the rate at which the source laser's flash lamps fire, as well as by the rate at which the klystrons deliver accelerating power. The beam rate can vary between 0.5 Hz and 120 Hz. Physics data, however, is taken almost exclusively at 60 Hz or 120 Hz. It should be noted, however, that the experiment was designed to be able to run concurrently with the BaBar experiment, which requires two entirely different beams during a PEP fill, one containing medium-energy

(9 GeV) electrons, and another containing low-energy (3.1 GeV) positrons. These electrons and positrons are extracted from the linac and stored in the SLAC PEP-II storage rings. The pulses containing these beams are collectively referred to as PEP pulses. Although the machine was upgraded in order to support so-called "interlaced" operation, wherein the linac could rapidly switch between the very different beam lattices necessary for the acceleration of either the E-158 beam or the PEP beam (and then either the PEP positron beam or the PEP electron beam), rising electricity costs have made running the machine concurrently with PEP too costly for the laboratory to afford. Hence, only Run I consisted of data taken concurrent with PEP operation. Both Runs II and III were taken in single beam mode.

As the electron beam passes through an accelerator cavity, it excites resonant cavity modes (this occurs whether or not the cavity is supplied with rf power) [80, 81]. Creating these electromagnetic fields costs the beam energy; this is called *beam loading*. Beam loading couples intensity to energy, since the total beam loading energy losses are proportional to the peak beam current. The total energy lost due to beam loading in a particular cavity is intimately related to the cavity's design and construction and depends, for instance, on its resistive wall losses, parameterized by its quality factor Q. For the constant gradient cavities employed at SLAC, the fractional beam loading is roughly 5%. This means that the intensity of the beam influences its energy at the 5% level. Hence intensity jitter of 0.5% will produce energy jitter of roughly 0.025%.

Beam loading energy losses generally increase over the course of the pulse, as the fields of the excited modes grow stronger as more bunches pass through the cavity. This plays a role in determining the beam's energy spread. Given a pulse with a flat intensity time profile, beam loading of 5% would cause the tail of the pulse to have 5% less energy than the head, leading to an energy spread value of 5%, where energy spread is defined to be $(E_{tail} - E_{head})/E_{head}$. Dispersion could then cause significant transverse position offsets in the tail of the pulse, leading towards poor beam delivery (such tail offsets can cause ion chamber trips in the accelerator and A-line) and possible systematic errors stemming from higher order effects in the asymmetry correction procedure. In the end, the pulse shaping capability afforded by the TOPS Pockels cell at the source allows for beam loading compensation. Thanks to this and other procedures performed by the machine operators (for instance, placing the electron bunches slightly ahead of the rf phase can help minimize

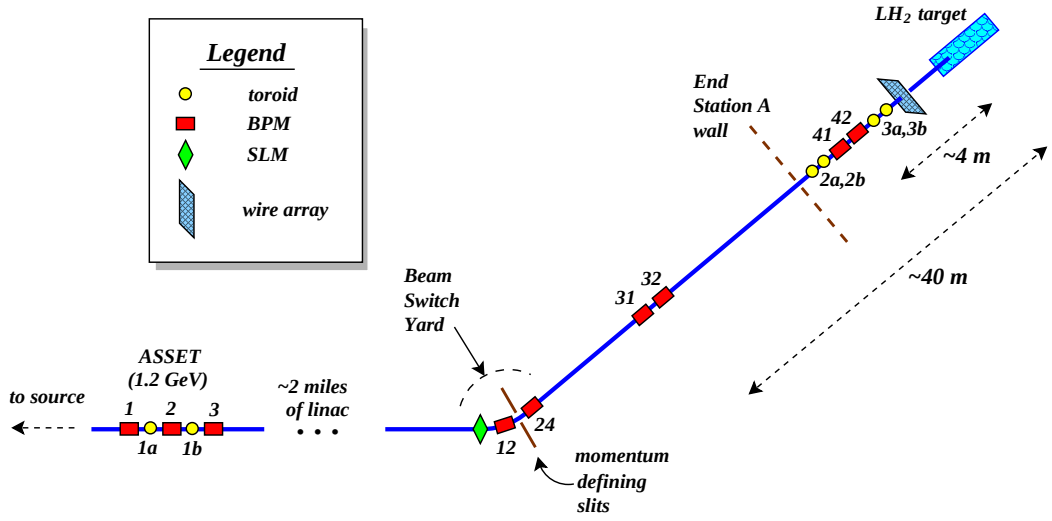

Figure 2.14: Location of beam monitoring devices.

the head–tail energy difference), energy spread is typically kept below 0.15% [82].

Towards the end of the linac, the beam encounters the dithering region. This consists of four horizontal and four vertical corrector magnets capable of imparting small (roughly 100 μrad) transverse deflections to the beam on a pulse-by-pulse basis. These magnets are used as part of the beam dithering procedure, described more fully in Section 3.2.3, which is used to measure the detector's sensitivity to the various beam parameters, so that false asymmetries due to beam helicity correlations can be subtracted out. The linac corrector magnets allow for the evaluation of the detector's sensitivities to transverse beam motions. At the end of the linac, a klystron's phase can be adjusted during in order to measure the detector's sensitivity to energy fluctuations. Adjusting the klystron's phase gives the experimenter the ability to change the energy of the beam by as much as $\pm$100 MeV on a pulse-by-pulse basis.

2.4 Beam Monitoring Devices

The experiment monitors all first-order and some second-order beam parameters in several locations, as shown in Figure 2.14. Toroids are sensitive only to charge, beam position monitors (usually abbreviated as BPM's) are used to indicate transverse beam position and, in areas of dispersion, beam energy, the synchrotron light monitor monitors both energy and transverse polarization, and the wire array measures the transverse position of the beam as well as higher order moments in its spatial profile.

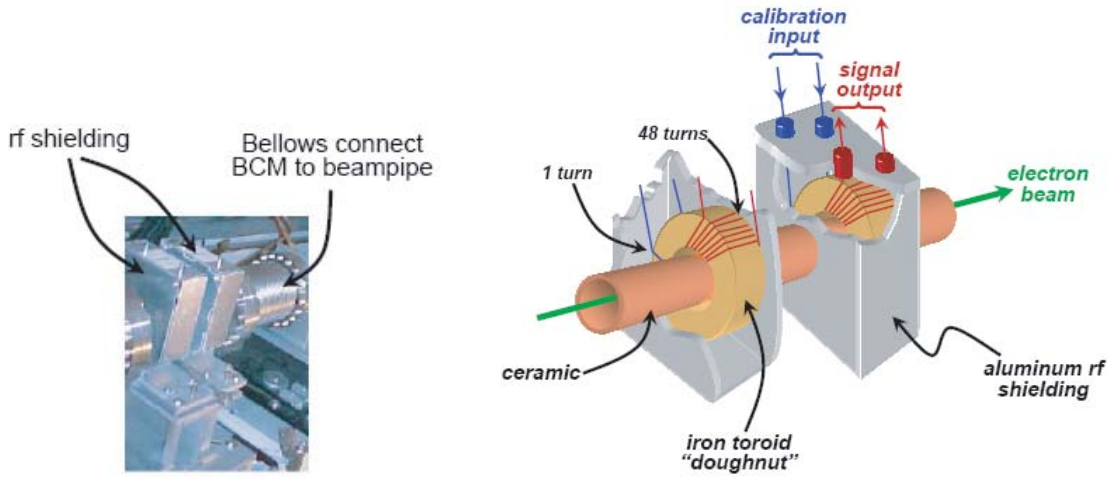

Figure 2.15: Beam current monitor (toroid) picture and schematic.

2.4.1 Toroids

Beam current monitors are referred to as toroids, due to their consisting of two iron doughnuts (toroids) surrounding the beamline, as shown schematically in Figure 2.15. Each toroid is wrapped with insulated copper wire, wound approximately 50 times. Ambient rf noise pick-up is reduced by an aluminum casing, which acts as a Faraday shield. The vacuum housing through which the beam passes, connected to the main beam line by flanges, is constructed of ceramic.

The toroids operate as current transformers, so that the electron beam passing through generates a pulse in the windings whose size is proportional to the beam current [81]. To maximize the signal-to-noise ratio, it is necessary to stretch the toroid signal out, from 300 ns (the beam pulse length) to something on the order of 1 ms. This is done by connecting the toroid windings in series to a resistor and a capacitor of appropriate size, forming an RLC circuit (typical toroid inductances are on the order of 20 mH). This produces a ringing signal that, after rectification, can be integrated in an analog-to-digital converter, as will be discussed further in Section 2.8.

In addition to the main signal windings, each toroid also has a single calibrator winding. Sending a known current through this calibrator winding simulates the beam, allowing for calibration of the toroid. Studies of the toroids' responses to a wide range of calibration currents reveals the toroid linearity to be better than 99.9%.

As can be seen in Figure 2.15, each beam current monitor actually houses two toroids,

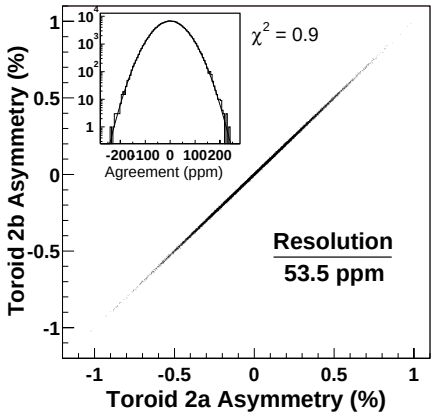

Figure 2.16: Typical toroid resolution plot. The charge asymmetry measurements of two neighboring toroids are compared, providing information as to each device's inherent resolution.

for redundancy. Subtracting the signals of the two neighboring toroids therefore removes beam jitter (which is highly correlated) and gives a measurement of the amount of random electronic noise (which is uncorrelated) in the two signals. This procedure yields the devices' resolution. The pulse-to-pulse toroid resolution σ_{res} is defined to be:

$$\sigma_{res} = \text{RMS of } \left(A_Q^1 - A_Q^2\right) \tag{2.17}$$

where A_Q^1 and A_Q^2 are the charge asymmetry measurements for toroid "1" and toroid "2," respectively. A typical resolution plot is shown in Figure 2.16. Resolutions of approximately 60 ppm are typically achieved. This can be compared to the typical beam intensity pulse-to-pulse jitter, which is roughly 5000 ppm.

The experiment utilizes three beam current monitors, each containing two toroids as mentioned above, for a total of six toroids. The toroids at ASSET are called toroids 1a and 1b. The toroids in the alcove sit a few meters upstream of the target. The two pairs of alcove toroids are called toroids 2a and 2b, and 3a and 3b. The distance between the two pairs is roughly 1 meter.

2.4.2 BPM's

The beam position monitors (BPM's) are devices comprised of three copper cavities, one cylindrical and two rectangular, through which the electron beam passes. An actual picture of one of the BPM's, along with a schematic drawing, is shown in Figure 2.17. The details behind the workings of the BPM's are somewhat complicated, with the mathematics rigor-

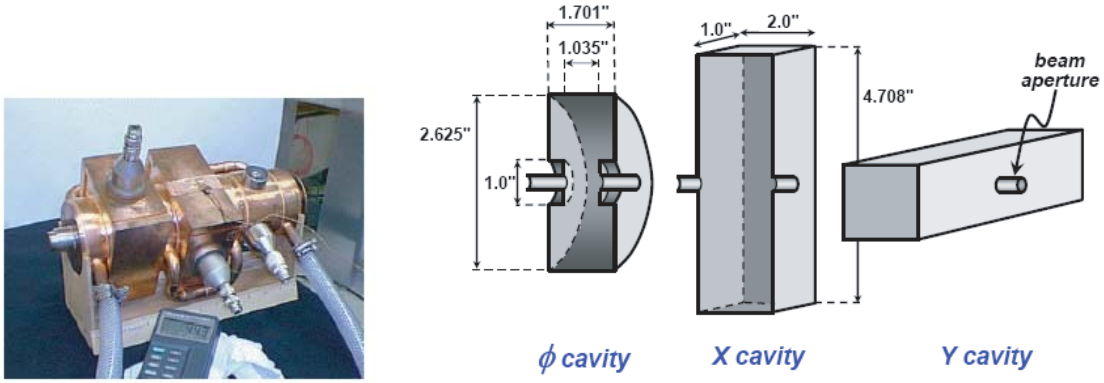

Figure 2.17: Beam position monitor (BPM) picture and schematic. Each BPM consists of three cavities, which provide sensitivity to charge, horizontal displacements, and vertical displacements.

ously derived in a number of outside sources [81, 83], but the general idea is easily grasped. As a charged particle beam passes through a copper cavity, it can excite TEM modes depending on the boundary conditions imposed by the cavity's geometry. The excited modes are known as *resonant modes*, and so the BPM's are often referred to as *resonant cavities*. These electric fields contained in the excited modes are detected by a built-in antenna, whose signal is transmitted by cable to processing electronics.

In every BPM, the cylindrical cavity has a resonant TM_{01} mode, also called the monopole mode. The pick-up voltage induced in the antenna by such a mode is, to first order, independent of the beam offset in the cavity, and sensitive only to the beam charge. This cavity is therefore referred to as the BPM's Q-cavity.

The two rectangular cavities have resonant TM_{210} or TM_{120} modes, depending on the orientation of the cavity. These modes are dipole modes, meaning that not only is the antenna's pick-up voltage proportional to the beam charge, but also to the beam's transverse displacement from the cavity's center. The TM_{210} and TM_{120} modes are sensitive to horizontal and vertical displacements, respectively, and are therefore known as x-dipole and y-dipole modes. These cavities are thus referred to as the BPM's X- and Y-cavities.

The pick-up voltage transmitted via coaxial cable to the BPM processing electronics, shown schematically in Figure 2.18, oscillates with some frequency ω. This frequency ω is determined not only by the natural frequency ω_{cav} of the cavity's resonant mode, but also by the frequency ω_{beam} at which the cavity, which acts like a damped oscillator, is being driven.[9] Ideally, the geometry of the cavity will be such that $\omega_{cav} = \omega_{beam} = \omega$. In this case

[9] As described in Section 2.3, the rf modulation of the beam means that $\nu_{beam} = \omega_{beam}/2\pi = 2856$ MHz.

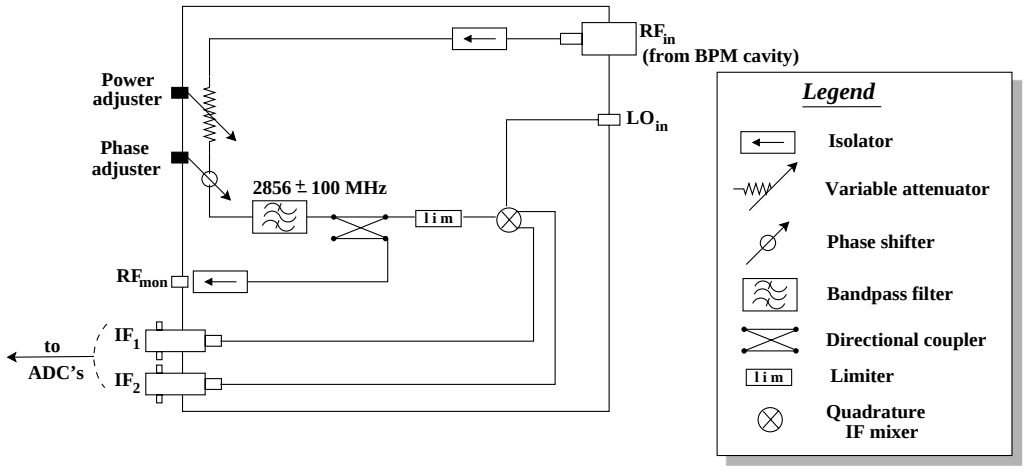

Figure 2.18: Schematic of BPM processor electronics.

the cavity is said to be "tuned." Tuning each cavity is accomplished by means of copper stubs whose position can be finely adjusted, thereby altering the cavity's geometry, and hence ω_{cav}, in controllable ways.

Each cavity signal decays with a time constant determined by the cavity's quality factor Q, which also determines the signal's bandwidth. On the one hand, a high Q is desirable because it lengthens the period for which the pick-up signal can be integrated, increasing the signal-to-noise ratio and thus the BPM's resolution. On the other hand, the Q has to be low enough in order to accommodate the inevitable "detuning" effects of temperature oscillations in the accelerator, which cause slow fluctuations in ω_{beam}. As a compromise, the quality factors are generally set (via the same tuning stubs mentioned above) to be $Q \approx 3000$, corresponding to a FWHM bandwidth of approximately 1 MHz.

The rf mixer in the BPM processing electronics removes the fast oscillations from the cavity signal by mixing it with a local oscillator signal possessing the same frequency ω_{beam} and phase ϕ_{beam} as the beam. This process produces two decaying, non-oscillatory outputs, whose amplitudes are proportional to $\cos(\phi_{cav} - \phi_{beam})$ and $\sin(\phi_{cav} - \phi_{beam})$, respectively, where ϕ_{cav} is the overall phase of the cavity signal. Both of these signals are then fed into integrating ADC's, to be described in Section 2.8. Normally each cavity is "phased" so that $\phi_{cav} = \phi_{beam}$, so that only one signal is needed in order to reconstruct all the information about the beam. For example, a properly tuned and phased X-cavity produces a signal that is directly proportional to $q \cdot x$, where q is the beam's charge and x is its horizontal offset from the cavity's center.

As can be seen in Figure 2.14, the experiment employs nine BPM's. Three of these are located in the 1.2 GeV ASSET region. These BPM's, known simply as BPM's 1, 2, and 3, are used as the measurement devices for the source position feedback. Two BPM's, called BPM's 12 and 24, are located in the Beam Switch Yard (BSY) where the beam goes through a bend in order to steer it into the A-line, which in turn leads to End Station A, where the target, spectrometer, and detectors are housed. The dispersion η in the BSY varies, but at BPM 12's position $\eta \approx -520$ mm, while at BPM 24's position $\eta \approx -410$ mm. Dispersion relates position to energy via the formula $(E - E_0)/E_0 = (x - x_0)/\eta$, where x_0 is the nominal reference trajectory for a particle of energy E_0. BPM's 12 and 24 can thus be used as energy measurements. Another pair of BPM's, 31 and 32, are located roughly 42 m upstream of the target, while yet another pair, 41 and 42, sit roughly 4 m upstream of the target, in the entrance to the End Station, known as the alcove. Because the more upstream pair possesses a long lever arm with respect to the target, whereas the more downstream pair possesses a short lever arm, measurements from these two pairs of BPM's can be used to determine both the angle and the position of the beam at the target.

Calibrating a BPM requires monitoring its signals, along with those coming from a device whose calibration is already known, while moving the beam in finite steps in both x and y. The wire array (described in Section 2.4.4) is used to calibrate BPM's 31, 32, 41, and 42. For the calibration of the other BPM's, nearby devices whose calibrations are known (such as stripline BPM's used by the machine operators but not recorded by the experiment's data acquisition system) are used. As mentioned in Section 1.3, the absolute calibration is not important as far as the overall asymmetry measurement is concerned, but it is useful for diagnostic purposes.

The resolutions of the BPM's are measured in an analogous manner as for the toroids. The measurements of neighboring BPM's are compared, thereby removing correlated beam jitter and leaving behind only random electronic noise. However, the procedure is a bit more complicated for BPM's than it is for toroids, since beam angle jitter will cause the two position measurements to disagree slightly (since even BPM's within a pair can be nearly 1.5 m apart). This is accounted for by using a third BPM, located as far from the two neighboring BPM's as possible, to remove angle jitter. As an example, the following

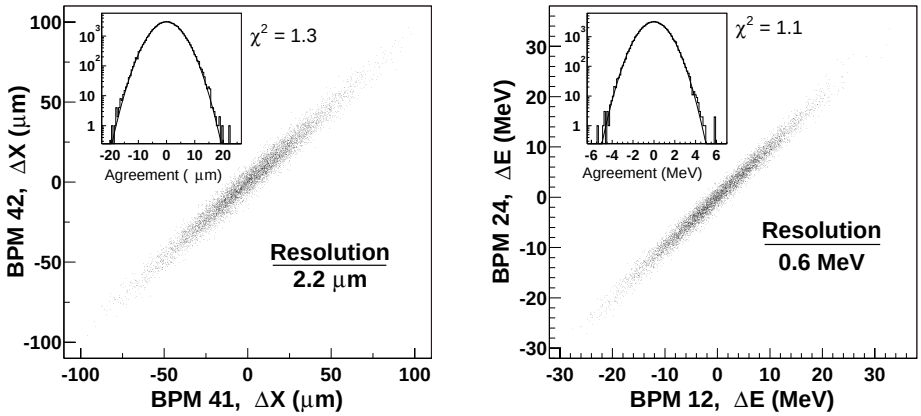

(a) (b)

Figure 2.19: Typical BPM resolution plots. The position measurements of two neighboring BPM's are compared, providing information as to each device's inherent resolution.

prescription is followed to find the pulse-to-pulse resolution σ_x of a BPM X-cavity signal:

$$\begin{aligned} \sigma_x &= \kappa \times \text{RMS of } \{\Delta x_1 - (\Delta x_2 - m \cdot d)\} \\ m &= (\Delta x_3 - \Delta x_1)/L \\ \kappa^{-1} &= 2\sqrt{1 + d/L + (d/L)^2} \end{aligned} \tag{2.18}$$

Here Δx_1 and Δx_2 are the position difference measurements for a given pulse pair for the two neighboring BPM's, d is the distance between the two neighboring BPM's, Δx_3 is the position difference measurement for the third BPM (used to take out angle jitter), and L is the distance between the first and third BPM's.

The above prescription works for all the BPM's except for BPM's 12 and 24, since these are located in regions of dispersion. Since the dispersions are somewhat different for these two BPM's, simply subtracting the measurement of one BPM from that of the other does not fully remove energy jitter. In fact, it doesn't even fully remove position and angle jitter, but may tend to *enhance* it, since the betatron function's phase advance is approximately 180° between the two energy BPM's.[10] As a result, a full regression procedure using the other A-line BPM's is needed to completely remove position and angle jitter, revealing the true resolution of the energy BPM's. Nevertheless, an upper limit can be placed on the resolution simply by comparing the right–left energy difference measurements taken by

[10] In fact, the momentum–defining slits lay between BPM's 12 and 24, so beam losses on the slits could further degrade agreement.

BPM's 12 and 24. In practice, this upper limit is usually very close (within 20%) to the result yielded by the full regression procedure, but can sometimes be high by as much as a factor of two.

The cavity signals must be attenuated before reaching the rf mixer, which begins to saturate at 300 mV. In order to set the attenuation, the beam is offset by ± 1 mm (which corresponds roughly to the maximum beam motions expected) and the attenuation is adjusted so that the BPM processor electronics remain well within their linear range. This defines the BPM's dynamic range. Having a wide dynamic range (i.e., high attenuation) is comforting, but it tends to worsen the BPM's resolution by decreasing the signal-to-noise ratio in the BPM's ADC, which is limited by fixed pedestal noise (see Section 2.8). A dynamic range of ± 1 mm provides adequate resolution, while still allowing for reasonable drifts in beam position. Linearity of better than 99% is maintained by implementing an analysis cut that flags all data where the beam has drifted past the limits imposed by the dynamic range. Typical resolution plots for two different pairs of BPM's are shown in Figure 2.19. Resolutions of 1 to 3 μm are typical for position, while resolutions of 1 to 2 MeV are typical for energy. These should be compared to typical beam jitter values of $\sim$50 μm for position and $\sim$10 MeV for energy.

2.4.3 Synchrotron Light Monitor

The synchrotron light monitor, shown in Figure 2.20, is situated in the BSY near a dipole magnet. As the beam particles bend in the dipole's magnetic field, every second they lose the following amount of energy due to synchrotron radiation [80]:

$$\begin{aligned} P_\gamma\,[\mathrm{GeV/sec}] &= 4.23 \times 10^3\ \mathrm{m^2\ GeV^{-3}\ sec^{-1}} \cdot \frac{E^4}{\rho^2} \\ &= 0.879\ \mathrm{GeV^{-1}\ T^{-2}\ sec^{-1}} \cdot E^2 \cdot B^2 \end{aligned} \tag{2.19}$$

Here ρ is the bend radius for a dipole with a magnetic field strength B, and E is the beam energy. The synchrotron photons, emitted in the MeV part of the spectrum, produce electrons and positrons through Compton scattering and pair production as they pass first through a 1-cm thick aluminum flange (separating the vacuum of the beam pipe from the outside) and then 1 mm of lead. These electrons and positrons then emit Čerenkov radiation as they pass through 3 cm of quartz. These Čerenkov photons, emitted in the

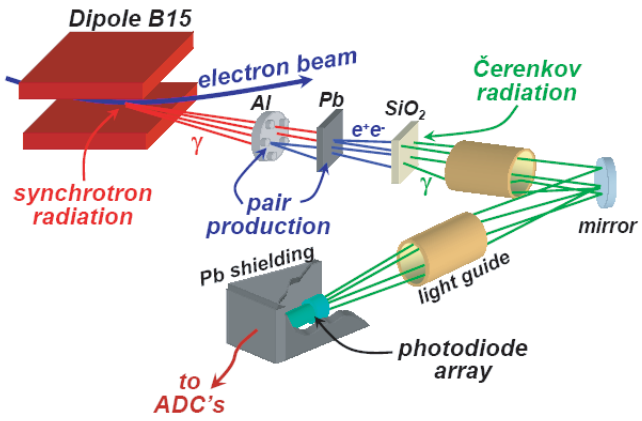

Figure 2.20: Synchrotron light monitor design.

visible part of the spectrum ($1-3$ eV), are directed into light guides (cylinders with highly reflective interiors), which eventually direct the light towards a photodiode array housed in lead shielding to protect it from stray photons and background radiation.

There are four photodiodes in the array. Each photodiode is 25 mm in diameter and is connected to a bias module providing unit gain. The photodiode signals are fed into 16-bit VME ADC boards, to be described in Section 2.8. Three of the photodiodes sit in the active region of the detector, while the fourth is blinded with black tape and sits outside of the active region, yielding a background measurement.

The synchrotron light monitor not only provides sensitivity to the beam's energy, via Equation (2.19), but also to the beam's vertical polarization component. This is because there is an asymmetry in the power spectrum of the emitted synchrotron radiation that is proportional to the electron beam's transverse vertical polarization, if any. Any asymmetry measured by the synchrotron light monitor therefore can be used to yield the beam's vertical polarization component via:

$$A_{\mathrm{SLM}} = P_y \cdot A_\gamma \tag{2.20}$$

where P_y is the amount of vertical polarization in the beam, and A_γ is the synchrotron radiation asymmetry for $P_y = 1$. Actually, A_γ is defined to be the asymmetry one would actually measure in the synchrotron light monitor, given flux-weighting appropriate for the photodiodes, as opposed to a power-weighting appropriate for a calorimeter. Given also an effective cutoff energy of roughly 3 MeV for the synchrotron light monitor (taken from GEANT simulations [84]), integration of the flux-weighted asymmetry spectrum indicates

that $A_\gamma = 60 \pm 30$ ppm [85].

Accurately knowing the amount of vertical polarization is important because, when the beam again goes through a bend, this time in the dipole chicane described in Section 2.6.1, it will emit synchrotron radiation carrying an asymmetry in its power spectrum. If this synchrotron radiation is allowed to reach the Møller detector undeterred, it will produce an asymmetry $A_{\rm SR} \approx P_y \times (100 \pm 50$ ppm), assuming a cutoff in the detector's photon efficiency at $2 - 3$ MeV. This background must be corrected for, as will be discussed in Section 4.8.2. The synchrotron light monitor therefore provides a crucial cross-check for the vertical beam polarization.

2.4.4 Wire Array

The wire array is a mesh screen of 32 horizontal and 32 vertical wires, each made of a copper–beryllium alloy. The operating principle is that of secondary emission. As electrons in the beam pass through a given wire, they ionize some of the atoms. The momentary effective charge build-up caused by the loss of some of the freed electrons is registered as a voltage across a resistor in series with the wire. Thus a voltage signal is produced for each wire that gets integrated in a CAMAC 11-bit ADC. Two aluminum foils positioned close in front and behind of the wires are each held at 100 V, which enhances the secondary emission (and thus the signal) from each wire by about a factor of two. These foils have holes in their centers so that they do not interact with the beam.

The spacing between the wires is 14 ± 1 mils, which roughly equals 356 ± 25 μm. The wires, each only $\sim$180 μm in diameter, are not strong enough to bear the full power of a $\sim$450 kW beam. Under prolonged exposure, they heat up, which can deform them considerably, or even cause them to snap. This happened during the first physics run, when the wire array was left in for a few weeks. During future physics runs, the wire array was only placed in the path of the beam for a few hours at a time, and then removed.

The wire array can be inserted or removed from the path of the beam by remote control. This is useful for getting a sense of the background signal it produces in the Møller detector via Mott and Rutherford scattering. Because the wire spacing is a well-known quantity, the wire array can be used to provide absolute calibration for the target and angle BPM's (BPM's 41 and 42 and BPM's 31 and 32, respectively). The wire array can also provide information regarding the beam's mean position on a pulse-to-pulse basis, though its inher-

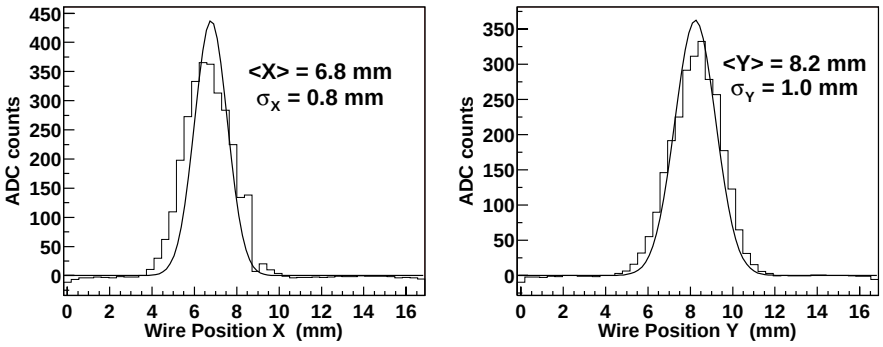

Figure 2.21: Typical wire array output, for a single pulse.

ent position resolution is poor (compared to that of the BPM's) at roughly 10 μm. More importantly, however, the wire array provides information as to the beam's spatial profile on a pulse-to-pulse basis.

A typical wire array output (after gain matching, discussed immediately below) from a single pulse is shown in Figure 2.21. From such plots, one can extract not only $\overline{x}$ and $\overline{y}$, the mean X- and Y-position of the beam, but also σ_x and σ_y, the horizontal and vertical RMS beam widths. Beam spot size jitter (where spot size is defined to be $S = \pi\sigma_x\sigma_y$) can result in density fluctuations in the liquid hydrogen target (described in more detail in Section 2.5.1), while helicity-correlated spot size differences can cause false asymmetries in the detectors. Correcting for this effect will be discussed in Section 4.5.

In order to interpret correctly the data provided by the wire array and retrieve correct values for $\overline{x}$, $\overline{y}$, σ_x, and σ_y, all of the wires must be "gain-matched." This is because the response of an individual wire to a given incident beam intensity is unique. Gain matching involves translating the wire array horizontally and vertically in discrete steps while keeping the beam still. The gains of all the wires are adjusted until the resulting beam profile shape is the same for every position of the wire array (this assumes that the beam shape remains fairly constant, which is a fair assumption). During this somewhat time-consuming calibration procedure, the wire array must be moved far enough so that all of the wires are sampled. Gain matching was performed approximately once for both Runs I and II.

2.5 Scattering Chamber

The scattering chamber, shown in Figure 2.22, is the cryostat housing the main liquid hydrogen (LH_2) target loop. It also houses a table on which there are several solid carbon

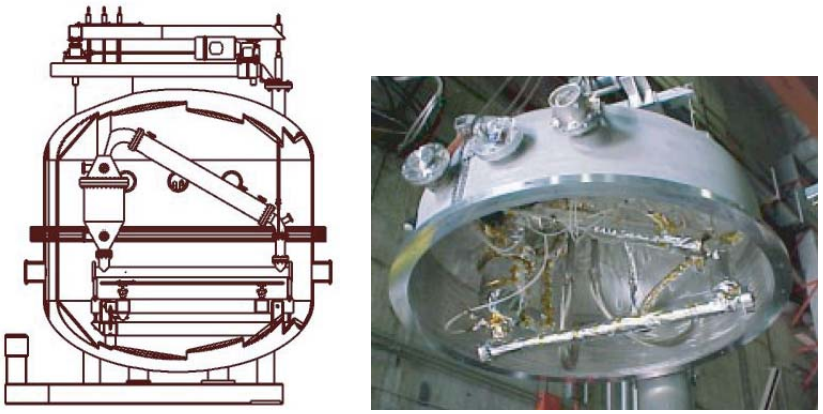

(a) (b)

Figure 2.22: Scattering chamber (a) schematic and (b) picture. In both figures, the beam would pass from the left to the right. In the picture in Figure (b), the top half of the scattering chamber, with the LH_2 target loop, is being lowered into place.

targets of varying thicknesses [86, 87].[11] The scattering chamber, roughly 2 m in diameter and 2 m in height, is connected to the beamline by flanges and is pumped out to a very high vacuum ($10^{-8} - 10^{-9}$ Torr). The scattering chamber itself remains at room temperature, with only the target loop being at cryogenic temperature.

Inside the LH_2 target loop, 55 liters of liquid hydrogen circulates through the main target tube, which is 154 cm long, 7.62 cm in inner diameter, and aligned with the beam axis to within 2 mm. The hydrogen flows through the system at roughly 10 m/s, driven by a 2 horsepower, brushless pump that operates at a repetition rate of up to 60 Hz. The loop's heat exchanger circulates cold helium gas and is connected to a Sulzer/CTI 4000 refrigerator capable of removing in excess of 1 kW of heat at 18 K, the approximate operating temperature of the liquid hydrogen. The beam can deliver up to 500 W of heat, depending on its current and energy. For this reason, the heat exchanger has a 1 kW heater included in its assembly. This heater can be controlled automatically by a target control program, running on a PC in the End Station A Counting House. The PC feedback keeps the heat load in the cooling system constant despite fluctuations in the delivered beam power due to beam drop-outs. Under normal operating conditions, the temperature of the LH_2 in the loop is constant to within 150 mK over the course of an hour [87].

One of the main requirements of the scattering chamber system was to allow for target

[11] Liquid hydrogen is the optimal target one can use for an electron–electron scattering experiment. Out of all materials, it offers the fewest number of backgrounds (many of which grow as Z^2) and the maximal radiation length per total number of Møller scatters.

changes remotely, quickly, and without breach of vacuum. Hence the target loop hangs from a frame whose vertical position is controlled by three motorized jacks. These jacks, hooked up to a remote controller, can raise the entire target loop by 15 cm, bringing it well out of the path of the electron beam. Raising or lowering the LH_2 target takes roughly 10 sec and does not affect either the loop temperature or the vacuum pressure.

The table holding the various solid carbon targets can be translated horizontally in and out of the path of the beam, once the LH_2 loop has been fully raised (microswitches prevent the table's movement should the target loop not be raised). The table's horizontal position determines which solid target, if any, interacts with the beam. The targets are held in place by aluminum frames and vary in thickness from 30 μm to 8 cm. These targets are used for initial spectrometer check-out as well as to make various background measurements, which are described in Section 4.8.2.

Working around 55 liters of liquid hydrogen (which contains the explosive energy of roughly 10 kg of TNT) demands special attention to safety [88]. In case of a sudden failure of the cryogenic system, leading towards a near instantaneous expansion of the liquid hydrogen, the system was designed to vent the hydrogen gas out the roof of End Station A through a large pipe. Hazardous Atmosphere Detectors (HAD's), capable of sensing small amounts of hydrogen gas, are located near the scattering chamber. Should the HAD's detect any hydrogen gas, they sound an evacuation alarm, notify the fire department, turn on special ventillation fans to full speed, close the hydrogen supply valves, and turn off all electrical power near the experiment.

To date, this liquid hydrogen target is one of the largest ever built for any physics experiment. It takes roughly 6 to 8 hours to completely fill the target loop with liquid hydrogen. Under normal conditions, the target remains stable (i.e., goes without venting) for a period of several weeks, with most ventings being attributable to occurrences such as power outages (resulting in a sudden loss of refrigeration in the helium plant) and computer freezes (which disable the computer program controlling the heater).

In addition to the main LH_2 target and the solid targets contained in the scattering chamber, there is an iron foil target available for use. This apparatus is only used while measuring the beam's polarization, and is therefore discussed along with the polarimeter detector in Section 2.7.5.

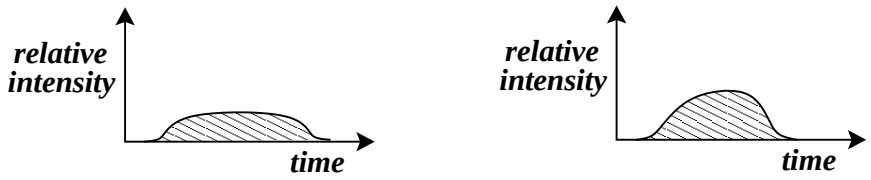

Figure 2.23: Pulse length effects. The charge time profiles for two different pulses are shown, with the total charge in each pulse assumed to be the same. The pulse on the right would deposit more heat in the target in a shorter period of time than the pulse on the left. This could produce target density fluctuations which mimic physics asymmetries.

2.5.1 Target Density Fluctuations

Though the temperature stability of the LH_2 loop is good on long timescales, temperature fluctuations on very short (a few ms) timescales can produce fake (i.e., non-physics-related) scattering rate asymmetries [86, 89]. The 154-cm length of LH_2 through which the electron beam travels under normal running conditions corresponds roughly to 0.18 radiation lengths. If, however, the target gets heated by the beam more for one pulse than it does for another, the liquid hydrogen will expand and its density will be different for the two pulses. Hence, the amount of hydrogen each pulse sees will be different. If the two pulses are the two in a helicity pulse pair, this will manifest itself as a scattering rate asymmetry. Such effects are known as target density fluctuations, and they can both degrade the detector resolution (by adding random noise which is not taken out by any correction procedure, since the pulse-to-pulse target density is not known) and potentially cause a helicity-correlated systematic, if the effect is consistent and stable with respect to helicity.

It is important to note that normalizing the detector signal to pulse charge, as must be done, does *not* correct for the problem, since it is the target density itself that is fluctuating. The detected scattered flux ϕ can be written as $\phi = N_{beam} \cdot N_{target} \cdot \sigma$, where N_{beam} and N_{target} are the number of particles in the beam and target, respectively, and σ is the cross section governing the process. Normalizing ϕ to N_{beam}, as is done by normalizing to beam current, is of course necessary, but it still does not account for fluctuations in N_{target}.

Target density fluctuations can potentially be caused in two ways. The first and probably more significant effect comes from variations in the beam's spot size. A smaller beam will deliver more heat to a smaller volume, consequently heating it more and causing density fluctuations. The second effect might arise from pulse length variations. There is a source feedback that keeps the beam current constant, but the length of the pulse may vary.

Such a scenario is illustrated in Figure 2.23, which shows two pulses having the same integrated charge, but different intra-pulse charge distributions. The pulse with the higher peak intensity will deposit more heat in the target in a shorter length of time, leading towards density fluctuations.

A limit on the size of the false asymmetries stemming from target density fluctuations can be estimated by looking for correlations between detector asymmetries and spot size differences, as will be discussed in Section 4.6. However, to ensure that this limit is as small as possible, it is important that there be sufficient transverse mixing of the hydrogen in between beam pulses. In this way the volume in which the beam's energy is deposited is maximized, minimizing the effects of possible density fluctuations. In order to encourage transverse mixing, eight wire mesh screens were distributed along the length of the target tube, positioned randomly every $4 - 8$ inches (this avoids resonances which might move fluid into the beam path in phase with the 120 Hz beam) [86, 87]. These wire mesh screens have 1.5-inch diameter holes cut into them, allowing the beam to pass through. These screens were designed to introduce turbulence on the $1 - 2$ mm scale (roughly equal to the beam spot size), as well as to introduce a transverse velocity component to the circulating hydrogen. This ensures that the liquid in the beam interaction region of the target loop gets mixed thoroughly throughout the whole volume, reducing the effects of target density fluctuations.

2.6 Spectrometer and Collimator System

After the liquid hydrogen target, all particles pass through the spectrometer and collimator system. The main goal of the spectrometer system is to select out from the wide spectrum of particles generated in the target (which includes not only Møller scatters of physical interest, but also Mott's, inelastic ep's, photons, positrons, neutrons, and pions) as clean a signal as possible for integration in the detectors. This is vitally important since the detectors have an integrated calorimetric response and cannot themselves differentiate between individual particles. Hence the spectrometer and collimator system is designed to provide a spatial separation of particles at the position of the detectors, while allowing the primary beam to pass unimpeded to the beam dump. Figure 2.24 contains an overview of the entire system.

The spectrometer can be broken down into three basic parts: the dipole chicane, the

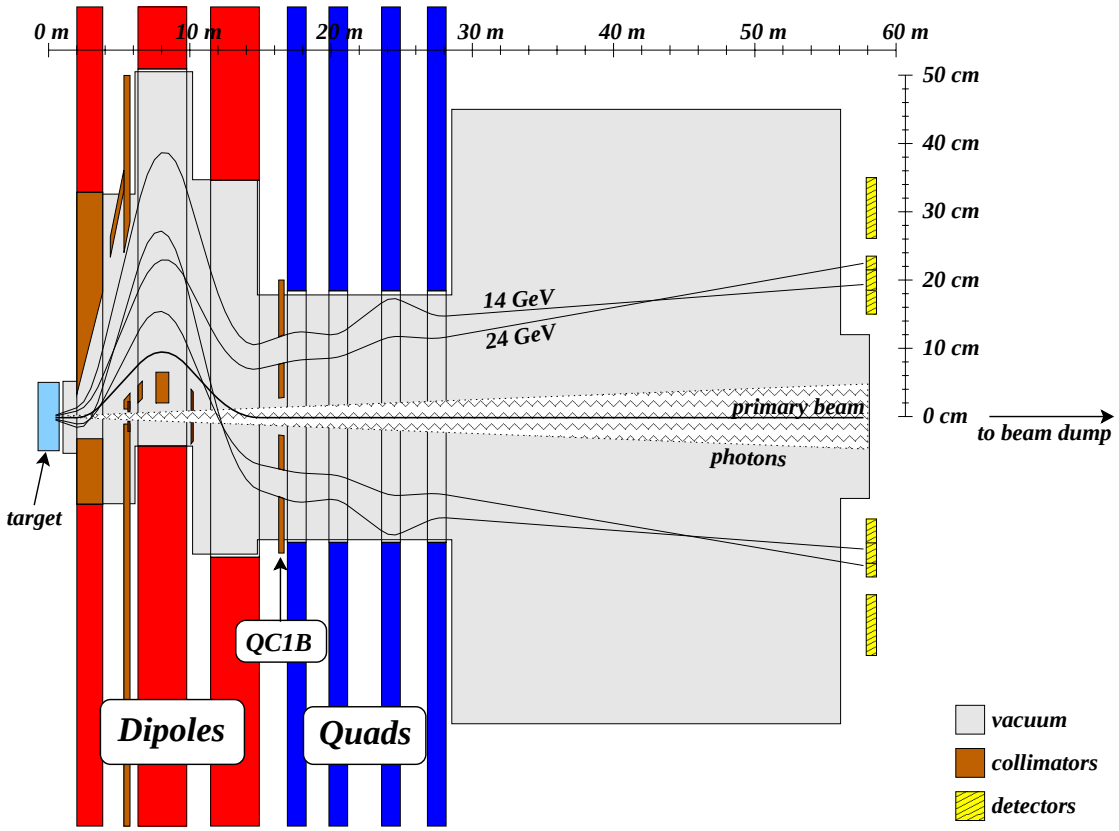

Figure 2.24: Spectrometer. Some collimators are omitted for clarity. Note the very different x and z scales (where z is understood to mean the beam direction, and x represents horizontal displacements). Sample ray traces are shown for the primary beam, as well as for 14 and 24 GeV Møller scatters.

main acceptance collimator (QC1B), and the quadrupole quadruplet. The beampipe snakes its way through the whole assembly, so that the entire volume through which both the primary beam and the scattered particles must pass, up until just a few meters before the detectors, is maintained at high vacuum. Almost all beampipes are made from aluminum, and all collimators have been fabricated either from copper or from "non-magnetic" tungsten, an alloy comprised almost entirely of tungsten, with a few percent of nickel and copper as balance elements. Iron is avoided in all regions of the spectrometer where the beam might potentially interact, as the asymmetry in the process of polarized electron scattering off iron (especially iron that has been partially polarized by stray magnetic fields) is many orders of magnitude larger than the Møller scattering asymmetry. The dipole and quadrupole magnets necessarily contain iron, but collimators prevent particles from scattering from these magnets and reaching the detectors.

2.6.1 Dipole Chicane

The dipole chicane serves two main purposes. First, it brings the detectors out of the line-of-sight of the target by creating what's called a "two-bounce" photon system [61]. The basic idea is that a photon must bounce twice off the vacuum pipe apertures or collimators in order to reach the detectors. Second, the chicane eliminates the large background of "soft" (low momentum) particles that would otherwise swamp the detectors. These are blocked by various strategically placed masks and collimators. For instance, the field strength of the first dipole is tuned so that low-energy electrons (with momenta less than 5 GeV) get "over-bent," colliding into its left (northern) inner wall. Similarly, because of their positive charge, all positrons with momenta less than 28 GeV smash into its right (southern) inner wall. The power deposited by these unwanted background fluxes is quite large (on the order of kilowatts) and would damage the magnet if left unprotected. For this reason, water-cooled copper plates have been installed on the insides of the magnet. These kinds of measures have been taken throughout the dipole chicane.

The dipole chicane is comprised of three dipole magnets with $\int B_\perp \cdot dl = 0$, so that it functions as an achromat.[12] Synchrotron radiation, however, prevents the chicane from functioning like a true achromat for all particles passing through it. That is, as particles travel through the chicane, they lose energy due to synchrotron emission. Higher energy particles (e.g., primary beam electrons) lose more energy than do lower energy particles (e.g., signal electrons), as stated in Equation (2.19). The consequence of this is that the path integral $\int B_\perp dl$ becomes slightly energy-dependent. The magnetic strength of the most downstream dipole is adjusted so that the integral remains zero, but the adjustment only works for signal electrons. Electrons of higher or lower energies will see a slightly non–zero $\int B_\perp dl$. The difference is very small, though. In fact, the third dipole's strength need only be adjusted by 0.07%, which moves the mean position of the electron fluxes at the detectors by a few millimeters.

In order to preserve the highest degree possible of azimuthal symmetry in the scattered particle flux profile, it is important that all three dipoles produce highly uniform magnetic fields. In addition, so-called pole face rotations must be taken into account. That is, since particles in the chicane do not enter and exit the dipoles normal to their magnetic pole

[12] The integral is performed along the beam particle path. The magnetic field strength $B_\perp$ is defined to be everywhere perpendicular to the particle path.

faces, field gradients at the entrances and exits of the dipoles produce a small quadrupole moment in the magnetic field. This quadrupole moment focuses horizontally while defocusing vertically, providing another way in which the azimuthal symmetry of the flux profile at the detectors can be ruined. If this aberrant quadrupole moment was not corrected for, the horizontal and vertical dimensions of the Møller signal flux at the detector would be different by 5 mm. A small adjustment in the strength of the fourth spectrometer quadrupole compensates for this effect.

As mentioned previously, photon collimation is one of the primary reasons for having the dipole chicane. When thinking about photon collimation, there are three types of photons to consider: those emitted by synchrotron radiative processes, those emitted by bremsstrahlung processes, and soft "splash" photons. As the electrons traverse the chicane, they emit synchrotron radiation. This synchrotron radiation is blocked partly by the main acceptance collimator, and further by the additional collimators described in Section 2.6.4. Bremsstrahlung processes in the liquid hydrogen target produce a very large flux of high energy photons that, if not collimated, would otherwise swamp the detectors and seriously degrade their statistical resolution. Wide-angle, high-energy photons are largely blocked by two collimators located between the first and second dipoles and between the second and third dipoles, respectively (see Figure 2.24). These so-called photon collimators absorb nearly 2 kW of radiation from this hard gamma flux.

Lastly, as particles hit the inner edges of collimators and beampipes, particularly just downstream of the quadrupoles, photons that have angles capable of bringing them into contact with the detectors may be produced by electromagnetic showers. Such edge effects are often referred to as "shower splash." To minimize this "splash," tungsten is used for collimators that have a direct line-of-sight to the detectors, since its Molière radius is nearly 40% less than that of copper. Nonetheless, some soft splash photons are still present. A special collimator designed to suppress these photons further, installed between Runs I and II, will be described in Section 2.6.4.

All three dipoles generate magnetic fields on the order of $\pm 1 - 2$ Tesla. They are connected to megawatt–capable power supplies housed in a building adjacent to End Station A. The high magnetic field strengths require currents on the order of 1 kA, so that the magnets need to be water cooled in order to dissipate the considerable heat generated. The overall stability of the chicane, important both for experimental efficiency as well as for personnel

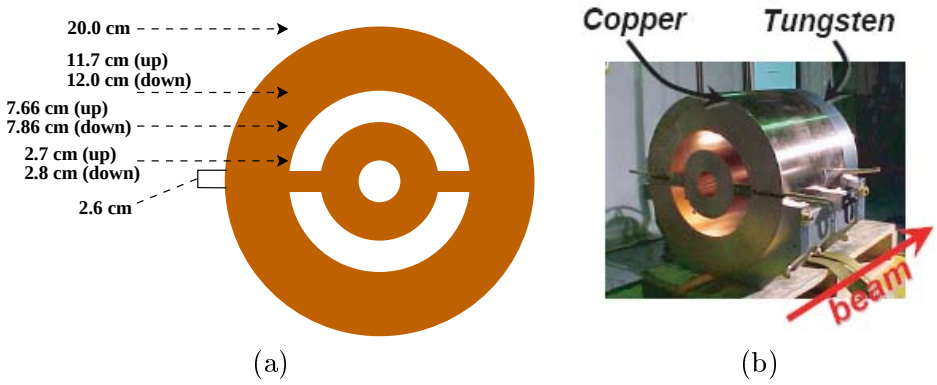

Figure 2.25: Acceptance collimator (a) design schematic and (b) picture. The dimensions quoted are for the upstream and downstream faces.

and mechanical safety, is monitored. The dipoles' power supply currents (which are logged automatically in the data stream) are connected to an alarm program which alerts the experimenter to any considerable deviation (typically greater than 0.1%).

2.6.2 Acceptance Collimator

At the end of the chicane, the acceptance collimator (QC1B) defines the energy spectrum of the particles that are detected. It does this by exploiting the correlation between energy and angle of the particles downstream of the target. This collimator is depicted in Figure 2.25. It consists of two annular rings positioned concentrically and connected by spokes, which serve both to support the inner annulus as well as to block synchrotron radiation (which carries a large asymmetry) produced in the dipole chicane. Along with the primary beam, forward-angle Mott and Møller scatters, as well as the high-energy photon beam (collimated by the dipole collimators), pass through the innermost aperture. Signal Møller electrons, in addition to wide-angle Mott and ep electrons, pass through the radial aperture.

By exploiting the correlation between the momentum of a particle and its scattering angle (and hence radial position at some point along the z-axis), and by using the fact that this correlation is different for ep scatters than for Møller scatters, the exact placement and size of the radial aperture can be used in conjunction with the quadrupoles discussed below to cleanly separate signal from background. Figure 2.26(a) shows the radial position of the scattered flux at the main calorimeter versus momentum without the acceptance collimator in place. In this configuration, one would be forced to either throw out many good Møller scatters (decreasing the signal) or accept more Møller scatters at the expense of including

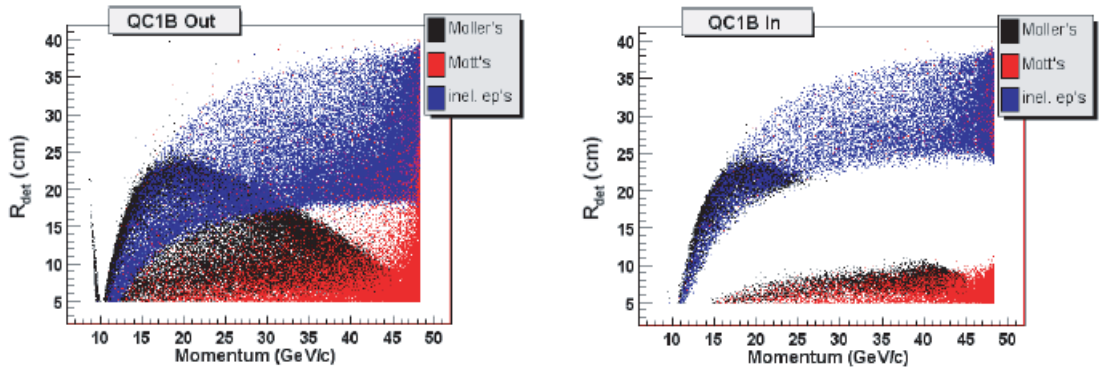

(a) (b)

Figure 2.26: Scattered flux phase space plots, for QC1B (a) out and (b) in. Here R_{det} is the radial position of the particles at the detector face.

more ep's and Motts (decreasing the signal to noise ratio, and possibly increasing the systematic error due to the large asymmetry in the inelastic ep's). The same scattered flux phase space plot is shown in Figure 2.26(b), except this time with the acceptance collimator in place. To a large extent, the Møller, Mott, and ep electrons have been cleanly separated. There is still some contamination of the Møller signal in the region $15 < r < 23.5$ cm, the active area of the main calorimeter. The acceptance collimator was designed to maximize the signal to noise ratio in the detector. The levels of background contamination is discussed further in Section 4.8.3.

The acceptance collimator consists of two sections brazed together to form a single, composite structure. The upstream portion (12 inches thick) is made of copper, while the downstream portion (3 inches thick) is made of tungsten, to avoid the photon "splash" described on page 78. To further suppress this potential source of background, by reducing the amount of edge material having a direct line-of-sight to the detectors, all edges (with the exception of the outermost one) flare radially inward or outward when viewed along the collimator's length, so that they remain more or less parallel to the passing particles' trajectories. The overall thickness of the collimator is 40 radiation lengths, divided roughly equally between the copper and tungsten, with the entire collimator being water cooled.

2.6.3 Quadrupole Quadruplet

The acceptance collimator alone cannot achieve the signal separation needed by the integrating calorimeters, which cannot use energy as a means of particle identification. Con-

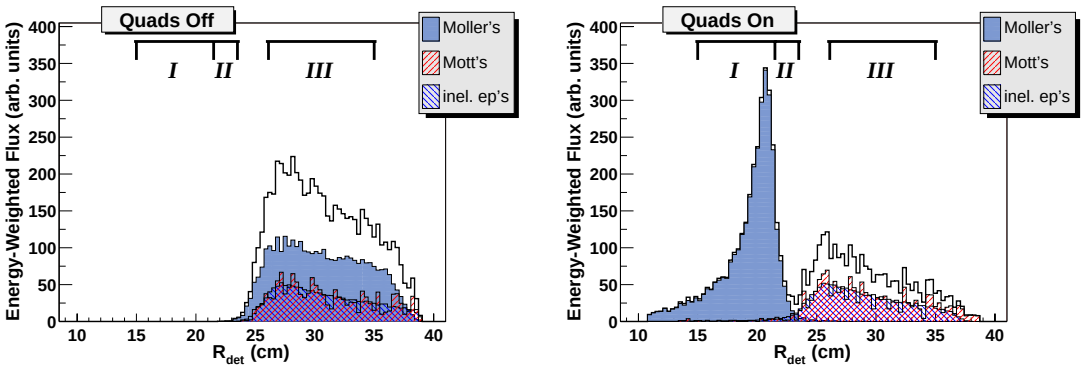

Figure 2.27: Effects of quadrupole focusing on signal flux. The radial, energy-weighted signal flux is plotted versus the radial position at the detector for the quadrupoles (a) off and (b) on.

sequently, the momentum–dependent focusing power of quadrupoles is used to separate the Møller and *ep* fluxes at the face of the detector. Four highly uniform quadrupoles are arranged in alternating $FDFD$ fashion, forming a quadrupole quadruplet. Here F and D refer to focusing and defocusing action, respectively, in the horizontal axis. The quadrupoles' field gradients vary by less than 0.1% in the region in which particles travel. When turned on, the quadruplet focuses particles in the energy range $14-24$ GeV (typical of the Møller electrons) and shifts them radially inwards by 10 cm or more, while tending to defocus particles of higher energy slightly and moving them radially inwards by only a few centimeters on average. This focusing–defocusing action can best be appreciated by looking at an illustration such as that in Figure 2.27. The data contained in the plots come from a simpler version of the Monte Carlo to be discussed in Section 4.8.3, and are meant only to provide a qualitative illustration, not to get quantitative estimates of signal rates. With the quadrupoles off, the Møller and *ep* distributions overlap considerably. When the quadrupoles are turned on, however, the Møller peak becomes tightly focused and moves radially inward, while the *ep* distribution becomes slightly defocused and does not move much. Thus the two signals are separated and can be integrated independently.

2.6.4 Additional Collimation

The previous sections described the salient features of the spectrometer in a fair amount of detail. However, there are a few items that were not mentioned, but which nonetheless merit description. These are: the calibration collimator (QC1A), the collimator masks ($CM1-8$), and additional synchrotron collimators. These collimators are pictured in Figure 2.28.

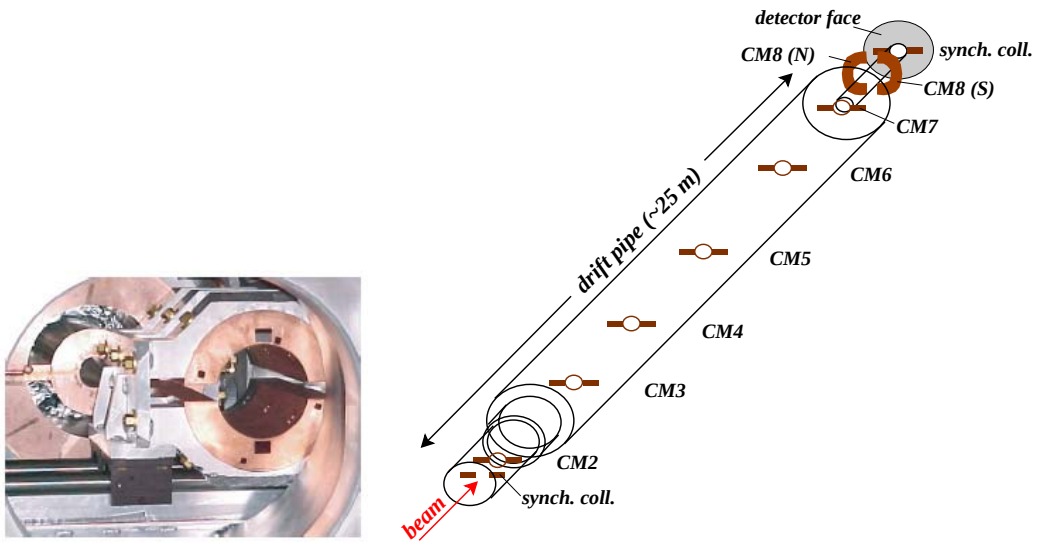

(a) (b)

Figure 2.28: Additional collimators. Figure (a) shows the calibration collimator QC1A in front of the main acceptance collimator QC1B. Figure (b) shows the collimator masks CM2–8 and the synchrotron collimators in the drift pipe downstream of the last spectrometer quadrupole. Note the two pieces of CM8 (north and south pieces).

The calibration collimator is basically the negative of the acceptance collimator. That is, it is designed to block all particles that make it through the radial aperture in the acceptance collimator. Actually, there are six holes bored through the collimator to allow some particles to pass. The collimator is remotely insertable and removable. For nominal physics running, it is completely removed from the beamline. It is only inserted for special "calibration" runs, hence its name. It is needed for polarimetry, in order to minimize the backgrounds for this measurement. The two larger holes centered on the vertical axis serve for this purpose. It is also needed for profile scans. The four smaller holes allow for a precise determination of the Møller and ep fluxes at the detectors. Like the main acceptance collimator, the calibration collimator is composed of 12 inches of copper followed by 3 inches of tungsten, for an overall thickness of 40 radiation lengths.

The collimator masks CM$1-7$ are mounted at various locations along the beamline (the first is located between the second and third quadrupoles, while the remaining are located between the last quadrupole and the detectors) and are intended to block stray photons. Without these masks, soft "splash" photons, especially from the inner edges of the main acceptance collimator and the photon collimator between the first and second dipoles, as well as energetic multi–bounce photons from the target would hit the detector. These seven

masks are tungsten rings, supported by horizontal spokes.

The last collimator mask (CM8) sits directly in front of the detector cart. The significant flux of high-energy ep scatters interacting in the calorimeter produces neutrons and other neutral hadrons that are capable of penetrating the shielding encasing the PMT's, thereby generating significant responses in the tubes' cathodes. Since the ep flux possesses a large asymmetry, this neutral background flux will also carry a large asymmetry. This 9-inch-thick tungsten and lead collimator encircles the beampipe and reduces this neutral background by an order of magnitude. The collimator is actually constructed out of two separate pieces that can be inserted and removed remotely. When inserted, the assembly suppresses the background for the Møller detector, at the cost of eliminating all signal in the ep detector. This unexpected background was discovered during Run I, after which time CM8 was designed, constructed, and installed. Hence, this collimator is only present during nominal Run II physics running.

The last additional collimation to discuss are two sets of synchrotron collimators. Of the collimators already described, the spokes of the main acceptance collimator and those of CM1 and CM7 are designed to block substantial amounts of background synchrotron radiation. Two more sets of synchrotron collimators are necessary to bring the total synchrotron background down to an acceptable level. Both sets are horizontal spokes, 20 radiation lengths (7 cm) thick and made of tungsten. The first set is located just downstream of the last quadrupole, while the last set is bolted directly to the face of the detector.

2.7 Detectors

The detector cart, pictured in Figure 2.29, is located 60 meters downstream of the target. It houses most of the detectors used in the experiment, with the exception being the luminosity monitor, which is located 8 meters further downstream. The first detector package in the cart is the profile monitor wheel. It is used to check collimator alignment and to tune the spectrometer during commissioning, as well as to make accurate measurements of the signal and background fluxes. After the profile monitor comes the main calorimeter, which integrates the Møller and ep fluxes separately. The polarimeter (not shown in the figure, but located between the profile wheel and the integrating calorimeter) is used to make periodic measurements of the beam's longitudinal polarization. Finally, there is a pion detector,

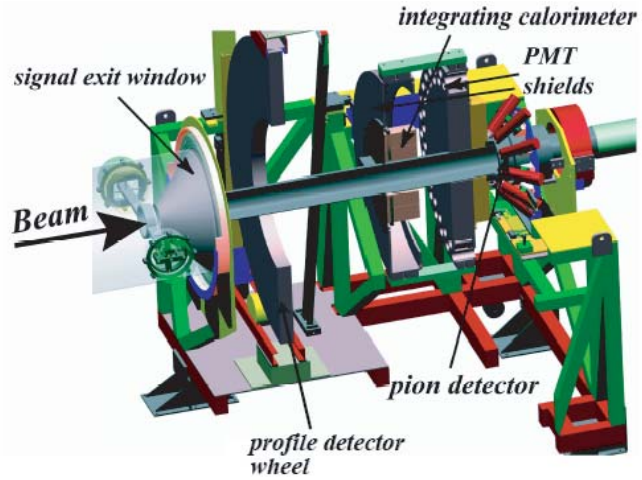

Figure 2.29: Detector cart. Note the shielding for the main integrating calorimeter PMT's. The collimator mask CM8 (not present for Run I) is not pictured, but would be between the signal exit window and the profile wheel.

which is used to measure the pion flux.

2.7.1 Møller and ep Detectors

The design concept of the Møller and *ep* detectors (which are the names given to the two parts of the main integrating calorimeter) was driven by the large flux of particles they would see [61, 90]. In order for the experiment to be feasible, very forward-angle Møller scatters ($\theta = 4 - 8$ mrad in the lab frame) need to be detected. At these angles the cross sections are rather large ($\sigma_{tot} \approx 10$ μbarn), especially given the very high beam luminosity ($\mathcal{L} \approx 4 \times 10^{38}$ cm^{-2} s^{-1}). The flux accepted by the calorimeter is 20 million electrons per pulse at 120 Hz, with an average energy of 17 GeV. This translates into a total radiation dose of about 100 Mrad for the entire experiment. One rather new detector technology particularly adept at dealing with such high dose rates is known as quartz fiber calorimetry [91].[13]

The basic idea behind quartz fiber calorimetry is illustrated in Figure 2.30. The detector is comprised of fused silica optical fibers sandwiched between 3-mm-thick copper plates. As charged particles travel through the copper, they produce electromagnetic showers. The charged particles in the showers then emit Čerenkov radiation as they pass through the silica fibers. The plates are angled at 45° with respect to the incident electron flux, so that

[13] This is a bit of a misnomer, since the "quartz" used is not crystaline, but amorphous (fused) silica.

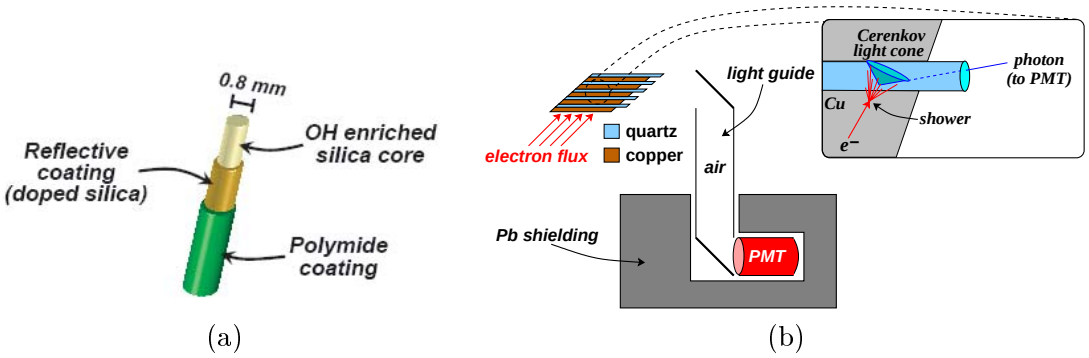

Figure 2.30: Working principle of quartz fiber calorimetry. Figure (a) shows the details of the quartz fibers used. Figure (b) shows the whole idea.

most of the Čerenkov radiation emitted will be in a direction parallel to that of the fibers. The fibers then act as optical guides, channeling the photons towards photomultiplier tubes (PMT's) for photon detection. Detectors employing this technology have been proposed for the CMS experiment at CERN, where the expected total dose exceeds 1 Grad [91]. Besides radiation hardness, another useful benefit of this technology is the quick response time associated with the Čerenkov process.

One feature of quartz fiber calorimetry is that the size of the signal generated in the PMT's by any given electron is proportional to the electron's initial energy, since the number of photons produced is directly proportional to the total charged path length of the electron's shower, which in turn is directly proportional to the energy of the electron. Simulations show that, for the chosen thickness of the calorimeter, each electron should produce roughly 20 photons.

The thickness of the main integrating calorimeter (both the Møller and the *ep* portions) was chosen to be thick enough so that energy fluctuations in the showers are as small as possible (shower energy fluctuations result in photon number fluctuations, which limit the statistical resolution of the photon detectors), while at the same time limiting its thickness to reduce its sensitivity to unwanted background particles, such as minimum-ionizing particles (abbreviated as MIP's) and pions [61]. The chosen length of 16 radiation lengths of copper was found to be a good compromise.[14] For example, the energy-weighted contribution from charged pions should be less than 0.1% of the signal from the electrons. The chosen length also ensures that the muon signal will be completely negligible. The 5% energy fluctuations

[14] Approximately 10% of the active volume of the calorimeter comes from the silica fibers. Their relative contribution to the detector's radiative thickness is therefore only 1%.

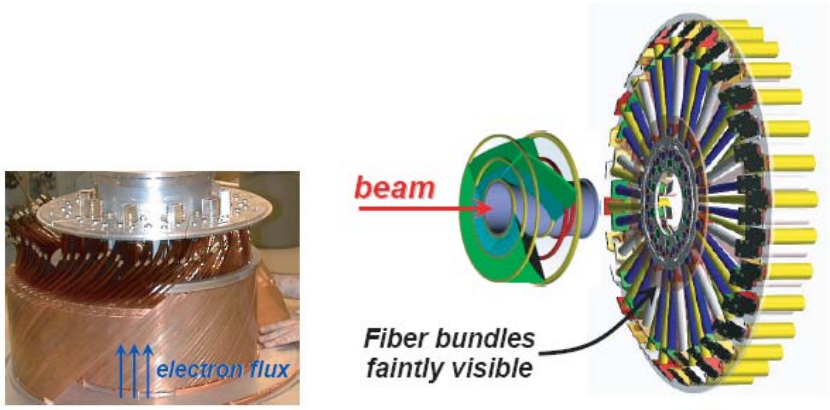

(a) (b)

Figure 2.31: Møller and *ep* detector design. The quartz–copper sandwiches, with the protruding fiber bundles, are clearly visible in (a). Those fiber bundles "feed" the collected photons into light guides, as shown in (b), which then bring the light to PMT's housed in protective shielding.

in the showers caused by this rather short length should result in only a minor resolution degradation.

The geometry of the calorimeter is shown in Figure 2.31. The whole calorimeter surrounds the beamline and is subdivided both radially and azimuthally. The Møller detector is actually comprised of three separate rings, referred to as the "in," "mid," and "out" rings. These rings are divided azimuthally into 10, 20, and 20 separate sections, respectively. The *ep* detector is one single ring, subdivided into 10 azimuthal sections. This radial and azimuthal subsectioning is summarized in Figure 2.32. Each section gets its own PMT, which receives light from the fibers sandwiched between all the plates contained in that particular section. The bundled fibers actually direct the photons to air light guides, consisting of two end mirrors and a mirrored cylinder, which direct the light into the PMT's (since the light is well collimated leaving the fibers, very few bounces are needed). The light guides allow for the PMT's to be distanced radially from the beamline, where the particle flux is significantly reduced. In addition, the PMT's are housed in a 7-inch-thick wall of lead shielding. These measures reduce the total integrated dose that each PMT may receive to roughly 10 ppm of their maximum dose level [90].

As mentioned previously, the signal flux at the calorimeter is roughly 20 million electrons per pulse, with each electron producing on average 20 photons for the PMT's. The radial profile of the flux and the segmentation of the Møller and *ep* detectors are such that any one PMT may face as many as 20 million photons on a pulse-by-pulse basis. At these high

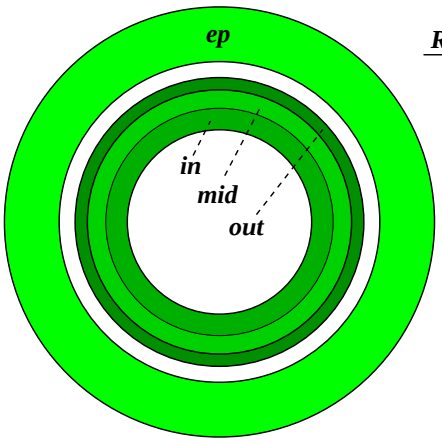

Region	# sections	R_{min}	R_{max}	
in	10	15 cm	18.5 cm	I (15 − 21.5 cm)
mid	20	18.5 cm	21.5 cm	
out	20	21.5 cm	23.5 cm	II
ep	10	26.1 cm	35 cm	III

Figure 2.32: Dimensions of Møller and *ep* detector rings. The rings each have different numbers of azimuthal sections (i.e., fiber bundles), each with a dedicated PMT. The "in" and "mid" rings together form Region I, while the "out" and "ep" rings constitute Regions II and III, respectively.

rates, both cathode and anode saturation could present problems, as both are capable of producing signficant nonlinearities in the detector PMT signals unless special measures are taken. To avoid cathode saturation in the R2154 Hamamatsu PMT's, a copper wire mesh screen rests on each PMT's cathode, assisting in the quick redistribution of charge over the cathode's surface in between beam pulses. To avoid anode saturation, the gain of each PMT's base is adjusted so that its output signal is no more than 0.5 V, given the highest expected beam intensity.

2.7.2 Luminosity Monitor

The signal for the luminosity monitor (located 8 m downstream of the detector cart) is comprised of very forward–angle scattering Mott and Møller electrons. An EGS4 simulation finds that the contributions from those two sources are 70% and 30%, respectively [92]. The average energy of these particles is roughly 40 GeV, with the average scattering angle being approximately 1.5 mrad. At these kinematics, the physics asymmetry present in the signal is expected to be roughly -10 ppb, just below the statistical precision that can possibly be achieved with the detector, given the flux (high though it is) and integrated running time for the experiment. The detector is therefore capable of providing a useful null cross-check for the main Møller asymmetry measurement, as will be discussed below.

The design of the luminosity monitor was driven primarily by the extraordinarily high flux the detector would be subjected to [93]. Under nominal running conditions, GEANT simulations predict an incident flux at the detector's location of approximately 350 million

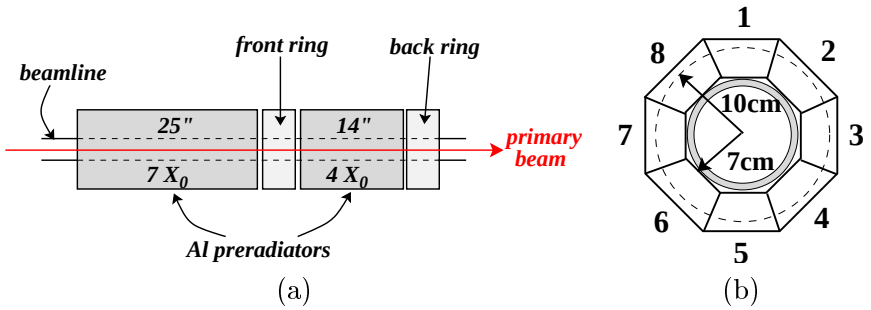

Figure 2.33: Luminosity monitor schematic, (a) side and (b) front view. Both rings are separated into eight chambers, each of which generates its own signal.

electrons per pulse, which, given their average energy, corresponds to nearly 270 W of power at 120 Hz. These electrons deposit roughly half of their total energy, for a total of 130 W of power deposited at 120 Hz. Over the course of the experiment, this works out to a total integrated dose on the order of a few Gigarads. Because of the challenges connected with such high rates and working in such a high radiation environment, an ionization chamber design was settled upon. Ionization detectors can be made very robust against radiation damage and have been shown to remain very linear up to high signal fluxes.

As shown in Figure 2.33, the luminosity monitor is divided into two rings, referred to as the front and back rings, respectively. Each ring surrounds the beam line, so that the considerable signal of primary beam electrons and high energy photons (roughly 500 kW of power at 120 Hz) passes harmlessly through its center. The inner and outer radii of the active area for each ring are 7 cm and 10 cm, respectively, large enough to accommodate the bulk of the signal.

Each ring is segmented into eight separate chambers (each "ring" is therefore actually octogonal in shape). Inside each chamber are eleven parallel plates, positioned transverse to the beamline and held alternately at ground or 100 V (and called cathodes and anodes, respectively). When a scattered Mott or Møller electron traverses the chamber, it ionizes gas molecules, freeing electrons and creating ions. The liberated electrons, and, more slowly, the ions, travel along the electric field lines towards the plates, where they are collected (the electrons to the anodes, the ions to the cathodes). Thanks to the plates' capacitance, this momentary charge build-up generates a voltage drop which is then detected.

A delicate balance exists between the choice of plate separation distance, plate voltage, and the size of the signal produced by the detector. In general, higher voltage and small plate separation is good because it minimizes the time that the ions and electrons spend

traversing the distance in between the plates, thereby minimizing the chances of ion recombination, which can produce non-linearities. However, very high voltages and very small plate separations can lead towards damaging (and potentially dangerous) arcing effects. In the end, a voltage of 100 V (as previously stated) and a plate separation of 1 mm was chosen as the best compromise. In addition, each chamber is filled with N_2 as a buffer gas (the pressure being held just above atmosphere). By displacing comparatively reactive oxygen from the insides of the chambers, the inert buffer gas further reduces the chances of electron and ion recombination enroute to the plates.

The electronics associated with the luminosity monitor are particularly simple. The chosen plate separation and plate voltage values produce a typical output signal from a single chamber between 4 and 8 V, allowing the raw detector signal to be fed directly into an ADC, without the need for any amplifiers. Two signals are drawn from each chamber. One is the high voltage signal supplying the anodes, with the dc-bias filtered out. The other is the high voltage signal supplying the cathodes. These two signals are transmitted via a single twin-ax cable to a 16-bit integrating ADC. Having the ADC channel read the signals differentially greatly reduces the sensitivity to transmission noise due to pick-up in the cables.

To counteract the effects of the high radiation environment in which the luminosity monitor must operate, a number of precautions have been taken. First, the entire detector is constructed almost entirely of aluminum. Second, where they are needed, the BNC connectors are made with a filling of highly radiation–resistant material known as rexolite, instead of the standard Teflon. Lastly, immediately upstream of the front ring is 7 radiation lengths (25 inches) of aluminum preradiator. This preradiator acts both as shielding for synchrotron radiation (largely emanating from the last chicane dipole in the spectrometer) as well as a showering medium for the signal electrons. The amount of degrader used was chosen to place the front ring close to shower maximum, while at the same time eliminating as much of the 17 W of incident synchrotron radiation as possible. Between the front and back rings is an additional 4 radiation lengths (14 inches) of aluminum, placing it well beyond the shower maximum. Simulations predict (and observations confirm) the signal size in the front ring to be three times higher than in the back ring. This enables the back ring to act both as a relative linearity check for the front ring, since its smaller signal size ensures that it will remain linear to higher currents, as well as a cross-check for asymmetry

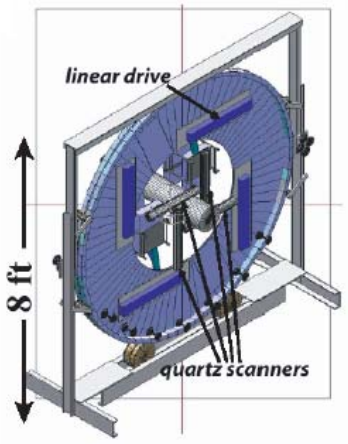

Figure 2.34: Schematic of profile wheel. The profile wheel supports four quartz Čerenkov scanners, used to map out the signal flux. The whole wheel can rotate, and each quartz scanner can be moved radially along linear drive rails.

measurements, since ultimately both rings detect the same source.

Overall, the luminosity monitor serves two main purposes. The first purpose, to look for the effects of target density fluctuations, has already been described in Section 2.5. The second purpose is to monitor false asymmetries. Because its expected physics asymmetry is so small, and because it is more sensitive to changes in beam parameters, the luminosity monitor is a valuable online monitoring device for many systematic effects.

2.7.3 Profile Monitor

Sitting directly in front of the main integrating calorimeter, the profile monitor, shown in Figure 2.34, consists of a large wheel which supports four movable quartz Čerenkov scanners. Each scanner consists of a 5-mm thick quartz block tilted 45° from the beam axis. Charged particles produce Čerenkov radiation in the quartz. The light travels down the length of the mirrored light guide into a PMT placed at the end. To prevent Čerenkov production in the light guide from contributing to the PMT signal, the light guides are evacuated to 0.6 Torr.

Because of abundant soft backgrounds, two of the quartz Čerenkov scanners are modified, as shown in Figure 2.35. In these two scanners, insertable 15-mm thick tungsten preradiators are placed directly in front of the quartz blocks. This $6X_0$ preradiator shields out low momentum background particles and amplifies high momentum signal particles, which shower in the tungsten. In addition, these scanners have insertable shutters in front

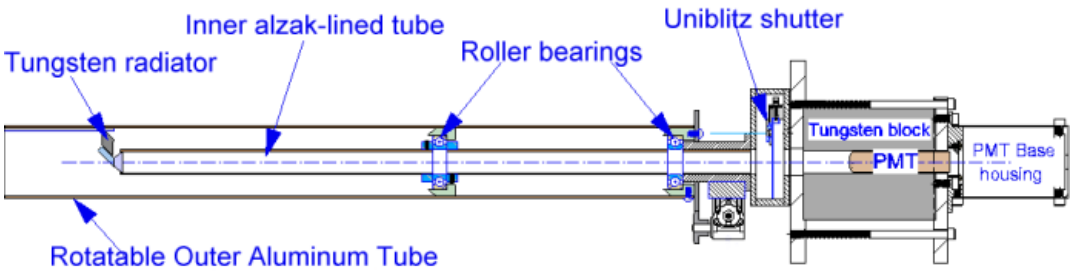

Figure 2.35: Detail of one modified Čerenkov scanner on the profile wheel. A tungsten radiator in front of the quartz block amplifies the signal. A remotely insertable/removable shutter can allow for background measurements.

of the PMT's, which allow for measurements of the background signals produced by energetic particles that penetrate the lead shielding of the PMT's.

The scanners are all attached to arms that can be moved radially inwards and outwards, between 15 cm and 55 cm from the beamline. In addition, the entire wheel itself can rotate by 180°. All movements are controlled remotely by a LabView program, and automated scans are performed to map out the particle flux both radially and azimuthally. The quartz blocks' radial and azimuthal dimensions are only 5 mm and 20 mm, respectively, giving the scanners adequate position resolution. In this way, fairly detailed profile maps of the Møller and *ep* fluxes can be obtained. All scanner data are read out by standard CAMAC 11-bit ADC's, high precision not being required (only pulse averages are taken, not pulse differences).

For nominal physics running, all quartz scanners are removed to their maximum radial positions, so as not to contaminate the signal flux. Special profile scans are conducted by rotating the wheel and moving the scanners radially in and out. The profile maps thus produced are useful in a number of ways. First, they provide a means of inspecting the alignment of collimators and of the detector, as well as checking that the spectrometer magnets are correctly set. This is important during the initial commissioning phase of each data taking period (i.e., at the beginnings of Runs I and II). Second, insertion of the calibration collimator described in Section 2.6.4 provides a very clean radial separation of the Møller and *ep* fluxes for certain azimuthal angles. Essentially, these profile maps help calibrate a detailed Monte Carlo simulation that calculates the total *ep* (both elastic and inelastic) contribution to the Møller asymmetry. In addition, this simulation verifies the experimental kinematics (i.e., $\langle Q^2 \rangle$), which is necessary when relating the final asymmetry

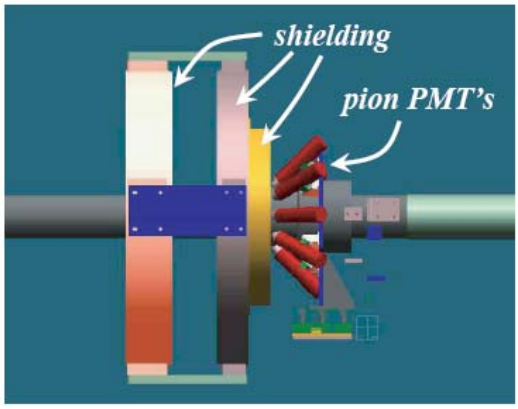

Figure 2.36: Pion detector schematic (side view). Note the lead shields, the middle one housing the PMT's for the main integrating calorimeter as described in Section 2.7.1.

to a value for the weak mixing angle.

2.7.4 Pion Detector

As described in Section 1.3.1, it is necessary to measure the pion flux seen by the Møller detector in order to compensate for the dilution this flux will produce as well as any significant asymmetry it may contribute to the Møller asymmetry. A pion detector is therefore placed directly behind the Møller detector, shielded by an additional 25 cm of lead. The amount of shielding was chosen in order to suppress the Møller electron flux completely (these electrons are considered a background for the pion detector, and their flux is expected to be several hundred times greater than the pion flux), while allowing for some enhancement of the pion signal due to hadronic showering [94]. The total amount of shielding in front of the pion detector corresponds to 60 radiation lengths and 0.3 nuclear interaction lengths. In addition to the shielding in front of the pion detector, there is some shielding between the pion detector and the beam pipe to reduce further electromagnetic background.

The active part of the pion detector consists of 10 fused quartz cylinders, each 4 cm in diameter and 10 cm in length. These cylinders are placed azimuthally about the beamline and are tilted by 45° with respect to the beam axis, as is shown in Figure 2.36, which contains a picture of the arrangement. Charged particles passing through each quartz cylinder produce Čerenkov radiation, which is detected by an attached PMT. The physical region the quartz cylinders define corresponds to the same radial coverage as the Møller detector's active region ($r = 15 - 23.5$ cm). However, due to the large angular spread of the

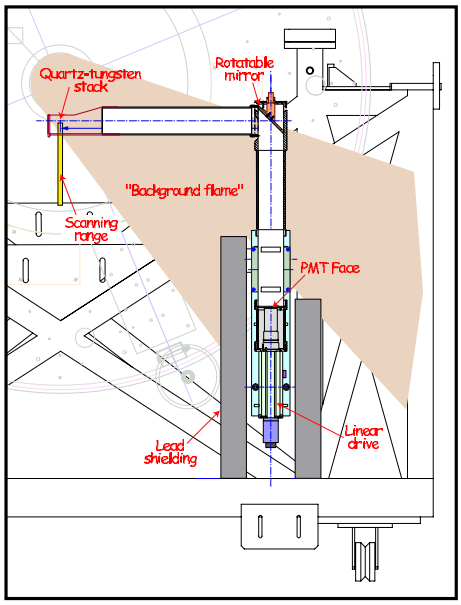

Figure 2.37: Polarimeter schematic. This is a view looking downstream in front of the profile wheel. The polarimeter is positioned between the profile wheel and the main integrating calorimeter. Note that the whole polarimeter assembly can move vertically 25 cm. The scanning range is indicated in the figure.

hadronic showers, the pion detector's acceptance is actually much larger than its physical size would suggest.

2.7.5 Polarimeter

The measured physics asymmetry involves a normalization to the beam polarization. The polarization is limited by strain relaxation in the source photocathode, and laboratory "bench" tests typically yield values of ∼80% or higher. However, these numbers are subject to modification due to a number of effects, some involving changes in the photocathode itself (for instance, zinc dopant levels could change due to migration during the heat cleaning process just before cathode activation in the gun) and others merely involving changes in operation (for instance, voltage changes in the CP and PS Pockels cells). Nevertheless, since these effects are unpredictable and fluctuate at some level over time, an *in situ* measurement of beam polarization is an absolute necessity. Furthermore, the error on this polarization measurement must be kept as small as possible, since it influences the final error on the physics asymmetry measurement.

Polarization is measured by inserting a remotely insertable/removable iron foil target

into the path of the beam.[15] There are several foils to choose from, with thicknesses of 20 μm, 50 μm, and 100 μm. In order to eliminate as much of the background flux as possible, all other targets are removed from the beam line, including the liquid hydrogen target and the wire array. To further suppress the background noise (radiative Mott's as well as inelastic *ep*'s), the calibration collimator is also inserted into the beamline, and special quadrupole magnet settings are used.

Helmholtz coils are then supplied with roughly 6 A of current, generating a magnetic field of ∼100 G that saturates the foil target, causing a known fraction of the atomic electrons to align their spins with the magnetic field. Some of the electrons in the beam scatter off the atomic electrons, pass through the spectrometer, and are detected by the polarimeter, described below. The cross section for this Møller scattering process is suppressed when the two electrons' spins are anti-aligned, so that the following asymmetry can be formed [95]:

$$A_{Fe} \equiv \frac{\sigma_{\uparrow\uparrow} - \sigma_{\uparrow\downarrow}}{\sigma_{\uparrow\uparrow} + \sigma_{\uparrow\downarrow}} = P_T \cdot P_B \cdot \frac{(7 + \cos^2\Theta)\sin^2\Theta}{(3 + \cos^2\Theta)^2} \cdot \cos\theta \tag{2.21}$$

where P_T and P_B are the polarizations of the foil target and beam, respectively, Θ is the scattering angle in the center-of-mass frame, and θ is the angle between the target and beam electrons' polarization axes. When writing $\sigma_{\uparrow\downarrow}$, for instance, the first arrow (in this case up) refers to the helicity of the target electron, while the second arrow (in this case down) refers to the helicity of the beam electron. Knowing P_T, Θ, and θ, the beam polarization P_B can be extracted from a measurement of A_{Fe}.

The Helmholtz coils are set up to provide a magnetic field longitudinal to the beam axis. The field lines will actually be parallel to the foil's surface, inducing the atomic electrons' spins to be aligned accordingly. Having the foils horizontal therefore maximizes the longitudinal analyzing power ($\cos\theta$) and hence the asymmetry, but increases the chances of multiple scattering, which dilutes the asymmetry. As a compromise, the foils are tilted by 20° with respect to the beam axis. Using rough numbers for the variables in Equation (2.21), $P_T \approx 8\%$, $P_B \approx 80\%$, $\Theta = 90°$ (corresponding to 4.6 mrad in the lab frame for a 45 GeV beam), and $\theta = 20°$, one finds $A_{Fe} \approx 0.047$.

The polarimeter, whose design is illustrated in Figure 2.37, sits between the profile monitor and the main integrating calorimeter. It is a quartz–tungsten Čerenkov calorimeter,

[15] All iron foil targets are actually composed of Supermendur, an alloy consisting of 49% iron, 49% cobalt, and 2% vanadium.

composed of six quartz plates (each 6 mm thick) sandwiched between seven tungsten plates (also 6 mm thick), the latter having reflective surfaces. These reflective surfaces aid in bouncing Čerenkov radiation created in the quartz plates towards a highly reflective cylinder that acts as a light guide. The light guide in turn brings the photons to a PMT, which then can be kept well away from the beamline and shielded by a lead enclosure. The quartz–tungsten plates are angled at 30° with respect to the beam axis to aid in light collection.

The holes in the calibration collimator allow for only a specific range of the Møller flux to pass, corresponding to 90° scattering in the center-of-mass frame, which maximizes the asymmetry. While this cuts down the overall rate, A_{Fe} is so large compared to the polarimeter's statistical resolution that this rate reduction doesn't matter. More importantly, the collimator suppresses the *ep* background. In order to measure the background that persists, the polarimeter is scanned radially, from 15 cm to 40 cm below the beamline. In addition, the signal contribution coming from background hitting the PMT is measured by rotating a mirror inside the light guide, which prevents any Čerenkov (i.e., signal) light from reaching the PMT. For polarimetry, the polarimeter sits at the Møller peak, roughly 18 cm from the beamline. During nominal physics running, however, the polarimeter sits at 50 cm from the beamline, so that it does not interfere with the primary data collection for the parity violation measurement.

2.8 Electronics and Data Acquisition

2.8.1 Description of ADC's

Custom 16-bit integrating analog-to-digital converters (ADC's) were designed and built specifically for this experiment [96]. The main design goal was to allow for the integration of signals coming from various types of devices (toroids, BPM's, PMT's, etc.) while introducing minimal electronic noise. The basic schematic for the ADC, designed to be fully VME-compatible, is shown in Figure 2.38. A single ADC board contains six channels each of which accepts twin-ax inputs.

The operational amplifier (op-amp) directly after the channel input can either be run in differential mode or single-ended mode. In differential mode, the difference between the two leads of the input twin-ax cable gets amplified. If one of these leads is in fact ground, then the ground is said to be left "floating." In single-ended mode, there are no floating

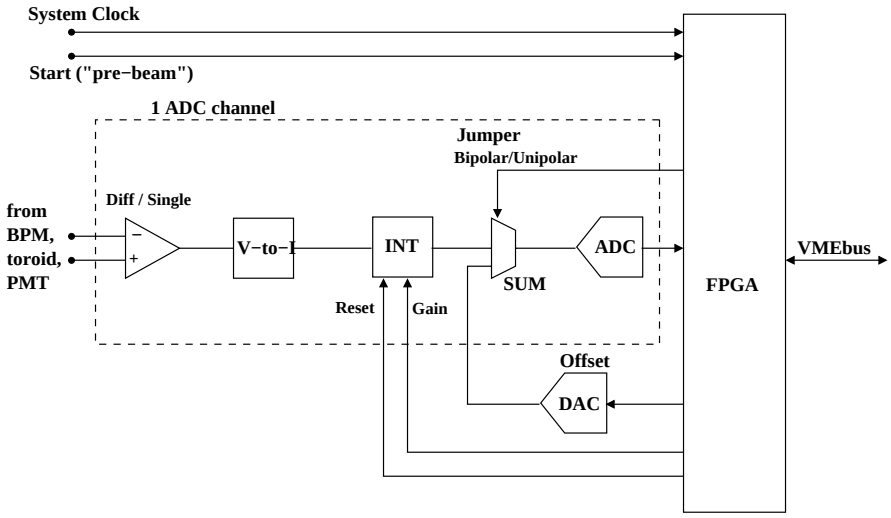

Figure 2.38: Schematic of a 16-bit VME ADC. The experiment uses ADC boards containing six ADC channels each. The details of a single ADC channel are shown. All of the channels on a given board share a common board memory and timing settings.

grounds. The second lead of the input is connected to the ADC's ground.

After the op-amp, the signal first goes through a voltage-to-current conversion resistor. The signal then encounters the capacitor of an integrator stage. This integrator output can be summed with the output from a digital-to-analog converter (the so-called DAC offset), which was designed to achieve sub-channel resolution in the ADC's. In reality, this feature was never needed, and even degraded the resolution slightly rather than improved it (due to a presumed design flaw), and so was never used.

A jumper on the ADC board can be toggled to select between running in bipolar and unipolar modes. Bipolar mode allows for the integration of both positive and negative signals by adding a constant offset voltage at the sum stage. This necessarily halves the amplitude of the maximum input voltage allowed, however. In unipolar mode, the input signal must be either all-negative or all-positive, which allows for the integration of larger input voltages.

After the sum stage, the signal is digitized by the ADC and stored in the Field Programmable Gate Array (FPGA) memory buffer. Along with the gain, the FPGA sends various timing triggers to the integrator stage; the gain and these triggers can all be changed remotely using the VMEbus system. The timing triggers are illustrated in Figure 2.39. The "start" trigger (also called the "pre-beam" trigger) defines the start of each pulse. Some time t_r later, the "reset" trigger for the integrator stage is released. This opens a

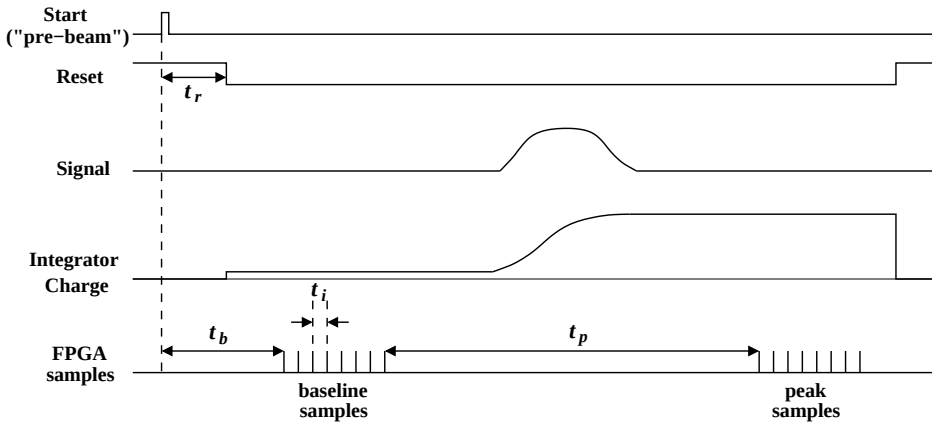

Figure 2.39: Definitions of ADC timing parameters. The timing should be set such that the signal arrives several microseconds after the last baseline sample has been taken.

relay allowing charge to accumulate on the integrator's capacitor.[16] The ADC then begins sampling and storing what are called "baseline" samples at time t_b. A total of N baseline samples are taken, where N can be 1, 2, 4, 8, or 16 (this too can be remotely set using the VMEbus system). After the signal has finished integrating on the capacitor, the ADC again begins sampling and storing what are called "peak" samples, at time t_p after the last baseline sample. The same number of peak samples as baseline samples are automatically taken. Both individual baseline and peak samples are separated by time t_i (almost always t_i is set to be 1 μs). After all peak samples have been taken, the reset trigger is again engaged, shorting the capacitor and releasing all accumulated charge. The FPGA then subtracts the average of the baseline samples from the average of the peak samples, storing the difference in memory which can be read out through the VMEbus system.

2.8.2 Types of Signals

The above description pertains to all of the VME ADC boards used in the experiment. The same basic boards are used to integrate signals from the toroids, the Møller and ep detectors, the BPM's, the luminosity monitor, the pion detector, and the SLM. The boards only need minor modifications (some of them in the hardware, some of them in the software) to accommodate all of these different signals. Essentially, the integration gate (defined by t_p), along with the integrating capacitor and/or the voltage-to-current conversion resistor, needs adjustment depending on how long the signal lasts. The board's gains can also be

[16] Merely opening the relay induces a small charge on the capacitor, even though the signal has not arrived yet. This is reflected in Figure 2.39.

modified to accommodate input signals of varying voltage levels. In addition, each board may be run in differential or single-ended mode, and in bipolar or unipolar mode, depending on the type of input.

The various modifications listed above, minor though they may be, greatly influence the noise performance of the board, i.e., the resolution it is capable of putting out for a particular input signal. This is because the boards possess a strong noise component in the MHz frequency range, but much lower noise at lower frequencies. Depending on the board's gain settings, this fast noise can end up getting amplified. For relatively weak input signals in the MHz frequency range, for example, the noise can end up being several channels. This noise is referred to as pedestal noise, since it is the natural noise of the electronics itself, and is present even in the absence of signal (a pulse with no beam is referred to as a pedestal pulse). Since the VME ADC's are 16-bit, this works out to a pedestal-limited resolution of approximately $3/2^{16} \approx 50$ ppm, if the ADC is operated at full range. If the ADC is not operated at full range (which is the usual case), the situation gets even worse, with the limit being set at 100 ppm, perhaps. In addition, if the ADC is being operated in bipolar mode, the resolution gets worse (higher) by a further factor of two.

Unfortunately, since the beam pulse length is on the order of a few 100 ns, most beam signals naturally have a significant signal component in the MHz range. Depending on the specific measurement, per-channel resolution of better than 100 ppm may not be an absolutely necessity, in which case the noise is not a problem. This is the case for the pion detector, the SLM, and the BPM's, for instance. For other signals, however, per-channel resolution of better than 100 ppm is important to have, in which case something must be done in order to improve the signal-to-noise ratio. This may include simply amplifying the signal before it is fed into the ADC, so that high ADC gain settings (which tend to amplify pedestal noise) are not needed.[17] It may also include shifting the signal out of the MHz frequency range through the use of transformers, as is done in the toroid electronics and, as will be seen below, in the Møller and *ep* detector electronics.

[17] This amplification, if needed, should be done as close to the signal's origin as possible, to avoid amplifying transmission noise picked up in the cables.

Toroid Signals

As was described in Section 2.4.1, the raw toroid signals are connected to a resistor and capacitor in series, forming an RLC circuit. The output therefore rings with characteristic frequency $\omega = 1/\sqrt{LC}$ and decays with a settling time τ that is proportional to the circuit's quality factor Q (explicitly, $\tau = 2Q\sqrt{LC}$).[18] Since the ADC's are limited by pedestal noise, making the Q as large as possible (i.e., stretching the signal out for as long as possible) improves the overall resolution of the measurement, by increasing the charge on the integrator and thus the signal-to-noise ratio. Stretching the signal out also shifts the signal into a lower frequency range where the noise performance of the ADC boards is better. In practice, the Q is adjusted so that the toroid signals last for several milliseconds, of which the first millisecond is integrated in the ADC's. The Q is so large, in fact, that the signal may have not died out by the time the next pulse starts, especially at 120 Hz, when pulses are only separated by 8 ms. Such "hysteresis" effects, whereby the charge measurement for one pulse reflects at some level the charge measurement of the previous pulse, would obviously be catastrophic. The simple solution is not to reduce the Q (which would degrade the resolution), but to suddenly overdamp the toroid signal, using a large resistor connected to the toroid electronics by a relay, at some point after the ADC has finished its integration, typically at 3 ms.

The toroid signals are amplified by preamplifiers physically located nearby the toroids themselves, avoiding amplification of transmission noise picked up over the ~100 feet of cable connecting the toroids to the actual ADC's. The ringing toroid signals must be rectified before entering the ADC's, which are operated in single-end mode. As modified for toroid signal input, the ADC's exhibit excellent signal-to-noise behavior, introducing pedestal noise of 1 channel or less. If the full range of the 16-bit ADC's were to be used, this could allow for toroid resolutions of roughly 15 ppm; resolutions nearly this good were actually observed in special beam tests. However, in reality the full range of the ADC's is not exploited, so that pulse-to-pulse resolutions in the 50 to 60 ppm range are more typical.

[18] Technically the circuit rings with frequency $\omega = \sqrt{1/LC - (R/2L)^2}$, but for small values of R (corresponding to the case of small damping), one can write $\omega \approx 1/\sqrt{LC}$ with negligible error.

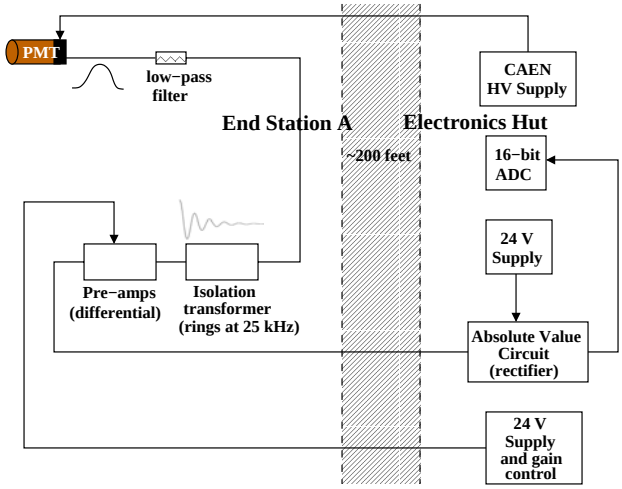

Figure 2.40: Electronics schematic for the main integrating calorimeter. The signals originate from the detector PMT's in the MHz range, but are shifted into the kHz range using isolation transformers.

BPM Signals

As described in Section 2.4.2, the BPM processing electronics produce decaying, non-oscillatory signals. High Q cavities lead towards better resolution, but detune quicker and hence need to be watched more carefully. The Q of the cavities are set such that the BPM signals die out several hundred nanoseconds after the end of the beam pulse. The signals are therefore in the MHz range. In addition, the signals can be positive or negative, so their ADC boards must operate in bipolar mode. Also, having a sufficient dynamic range (e.g., ± 1 mm) is important, so a signal produced when the beam is at 100 μm, for instance, cannot saturate the ADC. All of these constraints ensure that the resolutions of the ADC boards are not too impressive, at the 100 ppm level. However, better resolutions are not needed, since at the very most the experiment only calls for pulse-to-pulse position resolutions of 1 μm, which is equal to 1000 ppm resolution if 1 mm is defined as "full range." Beam motions greater than ± 1 mm are not expected.

Møller and ep Detector Signals

The Møller and *ep* detector signals originate from the PMT's in the MHz range and with values less than 0.5 V (this avoids anode saturation nonlinearities, as discussed in Section 2.7.1). The detector signals require board settings that produce unacceptable levels of noise (at the 100 ppm level, as with the BPM ADC's). To reduce the noise, the calorimeter

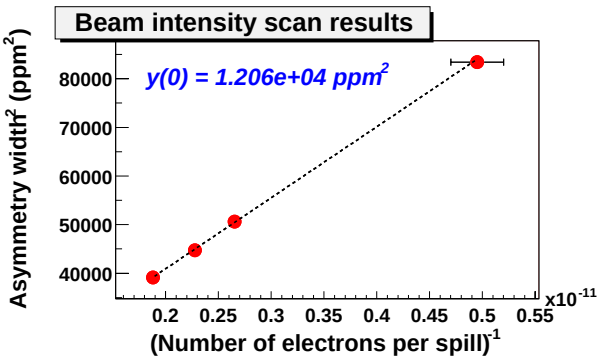

Figure 2.41: The width of the Møller detector asymmetry distribution is squared and plotted versus the inverse of the beam intensity. The y-intercept reveals the square of the inherent electronics noise. The beam energy was not changed for this intensity scan.

PMT signals are made to look more like the toroid signals through the use of the electronics shown schematically in Figure 2.40. Isolations transformers (RLC circuits) sit in a chassis located near the main integrating calorimeter. Connecting each PMT output to a transformer shifts the signal to a lower frequency range, producing an oscillatory, decaying signal like that of a toroid. It also results in improved noise performance in the ADC boards. In addition, it prevents the formation of ground loops, since the grounds of all PMT channels are connected via the tubes' high voltage supply.

After the isolation transformers, each PMT signal gets amplified by a set of preamplifiers whose gains are remotely adjustable. Typically an overall gain of 100 is used. After traveling along roughly 200 feet of twisted pair cable, the signal gets rectified in an absolute value circuit before being fed into one channel of an ADC board identical to those used for the toroid signals.

The width of the measured detector asymmetry distribution receives contributions from a number of sources, including beam jitter, counting statistics, and electronics noise. Regression (or beam dithering) corrects for the effects of beam fluctuations, as will be discussed in the following chapter. The contribution from electronics noise can be observed by plotting the square of the width of the regressed detector asymmetry distribution versus the inverse of the beam intensity. Such a plot is given in Figure 2.41. Data was taken at a wide variety of beam intensities, from roughly 2×10^{10} e^- per pulse to over 5×10^{11} e^- per pulse. As one can see, for higher intensities (moving left along the horizontal axis), the width of the asymmetry distribution gets smaller, since the contribution from statistical fluctuations is reduced. By extrapolating the data to the y-intercept, one takes the limit of infinite

beam current, which leaves only the contribution from electronics noise. From the plot, the electronics noise contribution to the detector width is found to be $\sqrt{12060} = 110$ ppm. The overall width of the Møller detector, as will be seen in Section 4.4, is typically $190-200$ ppm, meaning that the electronics noise contribution, while smaller than for the BPM ADC's, is still a non-negligible fraction of the total width.

While the Møller and *ep* detector electronics have not quite met the design goal of achieving a noise performance of roughly 40 ppm, they do meet the linearity goal set at 99%. Because of the measures taken to suppress cathode and anode nonlinearities, the linearity of the entire system (including the electronics) is found to be $99 \pm 1\%$ [97]. This result applies to the local signal region relevant for nominal physics running. For lower flux levels (for instance, those produced by the special calibration runs discussed in Section 4.8.2), very small signals are generated, and the linearity of the system is much worse, at roughly 80% [98]. The main reason behind these low-level nonlinearities is that the rectifier circuit tends to distort small signals.

Luminosity Monitor Signals

Like the BPM signal, the luminosity monitor signal sits in the MHz frequency range. However, as mentioned in Section 2.7.2, its raw signal is of a sufficient size such that amplifiers are unnecessary. In addition, its anode and cathode voltages are read in differentially, so that transmission noise pickup is further suppressed. Its signal is unipolar and does not fluctuate much, so that its ADC's can be run nearly at full range. All of these factors combine to give the luminosity monitor ADC's surprisingly good noise performance, similar to that of the toroids.

2.8.3 Data Acquisition

The experiment employs a data acquisition (DAQ) system, developed and implemented specifically for the ESA fixed target program at SLAC [99], consisting of numerous VME crates distributed over the entire SLAC site, communicating with one another through a combination of both network (TCP/IP) and fiber optic connections. A general outline of the DAQ system is given in Figure 2.42. Two "master" VME crates, one located at ASSET and the other in ESA, gather data from all of the 16-bit VME ADC's (described in the previous section) along with all of the 11-bit CAMAC ADC's. In addition, these

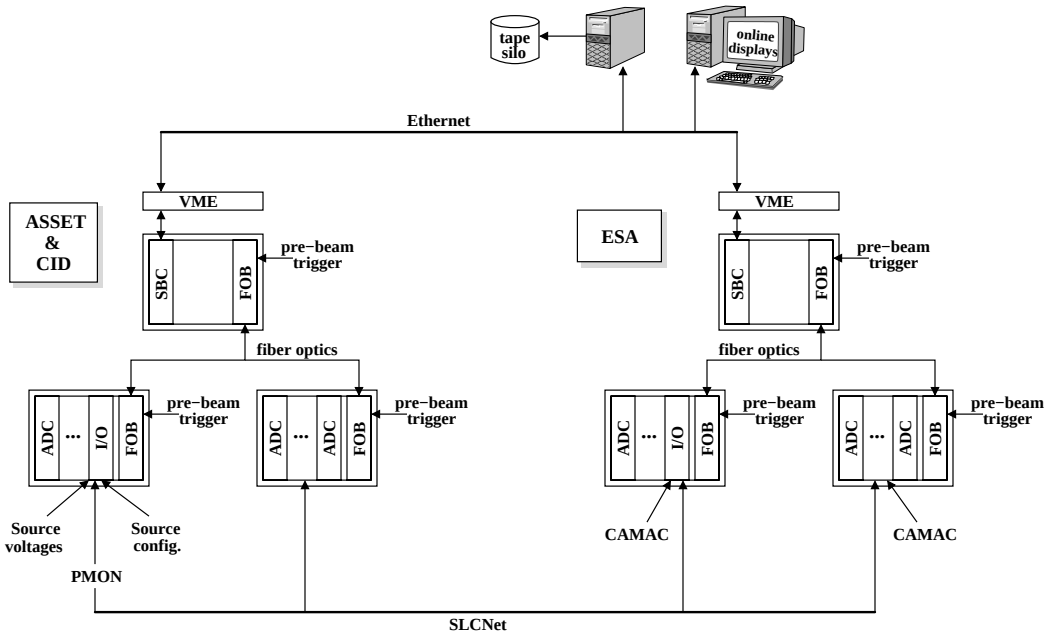

Figure 2.42: Overview of data acquisition (DAQ) system. The two "master" crates at ASSET and in ESA communicate to other VME crates via fiber optic cabling. For purposes of clarity, only a few VME crates are pictured. Each crate receives data from VME modules within, which may be I/O modules, 16-bit ADC modules, or modules for communicating with CAMAC crates (to name just a few examples). The helicity bits are generated by the PMON controller at CID and broadcast to other receiver units over the SLCNet.

crates receive a wealth of other information, as well. For instance, the status of the various insertable/removable optics components (e.g., half-wave plates, asymmetry inverter) as well as the Pockels cell voltages at the source are sent to an I/O module in a VME crate at CID. The helicity of the beam, as determined by PMON, is rf modulated and broadcast (after a one-pulse delay) to the two master crates.

Data from the two master crates are merged, loosely synchronized, and logged to SLAC's tape silos for storage by an ESA computer. The total amount of data logged for each 120 Hz event is roughly 2.5 kB. During nominal physics data taking, typically 20 GB are logged each day. Several real-time analysis processes perform simple online monitoring functions, in order to ensure good data quality. For instance, one process displays detector signal levels (useful for checking for dead PMT's and/or power supply problems), while another displays mean beam intensity and position, along with beam jitter numbers. This latter display is simultaneously monitored both by ESA experimenters as well as machine operators, and is indispensable for maintaining good beam quality.

The DAQ system was designed to operate at up to 120 Hz, the maximum beam rate.

The whole system is triggered by the beam trigger. The DAQ system is even triggered for PEP pulses (when BaBar is running concurrently), although these are not used for physics analysis. In addition to normal physics pulses and PEP pulses (also called "witness" pulses), there are pedestal pulses. During pedestal pulses, a special source permissive is disabled, so that the source laser flash lamps fire 50 μs out of time with the accelerator rf. Hence, no electrons are accelerated down the linac. Pedestal pulses therefore contain no electrons, but otherwise the machine is operating nominally. As a result, they are used to calibrate the responses of all read-out devices (toroids, BPM's, detector PMT's, etc.) when everything in the experiment is up and running but the actual beam. Pedestals are averaged for each timeslot separately, since the two timeslots are sensitive to different phases of the 60 Hz line noise. For each timeslot, the average of the five previous pedestal pulses is subtracted from every physics event.

The next step after pedestal subtraction is to generate helicity-correlated differences and asymmetries for all of the various read-out devices. As will be seen in the next chapter, it will be necessary at this time to find a way to correct for the effects of beam fluctuations, which not only are capable of generating false asymmetries, but also enlarge the widths of the detector asymmetry distributions.

Chapter 3

Data Analysis

3.1 Overview

In order to minimize sensitivity to slow drifts, the experiment must switch rapidly between helicity states, forming helicity pulse pairs. This rapid helicity switching is done in a pair-wise pseudo-random manner, forming so-called "timeslots" to remove sensitivity to the 60 Hz line noise. Approximately every 30 ms then (depending on the beam rate), another right–left detector asymmetry can be formed. The goal of the experiment is to achieve a measurement of the parity–violating cross section asymmetry $A_{\rm PV}$ in Møller scattering, which can be written as:

$$A_{\rm PV} = \frac{\sigma_R - \sigma_L}{\sigma_R + \sigma_L} \tag{3.1}$$

where σ_R and σ_L refer to the cross sections (proportional to the scattering rates) for right and left helicity beams, respectively. In order to obtain the correct cross-section asymmetry, false asymmetries due to helicity correlations in the beam parameters must first be removed. The main purpose of this chapter is to describe how this is done, as well as to discuss the means by which a systematic error can be ascribed to this correction procedure.

3.2 Beam Fluctuations

Assuming the target and main integrating calorimeter are perfectly co-aligned, a properly centered beam will produce little or no azimuthal variation in the signal levels of the detector PMT's. In other words, each PMT in a given detector ring (in, mid, out, or ep) should produce nearly the same size signal for the same high voltage. If, however, the beam were to shift slightly in any direction, the signal sizes in some of the PMT's will increase while

those in others will correspondingly decrease. The calorimeter therefore takes advantage of the azimuthal symmetry of the scattered electron flux, cancelling out (to a large extent, at least) the effects of beam fluctuations.[1]

The cancellation of beam effects due to azimuthal symmetry is not perfect and cannot be taken for granted, since at any time an unforeseen problem could affect one or more of the detector channels, rendering them unusable and thereby destroying the symmetry. For this reason, it is still necessary to monitor the sensitivity of the detector to beam fluctuations. Since the Møller and *ep* detectors consist of a total of 60 channels, one must first decide how to handle the data. For instance, when correcting for beam-related false asymmetries, should all "good" channels (meaning those not affected by hardware failures) first be averaged, taking full advantage of the cancellation due to azimuthal symmetry? Because each channel actually requires a different statistical weight, as will be made clear in Section 3.3, the channels must not be averaged first. Instead, beam fluctuations should be corrected for on a channel-by-channel basis, the asymmetry in each detector channel being computed separately. Averages of channels, using the appropriate channel weights, can then be calculated later. In fact, various linear combinations of channels, some being very sensitive to fluctuations in individual beam parameters, can be studied in order to attribute an overall systematic uncertainty to the correction procedure.

3.2.1 General Formalism

As long as the dependence of a given detector channel's signal d_i on each of the beam parameters is approximately linear over the beam jitter range, one may write the following Taylor expansion for that channel's measured asymmetry A_i:

$$A_i = \frac{d_{i_R} - d_{i_L}}{d_{i_R} + d_{i_L}} = A_i^{phys} + A_Q + \sum_{\xi} \alpha_{\xi i} \Delta\xi \tag{3.2}$$

where ξ runs over all beam parameters except charge, namely $\xi = \{E,\, x,\, y,\, x',\, y'\}$, and the α parameters are correlation coefficients (called "detector slopes") relating the sensitivity of each detector channel to each of the beam parameters. In other words, $\alpha_{\xi i} = \partial A_i / \partial \xi$. The units of these detector slopes are therefore ppm/MeV, ppm/μm, and ppm/μrad, or

[1] This reasoning is most applicable when the spectrometer quadrupoles are off. Even with the quadrupoles on, however, so that the signal flux is radially focused, cancellation of beam fluctuations' effects will still occur.

something equivalent. The quantities A_Q and $\Delta\xi$ are the charge asymmetry and the helicity-correlated difference in beam parameter ξ, respectively. Equation (3.2) says that the asymmetry actually measured by detector channel i is equal to the true underlying physics asymmetry in the scattering process being studied, plus false asymmetries due to helicity correlations in the beam parameters.

The first step in removing false asymmetries due to beam fluctuations is to normalize each detector signal to the beam charge q as measured by one of the alcove toroids. To first order, this is equivalent to subtracting out $A_Q \equiv (q_R - q_L)/(q_R + q_L)$ in Equation (3.2) above. This is done before all other data processing, so that Equation (3.2) should really be rewritten as:

$$A_i = \frac{d_{i_R}/q_R - d_{i_L}/q_L}{d_{i_R}/q_R + d_{i_L}/q_L} = A_i^{phys} + \sum_{\xi} \alpha_{\xi i} \Delta\xi \tag{3.3}$$

These detector asymmetries are referred to as charge-normalized asymmetries. Note that to first order, normalizing the detector signals does not change the detector slopes $\alpha_{\xi i}$.

The second step in removing false asymmetries due to beam fluctuations is to somehow measure the detector slopes $\alpha_{\xi i}$. This is done two different ways: through linear regression and by the process of beam dithering. Each of these methods will briefly be described in the sections immediately following. Once the detector slopes are known, the true physics asymmetries A_i^{phys} can be uncovered and used for further processing.

3.2.2 Linear Regression

In its simplest form, linear regression involves plotting the charge-normalized detector channel asymmetries versus the helicity correlated differences in each of the beam parameters. An example of this is illustrated in Figure 3.1. First the charge-normalized asymmetry in a given detector channel is plotted against ΔE. The slope is found, and then $A_i - \alpha_{iE} \cdot \Delta E$ is plotted against Δx. This slope is found, and then $A_i - \alpha_{iE} \cdot \Delta E - \alpha_{ix} \cdot \Delta x$ is plotted against Δy. This procedure continues so on and so forth until all slopes $\alpha_{\xi i}$ have been found. In fact, the procedure has to iterate and begin again, this time beginning with all the slopes that have already been found, and looking for small deviations from these slopes. That is, once all the slopes have been found, $A_i - \sum_{\xi} \alpha_{\xi i} \cdot \Delta\xi$ is plotted against ΔE, and a small deviation $\delta\alpha_{iE}$ from the slope α_{iE} already found is obtained. This is used to get the *new* slope, $\alpha_{iE} \rightarrow \alpha_{iE} + \delta\alpha_{iE}$. The same procedure continues for all of the beam parameters,

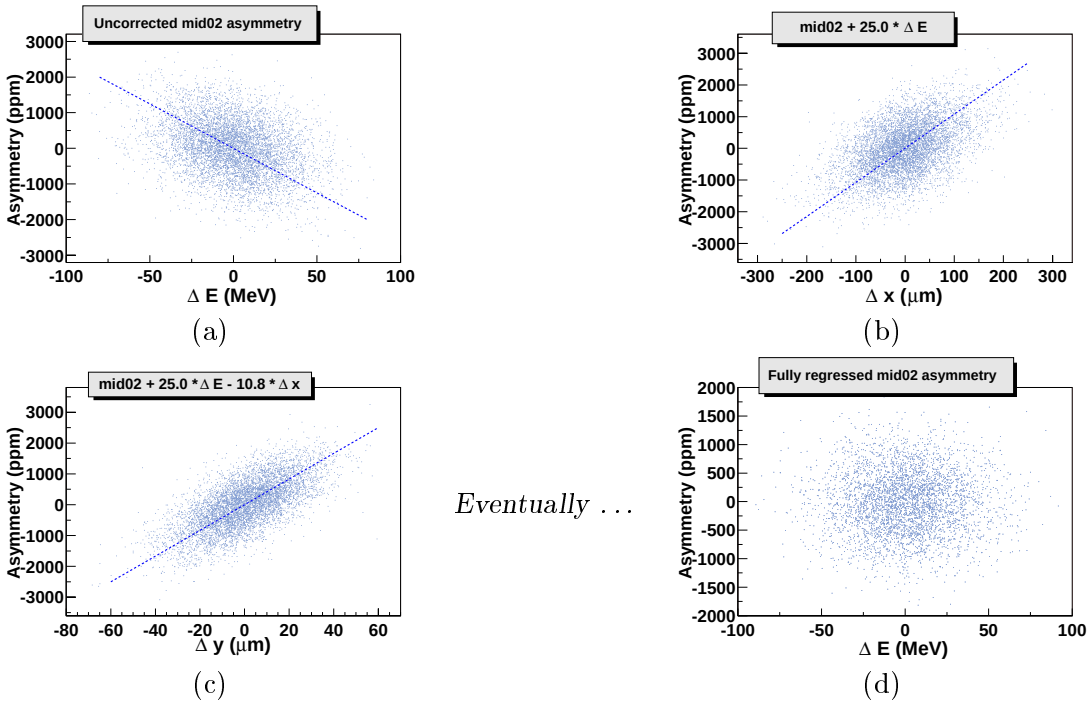

Figure 3.1: Linear regression example, for the charge-normalized asymmetry in the third PMT of the "mid" ring (mid02). In Figures (a)–(c), the asymmetry is corrected for an increasing number of beam parameters. Note the scatter of the data becoming smaller and smaller as regression proceeds. Figure (d) shows the fully regressed asymmetry plotted against ΔE, showing no dependence. After regression, the asymmetry should not depend on *any* beam parameter.

updating the slopes as the necessary deviations are found, until during one iteration none of the slopes changes by more than some tolerance limit. At this point, the slopes have all converged to their final values, and the detector channel asymmetry exhibits no further dependence on any of the beam parameters, as illustrated in Figure 3.1(d).

In addition to being tedious and inefficient, the above procedure yields unphysical slope values, since correlations among the beam parameters ensure that the values obtained for the slopes depend on the order in which they are found. It would be much better to perform a one–pass fit in which all of the slopes are found simultaneously. Not only would this be much faster and more economical with CPU time, but it would yield slopes whose values actually reflected some sense of physical reality. That is, suppose there existed a correlation between E and x, such as there exists due to dispersion. With the iterative procedure described above, if the α_x slope happened to be found first, it would tend to "absorb" some of the value of the α_E slope. The exact amount it would absorb would depend on the particulars of the scenario, and could range from all of it to just part of it. Nevertheless,

with the non-iterative procedure about to be described, the correct, physical slopes will be found.

The non-iterative procedure uses matrix inversion to find the best fit for the detector slopes $\alpha_{\xi i}$. This is done by first defining a χ^2 for the detector slopes' solutions:

$$\chi^2 = \left\langle \left(A_i - \sum_{\xi} \alpha_{\xi i} \cdot \Delta\xi \right)^2 \right\rangle \tag{3.4}$$

Here A_i is the charge-normalized asymmetry in detector channel i, the sum runs over all beam parameters except charge, and the brackets denote averaging over all pairs. Minimizing the function χ^2 with respect to $\alpha_{\xi i}$ gives the following matrix solution:

$$\begin{pmatrix} \langle A_i \cdot \Delta E \rangle \\ \langle A_i \cdot \Delta x \rangle \\ \vdots \\ \langle A_i \cdot \Delta y' \rangle \end{pmatrix} = \begin{pmatrix} \langle \Delta E^2 \rangle & \langle \Delta E \cdot \Delta x \rangle & \cdots & \langle \Delta E \cdot \Delta y' \rangle \\ \langle \Delta x \cdot \Delta E \rangle & \langle \Delta x^2 \rangle & \cdots & \langle \Delta x \cdot \Delta y' \rangle \\ \vdots & \vdots & \ddots & \vdots \\ \langle \Delta y' \cdot \Delta E \rangle & \langle \Delta y' \cdot \Delta x \rangle & \cdots & \langle \Delta y'^2 \rangle \end{pmatrix} \cdot \begin{pmatrix} \alpha_{iE} \\ \alpha_{ix} \\ \vdots \\ \alpha_{iy'} \end{pmatrix} \tag{3.5}$$

This can be re-written using the following simplified notation:

$$\vec{A}_i = \mathbf{M} \cdot \vec{\alpha}_i \tag{3.6}$$

where $\vec{A}_i \equiv (\langle A_i \cdot \Delta E \rangle, \langle A_i \cdot \Delta x \rangle, \ldots, \langle A_i \cdot \Delta y' \rangle)$, $\vec{\alpha}_i \equiv (\alpha_{iE}, \alpha_{ix}, \ldots, \alpha_{iy'})$, and the covariance matrix $\mathbf{M}$ has components $M_{ij} = \langle \Delta\xi_i \cdot \Delta\xi_j \rangle$. The solution is of course:

$$\vec{\alpha}_i = \mathbf{M}^{-1} \cdot \vec{A}_i \tag{3.7}$$

Thus by keeping track of the correlations between A_i and all of the beam parameter helicity correlations $\Delta\xi$, it is possible to solve for all of the detector slopes $\alpha_{\xi i}$ simultaneously. Since regression works by calculating average values for various covariance matrices, it obviously needs a finite statistical sample with which to work. A regression cycle of 10k pairs was chosen as a good compromise between wanting to reduce statistical fluctuations on the computed slopes, while still retaining sufficient sensitivity to slow drifts and other possible effects that could fluctuate over time.

In writing Equations (3.4)–(3.7), one minor point has been glossed over. Namely, the

quantities A_i and $\Delta\xi$ should actually have their averages subtracted out. In other words, Equation (3.4) should technically be written as:

$$\chi^2 = \left\langle \left(\left\{ A_i - \overline{A_i} \right\} - \sum_{\xi} \alpha_{\xi i} \cdot \left\{ \Delta\xi - \overline{\Delta\xi} \right\} \right)^2 \right\rangle \tag{3.8}$$

where $\overline{A_i}$ is the average charge-normalized asymmetry in detector channel i for that particular regression cycle, and $\overline{\Delta\xi}$ is the average helicity-correlated difference in beam parameter ξ. This will introduce terms like $\overline{A_i} \cdot \overline{\Delta E}$ and $\overline{\Delta E} \cdot \overline{\Delta x}$ into Equation (3.5). Because all helicity-correlated differences in the experiment are small compared to the beam jitter, such terms will only result in very minor corrections to the detector slopes $\alpha_{\xi i}$. Thus, in the interests of clarity, these higher-order terms have been dropped from the equations.

In principle, there are a number of effects which can interfere with the successful implementation of linear regression. In particular, finite device resolution can cause additional noise in the diagonal elements of $\mathbf{M}$, effectively increasing them, while leaving $\vec{A}_i$ uncompensated. In theory, this can result in a systematic reduction in the sizes of the calculated slopes. In other words, for a given beam parameter ξ, finite resolution in the device monitoring ξ will always result in the calculated slope $\alpha_{\xi i}$ being smaller than the true slope. This can produce a systematic error, though studies have shown that such effects are probably negligible. Nevertheless, an additional method of arriving at the detector slopes would provide a useful check.

Another important reason why some other method of getting $\vec{\alpha}_i$ is necessary is that linear regression fails when one or more parameters begin to dominate over all others. For instance, suppose the jitter in y' suddenly falls dramatically compared to all of the other beam parameters. Effectively, then, the last column and the last row of $\mathbf{M}$ will be small compared to all other columns and rows. As the jitter in y' gets smaller (or, equivalently, as the jitter in all of the other beam parameters gets larger), the last column and row will look more and more like a zero column and a zero row. Thus $\mathbf{M}$ will grow more and more singular. As this happens, the matrix inversion will blow up, and the whole process will fail. What is happening in this scenario is that the beam jitter is failing to span the whole parameter space. While data cuts can mitigate this potential problem, an independent method of finding the detector slopes, one that does *not* rely on beam jitter to span the

parameter space, is still a good idea.

3.2.3 Beam Dithering

Beam dithering implements much the same procedure as linear regression, but in a much more controlled way. It uses the magnets and klystron phase described in Section 2.3 to introduce small perturbations in the beam's position, angle, and energy, while simultaneously monitoring the responses of the BPM's and detectors. The goal of this analysis is to arrive at the detector slopes $\vec{\alpha}_i = \partial A_i/\partial \xi_j$, which is accomplished by again defining a χ^2 function for minimization:

$$\chi^2 = \sum_{j=1}^{9} \left[\frac{\partial d_i/q}{\partial C_j} - \sum_{\xi} \left(\frac{\partial d_i/q}{\partial \xi} \frac{\partial \xi}{\partial C_j} \right) \right]^2 \tag{3.9}$$

where C_j refers to the j^{th} dithering object. Since there are eight magnets and one klystron phase, $j = \{1, 2, \ldots, 9\}$. Here i is a free index, referring to a particular detector channel. Notice that $\partial(d_i/q)/\partial\xi$ is written, and <u>not</u> $\partial A_i/\partial\xi_j$. The former can be referred to as "unnormalized" detector slopes, since they are related to the actual detector slopes by a simple normalization constant (see below). Minimizing χ^2 with respect to the choice of unnormalized detector slopes $\partial(d_i/q)/\partial\xi$ yields the following solution:

$$\sum_{k=1}^{9} \left(\frac{\partial d_i/q}{\partial C_k} \frac{\partial \xi_j}{\partial C_k} \right) = \sum_{l} \frac{\partial d_i/q}{\partial \xi_l} \sum_{k=1}^{9} \left(\frac{\partial \xi_l}{\partial C_k} \frac{\partial \xi_j}{\partial C_k} \right) \tag{3.10}$$

which can be rewritten as $\mathbf{D} = \beta \cdot \mathbf{B}$ if the following matrices are defined:

$$\begin{aligned} D_{ij} &= \sum_{k=1}^{9} \frac{\partial d_i/q}{\partial C_k} \frac{\partial \xi_j}{\partial C_k} \\ B_{ij} &= \sum_{k=1}^{9} \frac{\partial \xi_i}{\partial C_k} \frac{\partial \xi_j}{\partial C_k} \\ \beta_{ij} &= \frac{\partial d_i/q}{\partial \xi_j} \end{aligned} \tag{3.11}$$

Solving Equation (3.10) by inverting the matrix $\mathbf{B}$ yields:

$$\beta = \mathbf{D} \cdot \mathbf{B}^{-1} \tag{3.12}$$

This is the fundamental equation for beam dithering, yielding the unnormalized detector slopes $\partial(d_i/q)/\partial\xi$. As stated earlier, these are related to the more useful normalized detector slopes $\partial A_i/\partial\xi_j$ by a simple normalization factor:

$$\frac{\partial A_i}{\partial \xi} = \frac{1}{2\langle d_i/q\rangle}\frac{\partial(d_i/q)}{\partial \xi} \tag{3.13}$$

where $\langle d_i/q\rangle$ is the average of the charge-normalized detector signal for channel i. Thus beam dithering can be used as a completely independent means of arriving at the detector slopes $\partial A_i/\partial\xi_j$, which can then be used to correct for false asymmetries. A new set of slopes is obtained at the end of every completed dithering cycle.

There is one main caveat to consider when contemplating the beam dithering procedure outlined above. Namely, the beam dithers C_j should completely span the beam parameter space. That is, at least five of the vectors $\partial\xi/\partial C_j$ (where ξ is understood to run over all of the beam parameters) should be linearly independent. If there are not five such beam dithers producing linearly independent responses in the beam parameters ξ, then the matrix $\mathbf{B}$ above will become singular. Thus the matrix inversion in Equation (3.12) will fail, and the detector slopes will remain unattainable. In practice, an analysis is performed that finds the optimal (i.e., most linearly independent) set of five dithers, and it is only these that are used for the full analysis. Thus the sums in Equations (3.9)–(3.11) only proceed over these five dither objects, and not the full set of nine.

Beam Dithering Example

The treatment above describes beam dithering mathematically. It is useful, however, to see an example of what beam dithering involves experimentally. One object (either one of the corrector magnets or the klystron phase) is dithered at a time. Dithering occurs on a very fast timescale. A voltage signal, called the dithering bit, encodes both which object is being dithered, and the amount by which the object is being dithered. Each dither follows the same basic three–step pattern shown in Figure 3.2. At any one step in the pattern, the dither value is held constant for roughly 10 pulses (or $\sim$0.1 sec at 120 Hz) before moving on to the next step in the pattern. The overall pattern repeats for roughly 2 seconds, at which point the beam ceases to be dithered for approximately 30 seconds. Then the next object in the dithering cycle is dithered. A dithering cycle is complete when all eight magnets

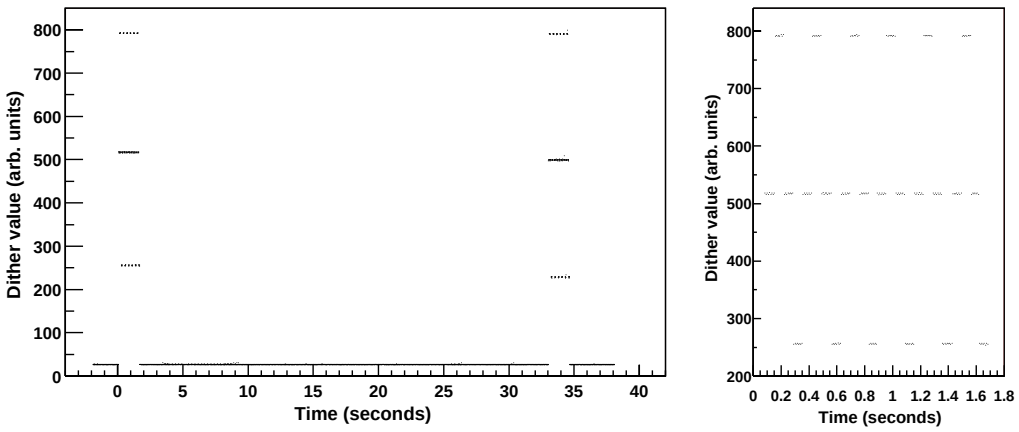

Figure 3.2: Beam dithering pattern. The plot on the left shows the "macroscopic" features of the overall pattern. One object is dithered for ∼2 sec, and then a different object is dithered ∼30 sec later. The plot on the right shows the "microscopic" features of the dithering pattern. The dithering occurs in a simple three–step pattern, with each step corresponding to ∼0.1 sec.

and the klystron phase have been dithered. The klystron phase actually gets dithered four times in a row, in order to reduce the uncertainty associated with that particular slope.

Note that the 2–seconds–on, 30–seconds–off pattern implies that the beam is being actively dithered about 7% of the time. During Run I, the beam dithering procedure was still being refined. While problems were being addressed, the amount of beam dithering was less, roughly 4%. As will be seen in the next chapter, all data for which the beam was actively being dithered is removed from the normal analysis chain. Such removal guards against the possibility that the large beam deviations required by dithering might produce systematic effects that would be hard to correct for.

To further illustrate the dithering procedure, Figure 3.3 shows the response of one BPM and one charge-normalized detector channel to the dither of a corrector magnet during a typical cycle. These data would be used to obtain one of the $\frac{\partial d_i/q}{\partial C_j}$ quantities and one of the $\frac{\partial \xi_i}{\partial C_j}$ quantities in Equation (3.11). A total of 30 such quantities (corresponding to the responses of five BPM's and one detector channel to five "optimal" dithers) are needed in order to fill the matrices contained in that equation. These plots could therefore be used along with the 28 other similar plots not shown to arrive at a complete set of detector slopes $\partial A_i/\partial \xi$, using Equations (3.12) and (3.13). These would only be the slopes for one channel, though. In order to get the slopes for all detector channels, this analysis would of course need to be repeated for the other 49 Møller channels, the 10 *ep* channels, and the 16 luminosity monitor channels.

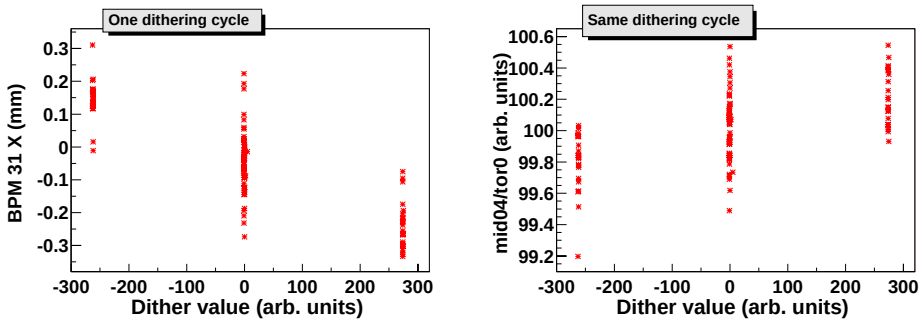

(a) (b)

Figure 3.3: The response of a BPM and a detector channel to a corrector magnet dithering during one cycle. BPM and detector signals are plotted versus a bit value representing the magnet's field strength. The detector channel (mid04) has been normalized to charge as well as to its average signal size. Dithering can be seen to affect it at the sub-1% level.

3.2.4 Beam Fluctuations Summary

Correcting detector channel asymmetries A_i consists of three basic parts, as can be seen from the plots shown in Figure 3.4. Figure (a) shows the raw asymmetry distribution for the entire Møller detector for one run (corresponding to roughly one hour's worth of data, or 200 kpairs at 120 Hz).[2] The data is very clean, so that the distribution appears as a Gaussian, but the distribution's width is large, nearly 3000 ppm. Recall from Section 2.7.1 that the entire Møller detector receives approximately 20 million electrons every pulse. Thus the width of its asymmetry distribution should be roughly 200 ppm, containing a ∼170 ppm contribution from statistical fluctuations and a ∼110 ppm contribution from the electronic noise in the ADC's.

The first step in correcting for the effects of beam fluctuations is to normalize all of the detector channels' signals to charge. This removes the ∼3000 ppm fluctuations due to intensity jitter (coming mainly from jitter in the source laser beam). Figure 3.4(b) shows the Møller detector asymmetry distribution after charge normalization. Its width has been greatly reduced to on the order of 300 ppm. This includes a ∼200 ppb contribution coming from energy jitter, a ∼100 ppb contribution coming from position jitter, and a ∼200 ppb contribution coming from statistical fluctuations and ADC noise.

The second step in correcting for beam fluctuations is to find the detector slopes $\vec{\alpha}_i$, by using either linear regression or beam dithering. Once these have been found, the third step is to subtract out on a pair-by-pair basis the false asymmetries $\alpha_{iE} \cdot \Delta E, \alpha_{ix} \cdot \Delta x, \ldots, \alpha_{iy'} \cdot$

[2] The precise way in which data from different channels are combined will be discussed in Section 3.3.

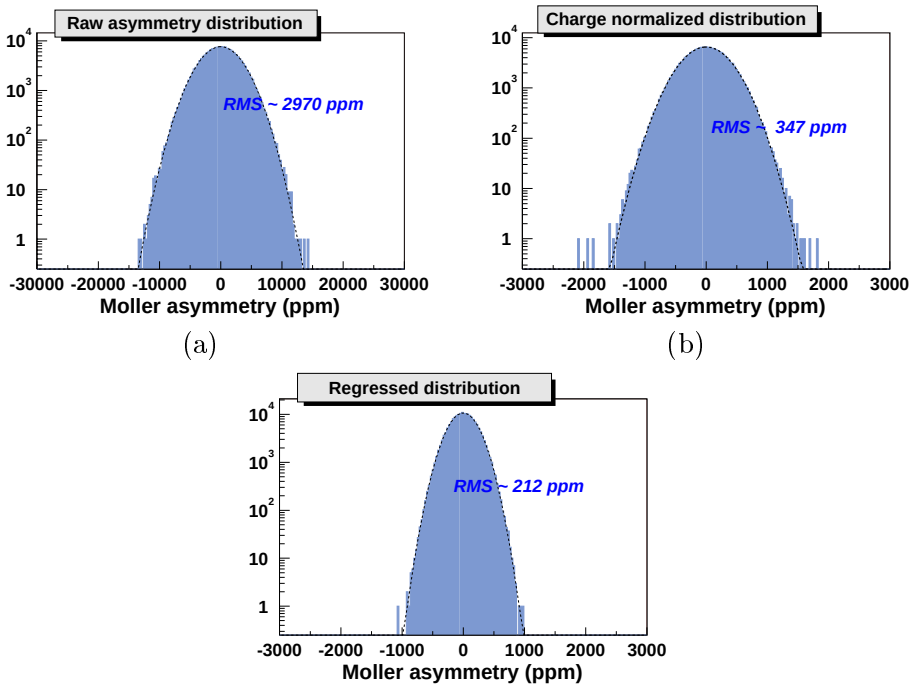

Figure 3.4: Example of Møller detector asymmetry correction procedure. Here the Møller detector consists of just the "in" and "mid" rings. Normalizing the detector's signals to charge and regressing against beam parameter differences drastically reduces the width of the asymmetry distribution.

$\Delta y'$ produced by helicity correlations in the beam parameters $E, x, \ldots, y'$. Figure 3.4(c) shows the Møller detector asymmetry distribution after false asymmetries due to beam helicity correlations have been subtracted out. The few slightly non-Gaussian tails evident in Figure (b) have disappeared, and the width is now just over 200 ppm.

Normalizing the detector signals to charge (equivalent to subtracting the charge asymmetry A_Q) and subtracting out false asymmetries on a pair-by-pair basis therefore not only shifts the mean of the detector asymmetry distribution (to the extent that the beam asymmetries are non-zero), but drastically reduces its width. There are two main caveats, however. First, the correction procedure only works for first-order beam effects. Higher order effects are discussed in Section 3.4. Second, this procedure obviously cannot correct for the electronic noise in the ADC's. Nevertheless, the procedure works well enough that, after approximately 200 million pairs, a statistical uncertainty of 200 ppm$/\sqrt{2 \times 10^8} \approx 14$ ppb should be achievable.

3.3 Data Weighting

Once the detector slopes $\vec{\alpha}_i$ are known, one can correct for false asymmetries due to helicity correlations in the beam parameters, in the following way:

$$A_i \equiv A_i^{meas} - \sum_{\xi} \alpha_{\xi i} \Delta\xi \tag{3.14}$$

Here A_i^{meas} are the raw, charge-normalized detector asymmetries, and the sum runs over all beam parameters except charge. This gives a corrected asymmetry A_i for every channel (note the change in nomenclature from that of Section 3.2). All channels must be averaged in order to form one grand asymmetry for the whole detector. Channel asymmetries are averaged with so-called channel weights:

$$A = \frac{\sum_{i}^{N} (A_i \cdot w_i)}{\sum_{i}^{N} w_i} \tag{3.15}$$

where w_i is the weight given to channel i. The choice now becomes how to weight the various channels? The simplest prescription is to give all channels equal weight, in which case the formula above becomes $A = \frac{1}{N} \sum A_i$. This is *not* the statistical average, however, since not every detector channel necessarily has the same statistical weight. As mentioned in the description of the Møller detector in Section 2.7.1, each channel is sensitive to a different azimuthal and radial range of the electron flux. Both because the signal flux is not absolutely uniform over the face of the detector, and because beam motion can cause the flux to be off from the center of the detector, different channels will see more or less scattered electrons per pulse. Channels with greater statistics need to be given more weight, and those with fewer statistics less weight. Since the detector is not precisely gain–matched, signal height cannot be taken as a strict indicator of a channel's statistical weight. Hence, a first–pass analysis is performed in order to obtain the width of each channel's asymmetry distribution after beam fluctuations have been removed via Equation (3.14). This width is inversely proportional to $\sqrt{N_e}$, where N_e is the number of signal electrons for that particular

channel. The channel's statistical weight can therefore be set equal to:

$$w_i = \frac{1}{\sigma_i^2} \tag{3.16}$$

where σ_i is the RMS of the distribution of A_i. Using the statistical weights in Equation (3.15) does a much better job of giving a true sense of the asymmetry in the underlying physics than the simple $\frac{1}{N}\sum A_i$ formula.

It is important to recognize that using improper (or, perhaps, "non–optimal") channel weights does not change the *asymptotic mean* of the final asymmetry, but only enlarges the statistical uncertainty. That is, no matter what weights one chooses to use (as long as their sum is non-zero, so that Equation (3.15) can be used), given an infinite amount of running time, the same physics asymmetry will be recovered. In other words, all weighting schemes converge towards the same final answer, just at different rates. Given a finite amount of running time, however, it can be very important to pick the weighting scheme that converges with the fastest rate (i.e., the one that produces the smallest final statistical uncertainty). Also, given a finite amount of running time, different weighting schemes *will* yield slightly different means.

The statistical weights given by Equation (3.16) are only correct if electronic noise contributions to all channels' widths are neglected. To recover the minimum statistical uncertainty, the channel weights must reflect the true statistics in each channel as closely as possible. Electronic noise can vary from channel to channel and can contain correlations between channels. By definition, the optimal set of channel weights w_i yields the smallest asymmetry widths allowed by counting statistics, and therefore minimizes the following covariance matrix:

$$\begin{gathered} M_{ij} = \langle A_i \cdot A_j \rangle - \langle A_i \rangle \cdot \langle A_j \rangle \\ \min \left| \sum_{i,j} M_{ij}\, w_i\, w_j - \left(\sum_i w_i - 1 \right)^2 \right| \Longrightarrow \text{ optimal set of } w_i \end{gathered} \tag{3.17}$$

where the brackets denote averaging over all pairs. The $\left(\sum w_i - 1\right)^2$ term constrains the sum of the weights to be close to unity. This minimization procedure is carried out for every run separately, so that each run has its own set of channel weights.

To appreciate the difference a particular set of weights can make, it is instructive to look

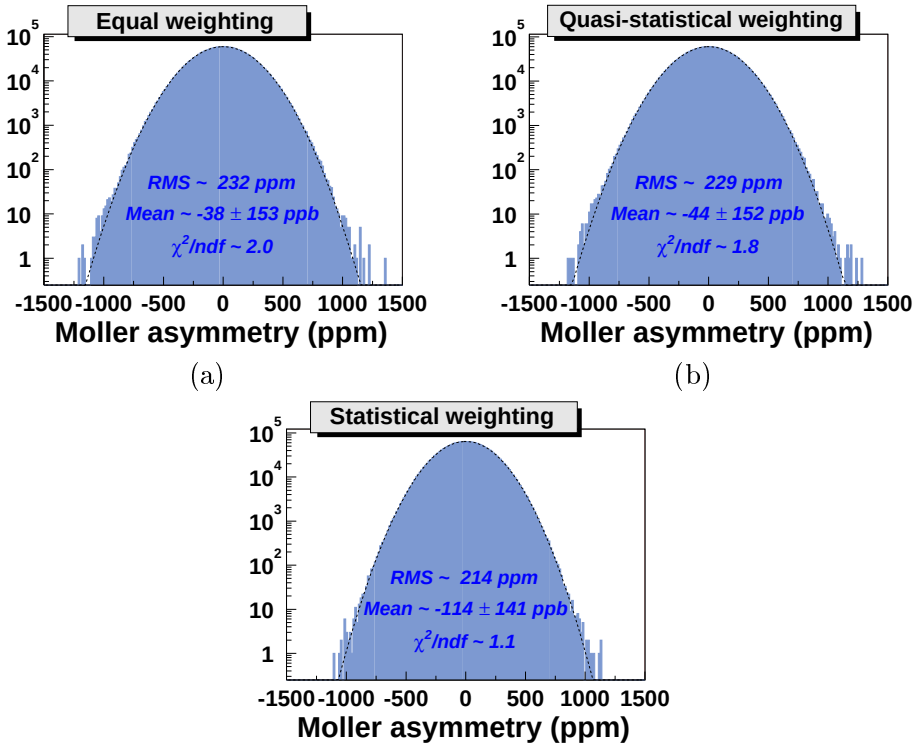

Figure 3.5: Effects of channel weighting on Møller asymmetry. Figure (a) uses the equal $1/N$ weights. Figure (b) uses the quasi-statistical $1/\sigma_i^2$ weights. Figure (c) uses the proper statistical weights given by Equation (3.17). The latter weights give the smallest statistical error as well as the most Gaussian profile, signified by the smallest χ^2.

at a sequence of plots as contained in Figure 3.5. Each plot shows the Møller detector asymmetry distribution for the same data set (roughly 2 million pairs) but uses different channel weights to generate the asymmetries. Figure (a) uses the simple $A = \frac{1}{N}\sum A_i$ (i.e., equal weighting) formula. Figure (b) uses the quasi-statistical weights given by Equation (3.16). Figure (c) uses the true statisical weights expressed by Equation (3.17), which take electronics noise contributions into account. By using the correct statistical weights, the RMS of the distribution has been reduced nearly 8%, from 232 ppm to 214 ppm. Accordingly, the final statistical error bar is smaller. In addition, the profile has become much more Gaussian, signified by the fact that the χ^2 per degree of freedom has been reduced from 2.0 to 1.1. Only by properly taking into account the various correlations between channels due to common mode electronics noise can the non-Gaussian tails be largely removed. The mean of the distribution has also changed somewhat, but this is to be expected. Different channel weights will tend to move the final answer around within the limits set by the statistics, but

as discussed before, as more data accumulates these differences become smaller and smaller, until eventually all weighting schemes return the same final answer.

The full asymmetry analysis is therefore a multi–step procedure requiring three separate passes over all the data. On the first pass, the detector slopes $\vec{\alpha}_i$ are computed either by regression or by beam dithering. A new set of slopes is computed either every 10k pairs (for regression) or every dithering cycle (for beam dithering). On the second pass, these slopes are used to correct for beam–induced false asymmetries via Equation (3.14). The corrected channel asymmetries A_i are used to calculate the covariance matrix $\mathbf{M}$. At the end of each run, this matrix is used to generate the channel weights w_i that yield the smallest asymmetry widths. On the third and final pass, these channel weights are used in conjunction with the detector slopes to compute the asymmetry for the whole detector, as expressed by Equation (3.15).

For any given detector, this procedure yields a separate asymmetry and statistical error bar, $A \pm \delta A$, for every run. These run–by–run asymmetries are combined in order to form a final grand asymmetry A_{raw} in the following manner:

$$A_{\mathrm{raw}} = \frac{\sum_{r}^{N} (A_r \cdot w_r)}{\sum_{r}^{N} w_r} \tag{3.18}$$

Here N is the number of runs in the data set, and $w_r = (1/\delta A_r)^2$ is the relative weight for that particular run. Choosing each run's weight to be inversely proportional to the square of its asymmetry's error bar ends up being equivalent to weighting by the total number of electrons detected during each run. This, in turn, ensures the smallest statistical uncertainty on the final grand asymmetry.

This "final" grand asymmetry A_{raw} actually requires further modification so that it can be compared to the Standard Model prediction for the parity violating cross section asymmetry A_{PV} in Møller scattering. Namely, normalizations (such as beam polarization) and physics backgrounds need to be taken into account. The corrections stemming from these sources are discussed in Section 4.8.

3.4 Treatment of Systematic Uncertainties

The previous sections described the methods by which beam–induced false asymmetries are subtracted from the data and the grand asymmetry $A_{\rm raw}$ is obtained. The correction procedure removes noise caused by beam fluctuations and false asymmetries caused by helicity correlations in the beam parameters, while the particular weighting scheme is chosen to yield the smallest statistical uncertainty possible. There remains the possibility, however, that the correction procedure fails to account for first-order beam effects properly, or that higher-order beam effects not accounted for by the correction procedure are in fact important. The first possibility reduces to the detector slopes obtained by regression and by beam dithering being somehow wrong, whereas the second possibility reduces to the detector slopes constituting an incomplete set. Both possibilities need to be addressed, and by doing so a systematic uncertainty can be ascribed to the asymmetry correction procedure. This section will not give a final number for this systematic uncertainty (that will actually be given along with the other results in Chapter 4). Instead, it will present the methods by which that uncertainty can be ascertained.

3.4.1 First-Order Asymmetry Systematics

That the detector slopes $\vec{\alpha}_i$ are not grossly wrong is attested to by the fact that they consistently succeed in reducing the width of the Møller asymmetry distribution from typically 300 ppm and above to roughly 200 ppm. If the slopes were wrong, the correction procedure would not subtract out noise in a coherent manner, meaning that the asymmetry distribution widths would at best stay the same, and at worst grow bigger, potentially a good deal bigger. This, then, is the first qualitative assurance one has that the correction procedure is handling the first order systematics correctly. In order to set a quantitative limit on how well the first-order correction procedure is working, however, it is necessary to look at different linear combinations of detector channels, some of which possess heightened sensitivity to fluctuations in specific beam parameters.

Each ring of the main integrating calorimeter consists of a given number of channels, each of which gets assigned a different weight. Up until now, the weights have been picked to reflect the statistics collected by each channel. These weights can be referred to as "monopole" weights, since the sum of the weights in a given ring is constrained to be nearly

unity. Using them therefore returns what is basically the "monopole moment" (i.e., the statistical average) of the detector's asymmetry data.

In contrast, there are also "dipole" weights, which return the "dipole moments" of the detector's asymmetry data, often simply referred to as dipole asymmetries. Dipole asymmetries are defined in the following manner:

$$\text{X dipole:} \quad A_x = \frac{2}{N_{ring}} \sum_i A_i \sin\left(\frac{2\pi\,(i - i_{top})}{N_{ring}}\right) \tag{3.19}$$

$$\text{Y dipole:} \quad A_y = \frac{2}{N_{ring}} \sum_i A_i \cos\left(\frac{2\pi\,(i - i_{top})}{N_{ring}}\right) \tag{3.20}$$

Here A_i is the corrected, charge-normalized asymmetry for channel i, and i_{top} is defined to be the channel number of the "top" of the ring. The dipole asymmetries are calculated for each ring separately. Note that finding the dipole asymmetries basically just boils down to using a different set of weights in Equation (3.15), ones that vary sinusoidally around the azimuth of each ring. However, the sum of the dipole weights is zero, so one cannot divide by their sum in Equation (3.15). Instead, one simply divides by N_{ring}, the number of channels in the ring.

To better understand what the dipole asymmetries really mean, consider a beam perfectly centered on both the target and the detector. Imagine a beam with no helicity-correlated differences in any beam parameter, i.e., $\Delta E = 0$, $\Delta x = 0$, ..., $\Delta y' = 0$. Under these conditions, the signal heights in all of the detector channels in a given ring will tend towards the same value for both beam helicity states. Consequently, the asymmetry in each detector channel will tend towards the same value, that given by the physics. Equations (3.19) and (3.20) will thus return zero. This makes sense, because in this case there is

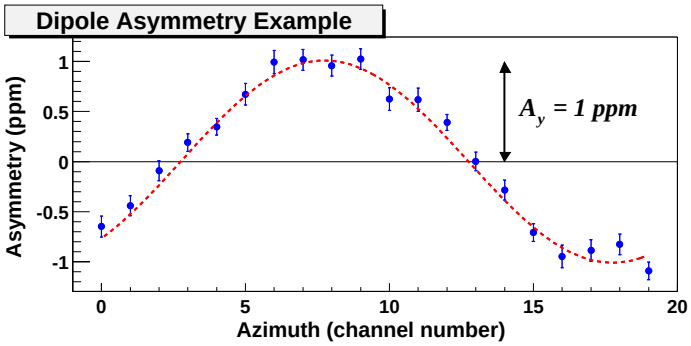

Figure 3.6: Example of a vertical dipole asymmetry.

perfect symmetry in the detector's asymmetry data. Now imagine that for the left helicity state, the beam obtains a small vertical displacement, i.e., $\Delta y \neq 0$. The signal sizes in some of the channels closest to the left-helicity beam will increase, while those directly opposite will decrease in a similar fashion. This will produce an azimuthally varying asymmetry, as depicted in Figure 3.6. Using the monopole weights will tend to cancel out these effects, as larger asymmetries on one side of the detector are compensated for by smaller asymmetries on the other side. Using the dipole weights, however, will tend to enhance these effects, uncovering the amplitudes A_x and A_y of any azimuthal variations in the detectors' asymmetries. As will be seen below, this can be useful when studying the systematic uncertainty associated with the asymmetry correction procedure.

Dipole asymmetries may only become statistically significant on certain timescales. On short timescales, beam jitter may completely swamp whatever beam systematic is causing the dipole asymmetry. On long timescales, however, things may tend to average out, thanks to the periodic systematics reversals performed at the source (see Section 2.2.8). On just the right timescales (long enough to average out jitter and reveal the systematic effect, but short enough that no systematics reversals are averaged over), the dipole asymmetries may become apparent as statistical outliers.

The most relevant timescale for studying such systematics is called the "slug." A "slug" is defined as a period of time for which the energy of the beam, the state of the source half-wave plate, and the state of the asymmetry inverter all remain unchanged. In other words, a slug is a period of data taken under nominally static conditions. Typically, a slug lasts only two days, at which point either the energy, the half-wave plate, or the asymmetry inverter is changed. For illustrative purposes only, a slug plot of the "out" ring's horizontal dipole asymmetry A_x is given in Figure 3.7. There are many statistical outliers readily

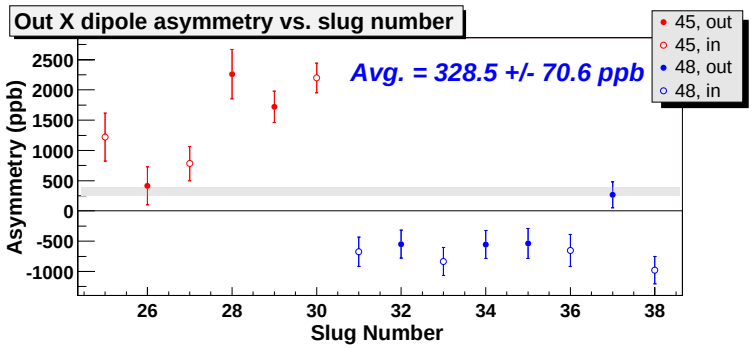

Figure 3.7: Slug plot of the X dipole asymmetry for the "out" ring, for Run II only.

apparent, and this is reflected in the very poor χ^2 per degree of freedom of the data. Note that this figure shows data that has been corrected for beam effects. As such, one would naively expect there to be no statistically significant asymmetry dipoles, since the correction procedure should have effectively removed all helicity-correlated systematics *on a channel-by-channel basis.* What, then, could be responsible for the statistical outliers?

As it turns out, dipole asymmetries can arise via two different mechanisms:

1. *Transverse beam polarization:* As mentioned in Section 2.2.8, the spins of the electrons precess as the beam passes through the bend in the BSY. If the energy is not precisely set to one of the two values that preserves longitudinal polarization (either 45 or 48 GeV), then the beam will acquire a small horizontal polarization component. In addition, imperfect beam steering through quadrupole magnets in the A-line can result in small horizontal and vertical polarization components. Transverse polarization components will create dipole asymmetries through two photon exchange processes, which are described further in Section 4.8.5. These processes couple horizontal beam polarization to a vertical dipole asymmetry, and vertical beam polarization to a horizontal dipole asymmetry.

2. *Improper beam corrections:* If the asymmetry correction procedure is somehow flawed (e.g., if the detector slopes are wrong due to a BPM malfunction), then it will effectively leave some of the beam helicity correlations unaccounted for in the data. These residual helicity correlations can then produce contributions both to the overall monopole asymmetry as well as to the dipole asymmetries.

The contributions stemming from the first mechanism can be estimated from specific measurements in a straightforward manner, as will be discussed in Section 4.8. Estimating the extent to which the second mechanism contributes, that is, the uncertainty on the first-order beam corrections, is a bit trickier. In the end, it relies on a useful (if lucky) feature of the data, namely, the fact that the data can be neatly divided into two timeslots, each of which happens to have very different beam systematics.

Regardless of the sizes of the beam systematics which they contain, two independent subsets of the same overall data set should produce statistically consistent final answers, if the asymmetry correction procedure works properly. The extent to which the two subsets' answers disagree, in fact, can be taken as a measure of how well the procedure works. The

larger the differences in the corrections between the two subsets, and the better the final agreement, the more the correction procedure can be trusted. The systematic error of a particular asymmetry correction ΔA is therefore equal to the following:

$$1^{\text{st}} \text{ order error } \delta(\Delta A) = \left| \frac{A_1 - A_2}{\Delta A_1 - \Delta A_2} \right| \cdot \Delta A \tag{3.21}$$

where A_1 and A_2 are the corrected asymmetries of data subsets 1 and 2, respectively, for a particular combination of detector channels that are especially sensitive to the correction being tested, and ΔA_1 and ΔA_2 are the sizes of the corrections themselves. For instance, it has already been discussed how the horizontal dipole patterns have heightened sensitivity to helicity correlations in x. Hence, one ring's X dipole will be used to place a limit on the systematic uncertainty of the Δx asymmetry correction. Likewise, one ring's Y dipole will be used to place a limit on the systematic uncertainty of the Δy asymmetry correction. This method of estimating the uncertainties attributable to each of the beam corrections will be fully exploited in Section 4.4.

3.4.2 Higher-Order Asymmetry Systematics

The method described immediately above only addresses the uncertainties associated with the first-order beam correction procedure. The two subsets into which the whole data set is divided turn out to have very large first order corrections (i.e., terms directly proportional to beam helicity correlations via the detector slopes $\vec{\alpha}_i$), which completely obscure any higher order effects there might be. Comparing the final asymmetries for these two data subsets can therefore place a limit on how well the first-order correction procedure works, but it cannot yield information as to the significance of higher order systematics. Higher order systematics could still be responsible for more subtle effects that do not necessarily scale with those of the first-order beam corrections. It is ostensibly these subtle effects, in fact, that are responsible for the statistical outliers in the data presented in Figure 3.7.

First of all, the meaning of "higher-order" asymmetry systematics should be clarified. One can begin by furthering the Taylor expansion of Equation (3.3), including higher order

terms as one goes along:

$$\begin{aligned} A^{\mathcal{O}(0)} &= A^{phys} \\ A^{\mathcal{O}(1)} &= \sum_{\xi} \frac{\partial A}{\partial \xi} \Delta\xi \\ A^{\mathcal{O}(2)} &= \sum_{\xi} \sum_{\zeta} \frac{\partial^2 A}{\partial \xi \partial \zeta} \Delta\xi \Delta\zeta \\ &\vdots \\ A^{meas} &= A^{\mathcal{O}(0)} + A^{\mathcal{O}(1)} + A^{\mathcal{O}(2)} + \cdots \end{aligned} \tag{3.22}$$

The sums always proceed over the beam parameters E, x, y, x', and y'. The $\mathcal{O}(2)$ terms included here can be shown to be small compared to the $\mathcal{O}(1)$ terms, since the former are suppressed by an additional factor of a beam helicity correlation $\Delta\xi$ (typically very small), and the second-order partial derivatives of the measured asymmetry are no larger than the first-order ones [100]. However, the formalism in Equation (3.22) does not include all of the relevant sources of higher order systematics.

One obvious higher order systematic that Equation (3.22) does not include, and which has already been discussed in Section 2.4.4, is the spot size induced asymmetry. This effect's contribution to the final experimental asymmetry will be calculated in Section 4.4. Perhaps the most significant higher order effects not contained in Equation (3.22), however, are intra-pulse beam effects. As will be seen below, systematics from such effects basically stem from the fact that no device in the experiment (no BPM, no toroid, and no detector channel) is sampled in discrete steps over the length of the beam pulse, with a flash ADC for example. The experiment uses integrating ADC's, so that only averages are measured. Consequently, it is possible that the beam parameters could be changing within a pulse in such a way that the computed first-order corrections do not correctly account for the actual false asymmetries observed by the detector.

An example of an intra-pulse beam effect is shown in Figure 3.8. This example shows all of the necessary ingredients to generate a systematic error in the asymmetry correction procedure. First, the beam needs to have a time dependence in at least one of its parameters ζ. In this case, the beam energy is changing over the course of the pulse. Second, the helicity correlation in another beam parameter ξ, not necessarily the same as ζ, must also have a time dependence. In this case, ΔE grows from 20 keV to 100 keV over the course of the

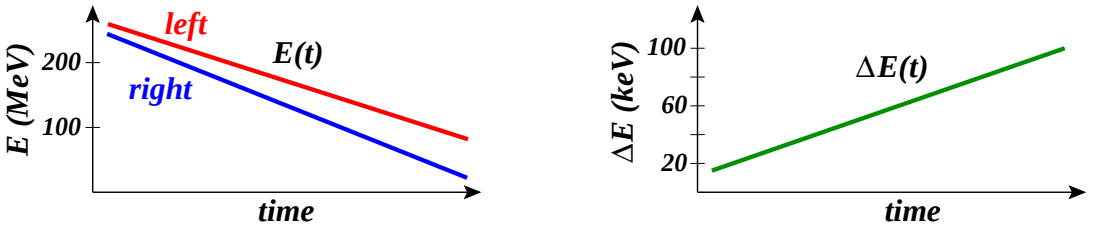

Figure 3.8: Example of intra-pulse beam effects. The detector slope α_E is assumed to be a function of energy. Since the pulse energy varies as a function of time, the slope $\alpha_E(t)$ effectively becomes a function of time. In this case, the false asymmetry observed by the detector, $\langle \alpha_E(t) \cdot \Delta E(t) \rangle$, would not equal the correction computed, $\langle \alpha_E(t) \rangle \cdot \langle \Delta E(t) \rangle$.

pulse. Third and lastly, the detector slope α_ξ for beam parameter ξ must itself vary as a function of ζ. The reason that these conditions create an error in the computed asymmetry correction can be easily understood by noting that the actual false asymmetry observed in the detector is equal to the following:

$$\langle \alpha_\xi \cdot \Delta\xi \rangle \equiv \frac{1}{\tau} \int_0^\tau \alpha_\xi \left(\zeta(t) \right) \cdot \Delta\xi(t)\, dt \tag{3.23}$$

where $\zeta(t)$ is the value of beam parameter ζ evaluated at time t, $\Delta\xi(t)$ is the helicity correlation in beam parameter ξ evaluated at time t, and τ is the total pulse length. Because α_ξ varies as a function of ζ, and ζ changes as a function of time, the detector slope α_ξ effectively becomes a function of time, a fact reflected in the above equation when writing $\alpha_\xi \left(\zeta(t) \right)$. Since the experiment only measures the averages of the detector slope and of the beam helicity correlation over the course of the pulse, an error in the computed correction will result, owing to the following simple fact:

$$\langle \alpha_\xi \cdot \Delta\xi \rangle \neq \langle \alpha_\xi \rangle \cdot \langle \Delta\xi \rangle \tag{3.24}$$

The detector actually sees $\langle \alpha_\xi \cdot \Delta\xi \rangle$, whereas the experiment only measures $\langle \alpha_\xi \rangle$ and $\langle \Delta\xi \rangle$, so that one can only use these quantities when computing a correction. Given a scenario similar to the simple one pictured in Figure 3.8, where both the beam parameter ζ and the helicity correlation $\Delta\xi$ vary linearly over the pulse, and the detector slope α_ξ is assumed to vary linearly as a function of ζ, the error in the correction procedure stemming from intra-pulse beam effects will be exactly equal to the following:

$$\text{higher order error} \;=\; \frac{1}{12} \cdot \frac{\partial \alpha_\xi}{\partial \zeta} \cdot (\zeta_{head} - \zeta_{tail}) \cdot (\Delta\xi_{head} - \Delta\xi_{tail}) \tag{3.25}$$

where ζ_{head} and ζ_{tail} are the values of beam parameter ζ at the head and tail of the beam pulse, respectively, and similarly $\Delta\xi_{head}$ and $\Delta\xi_{tail}$ are the helicity-correlated differences in parameter ξ at the head and tail of the pulse. In fact, even if a time dependence much more complicated than that pictured in Figure 3.8 is assumed, one can proceed by dividing the beam pulse up into an arbitrary number of roughly linear segments, and then using Equation (3.25) on each of these segments individually, obtaining a separate error for each segment. These errors can give a feel for the rough size of the effects capable of being produced, as well as actually be added up (either linearly or in quadrature, depending on how conservative one wants to be) to get an overall systematic uncertainty estimate for these important "higher order" effects.

Comparing Equation (3.25) to Equation (3.22), one sees that Equation (3.25) is indeed "higher order" in the sense that it contains second-order partial derivatives of the measured asymmetry $\partial^2 A/\partial\xi\partial\zeta$. However, it only contains terms first-order in $\Delta\xi$, so it is not higher order in beam asymmetries. Thus, depending on the exact shape of the beam and on the functional form of α_ξ, these error terms could potentially be as significant as the $\mathcal{O}(1)$ terms in Equation (3.22).

As will be seen in Section 4.4, the detector slopes and beam shape are such that the higher-order error terms of Equation (3.25), while non-negligible, do not appear to be larger than 10 ppb. The evidence for this comes from two sources. First, a thorough characterization of the beam was performed, from which beam shape data could be gathered [101]. A detailed Monte Carlo simulation of the target and spectrometer system was also used to ascertain the detector slopes' dependences on each of the beam parameters. From these two pieces of information, it is possible to construct "worst-case scenarios" that are useful in quantifying the likely systematic error stemming from higher order effects. Second, during Run III several of the BPM signals were actually divided into four time slices, which could then be used in the regression analysis. Analysis of this "sliced" data is ongoing, but so far it has proven invaluable in testing various hypotheses and providing a fuller sense of the processes behind the behavior responsible for the statistical outliers evident in Figure 3.7.

The sliced data from Run III, in fact, is largely responsible for the decision to divide the Møller detector into a Region I and II. As will be seen in the next chapter, Region II (the "out" ring) shows evidence for being particularly sensitive to beam parameters that exhibit strong intra-pulse variation, whereas Region I does not show corresponding behavior.

Analysis of the sliced data supports the conclusion that it would be better to drop Region II from the actual physics analysis, treating it instead as a sensitive monitor for beam effects (much like the luminosity monitor). In doing so, one loses some statistical power, but this loss is more than made up for by the accompanying decrease in the systematic uncertainty due to higher-order beam effects.

Perhaps the most important observation to make regarding higher-order asymmetry systematics is that such terms will undoubtedly be correlated in some manner to charge. This is simply because all aspects of the beam, even higher order moments and shapes, are to some extent affected by the overall beam intensity. For instance, wakefield effects, which are the mechanism responsible for inducing so-called "banana shapes" in the beam, scale with beam intensity [81, 102, 103]. Therefore in addition to normalizing the detector data to beam intensity (equivalent to subtracting out the first-order dependence on the charge asymmetry), the full asymmetry correction procedure should include regression against charge. This regression slope α_Q will absorb at least part of the detector's dependence on all higher-order beam systematics, including intra-pulse effects, as well as account for imperfections in the normalization procedure due to toroid and detector nonlinearities. The summation over the beam parameters in all equations in this chapter analogous to Equation (3.14) should therefore include an extra term $\alpha_Q \cdot A_Q$.

The "extra" regression against charge can be thought of as an acknowledgement that the precise dependence of the measured asymmetry on the charge asymmetry is a bit more complicated then the standard $A^{meas} = A^{phys} + A_Q$ formula would suggest. In reality, the deviation from this simple formula should be small, corresponding to charge slopes α_Q on the order of 1% or less. Indeed, the size of the charge slopes can be taken as an indication of the relative importance of higher-order beam effects. This approach is more fully explored in Section 4.4. This, then, offers a third way in which an upper limit can be placed on the contributions of higher-order asymmetry systematics to the final experimental error bar.

Estimation of systematic uncertainties is always a difficult task with a somewhat questionable outcome. Estimating the systematic uncertainty in the asymmetry correction procedure is especially difficult, since so many higher-order beam parameters (mostly relating to beam shape) went essentially unmeasured during Runs I and II. Nevertheless, this section has motivated three independent methods by which an upper limit can be placed on the effects of higher-order beam systematics. First, the sizes of the effects can be estimated

using Equation (3.25). Second, analysis of the "sliced" data in Run III can be used to show the extent to which failure to regress against beam shape information affects the final answer. Third and lastly, the charge slopes can be used to indicate the relative importance of higher order effects. While none of these methods is absolutely foolproof by itself, and all are to some extent limited by statistics, together they provide reasonably convincing evidence for an upper limit on the beam-related systematic error in Runs I and II.

3.5 Asymmetry Blinding

In order to avoid bias, particularly when applying cuts to the data, the Møller detector asymmetries are "blinded" by a random offset whose value remains hidden to all experimenters. This offset is chosen by a pseudo–random algorithm having a flat probability distribution between a given set of limits. The limits are chosen to be the same order of magnitude as the Standard Model prediction for the size of the asymmetry. Thus the random offset can be any number between -200 ppb and $+200$ ppb.

Separate blinding offsets are chosen for Runs I and II, in order that each result can be unblinded independent of the other. The blinding offset is chosen at the beginning of each data–taking period and kept encrypted until the time of the unblinding. The same offset is applied to every Møller detector channel asymmetry for every pair, and is applied during the "prompt" processing of the data. Prompt processing generates processed pair–wise data structures from raw event–wise data structures. The raw event-wise data structures only contain raw signal data from all of the ADC's (plus other devices such as VSAM's and I/O modules), and occur for every beam pulse. The processed pair-wise data structures, however, contain asymmetry information for all of the signals, and so only occur for every pulse pair. The raw data structures are useful for quickly looking at the gross characteristics of the detector and beam (signal heights, signal noise, beam jitter, beam position, etc.), but are not used for much more than that. The real detector analysis, involving corrections for beam-related false asymmetries and averaging over many channels and pairs, uses the processed data structures.

The Møller detector asymmetries remain blinded until the analysis of the data has been finalized. That is, all cuts are decided upon while the data is blinded. This avoids taking the effect a particular cut has on the final answer into account when evaluating the cut's

worthiness. As will be seen in Section 4.3, however, there are relatively few cuts, and all of them are designed to merely guard against hardware failures of one type or another. Nevertheless, a "blind" analysis helps provide an additional layer of protection against the possibility of experimental bias.

It should be emphasized that blinding only applies to the Møller detector. For instance, beam asymmetries, *ep* detector asymmetries, luminosity monitor asymmetries, and pion detector asymmetries all remain unblinded. One reason for this is that all these quantities only become interesting as they relate to the asymmetry in the Møller detector. That is, most of these quantities should be zero (or nearly zero) within the statistical precision of the measurements, and any departure from zero does not signal "new physics," but rather some systematic problem that needs to be identified and addressed. For example, a large asymmetry in the luminosity monitor, after correcting for beam effects, could signal some systematic problem with the regression and/or beam dithering procedure. Likewise, the pion detector asymmetry is only useful as a background measurement that then needs to be subtracted from the Møller result. The *ep* asymmetry, on the other hand, is predicted to be large (on the order of 1 ppm, although more precise predictions are hard to make) and does in principle contain interesting physics. However, the *ep* detector was not designed to make a clean measurement of any one physical parameter. As with the pion detector, it is useful primarily as a background measurement. Hence, it does not make sense to blind these quantities, since their knowledge does not threaten to bias the main physics measurement.

Chapter 4

Experimental Results

4.1 Data Organization

This chapter presents the "unblinded" results from the first two periods of physics data collection, referred to as Runs I and II, for SLAC experiment E-158. The data can be organized in a variety of ways, depending on the timescale one wants to probe. On the shortest timescales, the data are divided into pulse pairs of opposite helicity. Depending on the beam rate, pulses within a pair are separated by either 17, 34, or (very rarely) 67 ms. On somewhat longer timescales, the data are grouped into individual runs, the length of a run being limited primarily by the size of the raw event data structure. In order to keep the size of each run's raw data file on disk on the order of a few hundred megabytes, a run typically lasts one hour. The longest timescale, and one that is particularly relevant when studying systematics, is the slug, which is the longest grouping that can be made of consecutive runs taken under nominally identical conditions.

Typically a slug consists of two days' worth of data, equal to several million pairs. Slugs can be combined to study systematics even further. For instance, all of the 45 GeV slugs taken with the source half-wave plate inserted into the source laser's path, and with the asymmetry inverter in the +2I position, can be grouped together to study the experiment's sensitivity to this particular configuration. The first run (3690) used for the analysis of Run I data started on 12 March 2002. The last run (5606) ended just after midnight on 28 May 2002. This period consisted of 24 slugs. The first run (6695) used for the analysis of Run II data began on 10 October 2002, while the last run (7855) ended on 13 November 2002. There were a total 14 slugs during this period.

4.2 Polarimetry

As soon as possible after the electron beam is produced, the PITA slopes m_{CP} and m_{PS} are measured using an ASSET toroid. The voltages of the CP and PS Pockels cells are then adjusted to null the measured charge asymmetry. Because this charge asymmetry is proportional to the terms Δ_{CP} and Δ_{PS}, which themselves generate linear polarization components in the source laser beam, nulling the charge asymmetry with the CP and PS cells should correspond to maximizing the electron beam polarization. To check this, so-called polarimetry scans are conducted soon after the beam reaches the ESA. Polarimetry scans involve tweaking the CP and PS cell voltages, one at a time, about their nominal operating points, while measuring the beam's polarization using the normal polarimetry procedure described in Section 2.7.5. The data from a typical scan are shown in Figure 4.1, yielding a set of Pockels cell voltages that maximizes the beam polarization. Polarimetry scans must be conducted for each state of the half-wave plate, because the half-wave plate introduces additional phase shifts that both change the PITA slopes and produce linear polarization asymmetries at the photocathode. These then need to be compensated for using Δ_{CP} and Δ_{PS}.

Once a particular set of optimal voltages has been obtained for a given state of the half-wave plate, its values are entered into the source control program. In addition to interacting with the PMON system to control the helicity of each beam pulse, the source control program implements the helicity-correlated feedback loops described in Section 2.2.6.

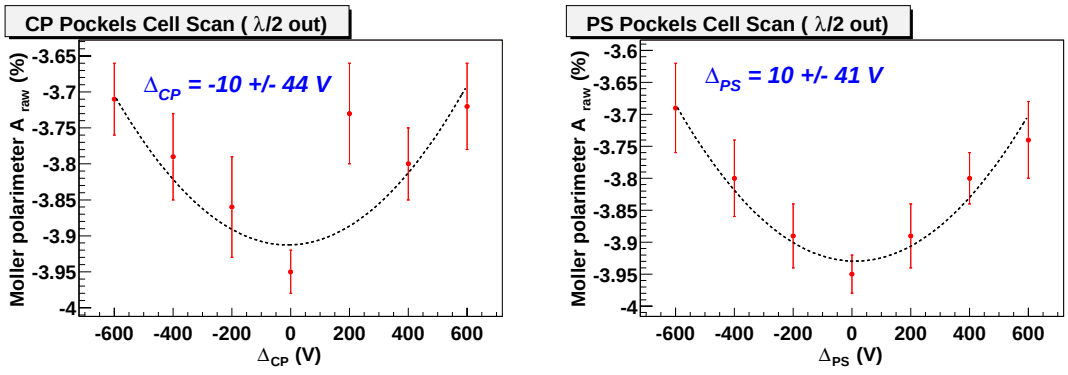

(a) (b)

Figure 4.1: Example Δ_{CP} and Δ_{PS} polarimetry scans. These data show that no offsets from the nominal Pockels cell voltage values are necessary in order to achieve maximal beam polarization. This is consistent with the fact that nulling the charge asymmetry with Δ_{CP} and Δ_{PS} simultaneously maximizes polarization.

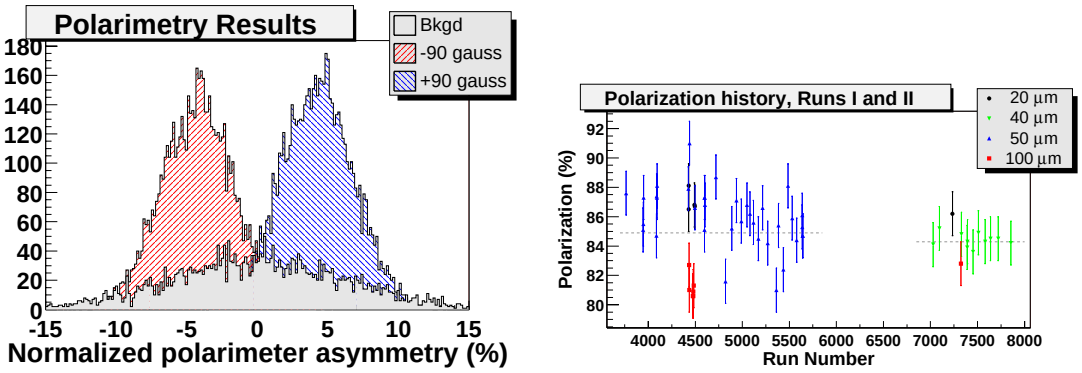

(a) (b)

Figure 4.2: Polarimetry measurements for Runs I and II. Figure (a) shows typical polarimetry measurements. The broad distribution centered about zero is a "background" run taken with no target. Figure (b) shows every polarimetry measurement for Runs I and II. The error bars are dominated by systematics, which are stable over time, and thus they have been reduced for purposes of clarity. The dotted lines show the average polarization for both runs.

In particular, the phase feedback constantly monitors the asymmetry correction the IA loop is making to null the charge asymmetry and adjusts Δ_{CP} and Δ_{PS} to keep this correction small. Typically these changes are very small (on the order of tens of volts) and fairly stable over the course of a slug. Thus, the beam polarization needs only to be monitored at the beginning and (sometimes) end of a slug.

A typical polarimetry analysis plot is shown in Figure 4.2(a). For redundancy, the polarimetry procedure is repeated with the foil-saturating magnetic field reversed, as can be seen from the two Gaussian asymmetry distributions on either side of zero. This redundancy check verifies that the observed asymmetry is a real physics effect due to polarized Møller scattering. A number of factors need to be taken into account in order to convert the raw asymmetry measured by the polarimeter into an actual measurement of the beam polarization [104]. First, a dilution factor due to the photonic background is measured by taking data with the iron foil target in and out, as well as with the spectrometer quadrupoles on and off. Second, the *ep* background (both radiative Mott scattering as well as deep inelastic scattering) needs to be subtracted out. This background is estimated by first scanning the polarimeter radially inwards and outwards, mapping the area from 15 to 40 cm. These scans are then compared to a GEANT simulation. From these studies the *ep* background in the region of the polarimeter is determined to be $8 \pm 3\%$ [104].

After background subtraction, the beam polarization can be obtained using Equation (2.21), assuming the target polarization and scattering kinematics are known. The

beam polarization is found to be very stable throughout Runs I and II, as can be seen from Figure 4.2(b). The plot in this figure shows all of the polarimetry measurements for boths Runs I and II. The error on each data point is completely dominated by systematics. The statistical error is negligible, since a 10 minute polarimetry measurement can easily reduce the statistical error to less than $\pm 1\%$. The major sources of systematic uncertainty are the 3% uncertainty in the background subtraction procedure, as well as the 3% uncertainty in the measurement of the absolute target polarization. Several other effects contribute each at the 1% level.

Most of the polarimetry measurements were made using either the 40 or 50 μm iron foil, but the 20 and 100 μm foils were also sometimes used. While the polarization values obtained using the 20, 40, and 50 μm foils all agree with one another, the value obtained using the 100 μm foil significantly disagrees. This is due to the Helmholtz magnetic field ($\sim$90 gauss) being too low to fully saturate the relatively thick foil. For Run III, a more powerful supply was installed for the Helmholtz coils, which were then capable of generating stronger magnetic fields that could fully saturate even the 100 μm foil. Polarization values obtained using all four foils now agree. Nevertheless, all polarimetry measurements for both Runs I and II are given equal weight, and the observed spread due to target thickness is taken as an additional contribution to the total systematic error. The final results are therefore:

$$\begin{array}{lll} & \underline{\textit{Run I}} & \underline{\textit{Run II}} \\ P_B = & 84.9 \pm 4.6\,(\text{syst})\ \% & 84.3 \pm 4.6\,(\text{syst})\ \% \end{array} \tag{4.1}$$

There is no reason to assume run-to-run (or even slug-to-slug) polarization fluctuations beyond the limits set by the systematics-dominated error bars. Therefore, these polarizations are taken to be effectively constant for both runs.

4.3 Data Cuts

Data cuts ensure that the data used for physics analysis is of optimal quality. In order to guard against potential bias, however, no helicity-related quantities are ever involved in any of the cuts. All cuts are merely designed to guard against hardware failures of one type or another. There are relatively few cuts, so each will be described below. Along with its description, the *approximate* data loss rate resulting from each cut, averaged over both

Runs I and II, will also be given (in parentheses):

`torLevel` **(2%)** Eliminates low-intensity data. Both pulses within a pair must contain at least 1×10^{11} electrons.

`diffTrigger` **(5%)** At a very small level, the gain of the toroid electronics depends on the amount of time since the previous beam pulse (the instantaneous beam rate). The variable `diffTrigger` records how much time has elapsed since the previous beam trigger. If the two pulses in a pair have different `diffTrigger` values, the toroids could have had *slightly* different gains during the time the two pulses occurred. In this way relatively large charge asymmetries (a hundred times larger than the normal beam jitter) could be erroneously recorded. It is important to eliminate these spurious asymmetry measurements by requiring both pulses within a pair to have the same `diffTrigger` values.

`timeslot` **(2%)** The `timeslot` variable records the phase of the beam pulse with respect to the 60 Hz line noise. Normally the helicity algorithm operates such that both pairs within a pulse occur within the same timeslot. Occasionally, however, a timing glitch occurs, and a pair will have "mixed timeslot" data, which must be cut in order to eliminate the substantial 60 Hz line noise.

`outliers` **(3%)** A simple version of the `GoodBeam` cut described below. It monitors all of the BPM and toroid measurements, looking for pairs in which the two pulses' measurements differ wildly from one another (more than six times the normal jitter). It does *not* look at any helicity-correlated quantity. It also includes the `torLevel` cut.

`nodither` **(4%)** Eliminates data for which the beam was actively being dithered. This data is only used for the dithering analysis.

$\longrightarrow$ `regUsed` **(10%)** This is basically a combination of the five cuts above. Its loss rate indicates the level at which the above cuts are correlated.

`trans90` **(0.4%)** Transmission through the entire linac and into ESA is normally very good (losses are typically 1%). Beam instability can cause significant losses due to beam scraping, so that the ratio of the alcove toroid average signal level to that of the ASSET toroids falls below 90%. It is essential to eliminate these occurrences, since they could

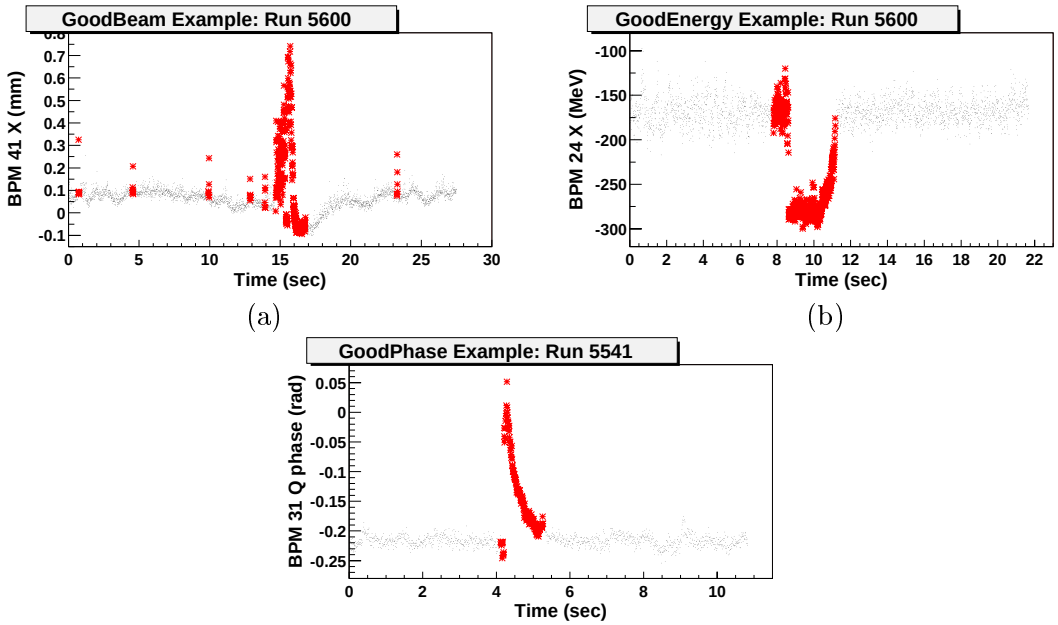

Figure 4.3: Examples of beam quality cuts, with the large red asterisks corresponding to data that is eliminated by each cut. Figure (a) shows an example of a sudden beam centroid "jump", eliminated by the `GoodBeam` cut. Figure (b) shows an example of a klystron cycle, eliminated by the `GoodEnergy` cut. Figure (c) shows an example of a sudden BPM phase spike, eliminated by the `GoodPhase` cut.

potentially cause detector asymmetry outliers that are impossible to correct for. This is a stretched cut (two pairs before and 4 pairs after the bad pair).

`toroidAgree` (0.1%) All four alcove toroids have resolutions typically in the range of 30 to 70 ppm. They should therefore show agreement roughly at this level. If two alcove toroids disagree by more than 1000 ppm, this indicates something unusual happened (i.e., a momentary hardware failure). Neither toroid measurement can be trusted. To be conservative, then, the pair is thrown out.

`GoodBeam` (6%) Includes the simple `outliers` cut above (stretched by ± 1 pair), but also implements a more sophisticated cut intended to eliminate severe beam instabilities. Regions of unusually large beam position jitter or very rapid beam centroid "jumps" are removed, as such behavior usually indicates poor beam stability and can adversely affect the data quality. The cut is stretched by ± 50 pairs for optimal performance. An example is given in Figure 4.3(a).

`GoodEnergy` (2%) Occasionally a klystron power supply fails, at which point the beam

loses roughly 200 MeV of energy. A spare klystron automatically compensates a few pulses later by turning on. Eventually, the first klystron recovers, resulting in an excess 200 MeV of energy, at which point the compensating klystron turns off, finally returning the entire system to the starting configuration. This sequence of events is commonly referred to as a "klystron cycle," and is handled automatically by the linac's energy feedback loop. Klystron cycles result in sudden beam motions, which adversely affect data quality. They are easily caught and eliminated by the `GoodEnergy` cut, with some overlap with the `GoodBeam` cut. Like the `GoodBeam` cut, the `GoodEnergy` cut is stretched by ± 50 pairs. An example of the cut is given in Figure 4.3(b).

`GoodPhase` **(3%)** By constantly monitoring the phase of each BPM's Q-cavity signal (which should be independent of beam position), a BPM feedback loop ensured that all BPM's remained properly phased. Occasionally, large phase adjustments resulted in sudden spikes to appear in the BPM's Q-cavity phases. These phase spikes can affect the BPM's X and Y measurements, inducing large noise in the readings. A typical spike signature is shown in Figure 4.3(c), which shows a fast rise time, followed by a slower decay. The cut is therefore stretched for five pairs before and 50 pairs after the bad pair.

$\longrightarrow$ `GoodAll` **(8%)** A combination of the previous three cuts. This loss rate indicates the level of correlation between the three cuts.

`rate` **(0.8%)** In addition to the `diffTrigger` cut, which basically ensures that both pulses in a pair were taken at the same instantaneous rate, there is a cut on the absolute rate. This `rate` cut only affects a small portion of the data, namely runs taken between 23 April 2002 and 27 April 2002, when the toroid damping hardware was being tested and finalized. An error in the DAQ triggering setup means that all non-120 Hz data in one of the timeslots (`timeslot==0`) should be rejected for this period. A typical example from one run (run 4724) during this period is shown in Figure 4.4.

`source` **(0.2%)** The CP and PS cell voltages are read out by an I/O module at the source. The `source` cut first checks that the Pockels cell voltage data is present in the data stream, and not missing due to intermittent data packet loss. Secondly, it checks that the CP cell's high voltage is good ($|V_{CP}| > 2000$ V), which ensures good beam

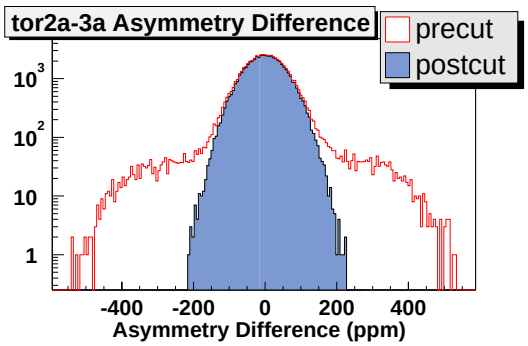

Figure 4.4: Example of `rate` cut for run 4724. The histograms show the toroid 2a-3a asymmetry agreement distributions before and after the `rate` cut is applied. The cut eliminates the non-Gaussian outliers in the agreement distribution.

polarization, essential for any asymmetry analysis. Note that the PS cell's voltage is always within a few hundred volts of zero. This cut is stretched by ±120 pairs, in order to account for a time delay in the method by which the source voltages are read out.

`daqError` (0.2%) Very infrequently the DAQ can suffer a problem whereby its triggers get mixed up, causing normal beam pulses to be mislabelled as pedestal pulses. Such an event is called a "mixed-up spill." The signature for a mixed-up spill is a pedestal with abnormally high signal levels in some (but not necessarily all) of the ADC's. As a precaution, all data is rejected starting 500 pairs before the first detection of a mixed-up spill.

`bpmlinear` (1%) If the beam moves beyond ±1 mm from cavity center in any of the BPM's, those BPM's signals may exceed the linearity range of the rf mixer in the BPM electronics. Nonlinear BPM measurements can produce systematic errors in the asymmetry correction procedure. The `bpmlinear` cut eliminates all data for which the beam has drifted beyond ±1 mm in any of the BPM's, ensuring better than 99% linearity. This cut is stretched by 2 pairs before and 4 pairs after the offending pair.

`mollerAdc` (0.04%) This cut removes all pairs for which any Møller detector ADC is saturated or close to it (greater than 50000 channels or less than −100 channels). The gains of the Møller ADC's were set such that typical signal levels were 20000 to 30000 channels, far below the saturation regime where nonlinearity problems begin to appear. A highly saturated ADC reads out as negative. A saturated Møller ADC

channel signifies a temporary problem with the Møller electronics, in which case none of the Møller ADC data can be trusted, so that the entire pair is removed from the analysis. This very conservative approach still only results in a negligible amount of data loss.

`mollerAsym` **(0.07%)** Occasionally the DAQ fails to read out a Møller ADC board, resulting in some detector channels with nonsensical values. These can result in nonphysical channel asymmetries larger than ± 0.9 (i.e., ± 900000 ppm). If this every happens, the entire pair's data is likely compromised, and so is rejected. This is not a helicity-correlated cut, as only the absolute values of asymmetries are monitored.

`regSlopes` **(0.2%)** The regression procedure treats both timeslots independently. Occasionally, due to beam drop-outs and rate limits, one timeslot might only have a few pairs' worth of data by the time the other timeslot has collected its full 10k pairs, signifying the end of the regression cycle. Whenever one timeslot has fewer than 100 pairs for a given regression cycle, the data for that timeslot is simply removed. The errors on the computed slopes would otherwise be much too large for them to be of any use.

The combined data loss rate for all the above analysis cuts is approximately 18%. The total number of pairs passing all of the cuts is 186,040,008. Since the cuts are carefully designed not to involve any helicity-correlated quantities, they do not bias the final answer in any way. Explicit checks are made that eliminating any particular cut does not significantly shift the measured asymmetry. By guarding against things like hardware failures and poor quality beam, however, the cuts do manage to reduce the noise in the asymmetry measurement, decreasing the final statistical uncertainty.

4.4 Møller Asymmetry

This section presents the results of the Møller detector asymmetry analysis. Unless otherwise stated, linear regression has been used to correct all asymmetries for beam effects, with beam dithering being used as a means of checking the regression procedure. Furthermore, since the analysis has been finalized, all results presented will be "unblinded," that is, without the random offset from zero meant to protect the analysis from experimental

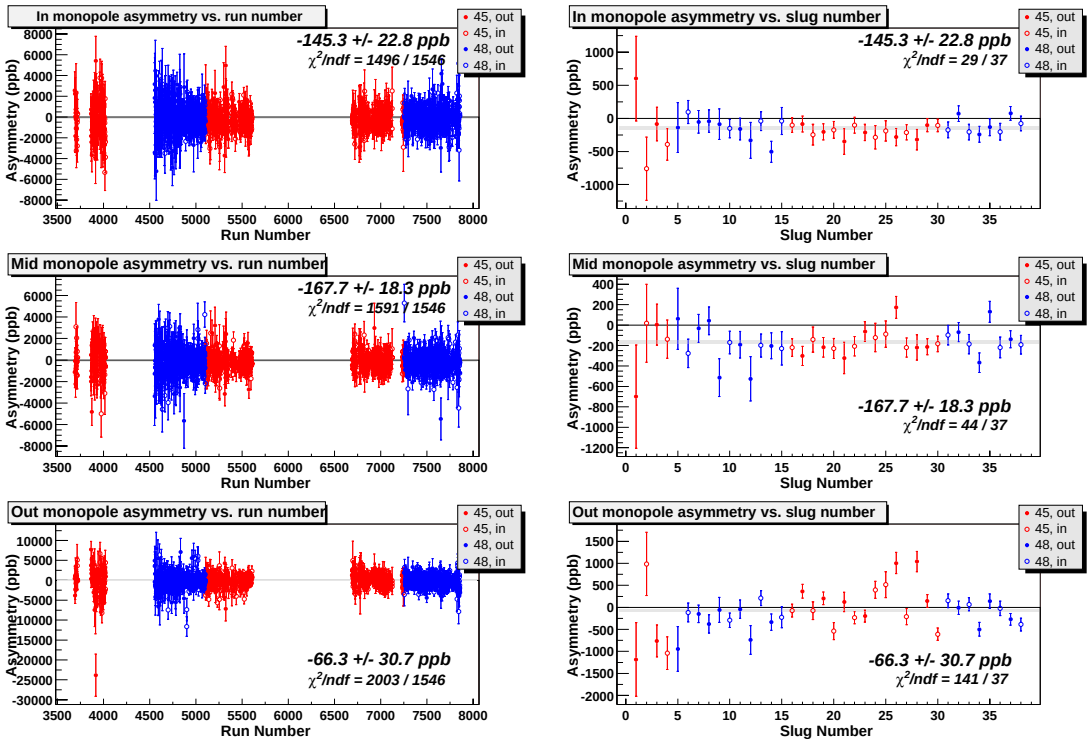

Figure 4.5: Ring monopole asymmetries for the Møller detector for Runs I and II, plotted for two very different timescales. In the left-hand plots, each data point represents approximately 1 hours' worth of data. The data is completely statistics-limited, and no outliers are observed. In the right-hand plots, each data point represents approximately 2 days' worth of data. Systematic effects now produce noticeable outliers, particular in the "out" ring's data.

bias. Finally, while beam effects have been corrected for, other corrections and dilutions due to background fluxes (e.g., photons and *ep*'s), as well as the normalization due to beam polarization, have not been taken into account. These will be discussed in Section 4.8.

In Figure 4.5, the asymmetries in each of the rings of the Møller detector are plotted separately as a function of time for two very different timescales. First, the results are given on a run-by-run basis, so that each data point represents approximately one hour's worth of data. Second, the results are shown on a slug-by-slug basis. The χ^2 per degree of freedom of the run-by-run plots are all fairly good at 0.96, 1.03, and 1.30 for the "in", "mid", and "out" rings, respectively. Big problems, however, are revealed by the "out" monopole slug plot, as evidenced by its reduced χ^2 value of 3.81. This is a good example of how subtle systematic effects only become visible after integration over long periods of time, typically days, whereas on shorter timescales of hours, they are completely hidden by statistical fluctuations.

This example shows that the run-by-run plots are of little use when diagnosing system-

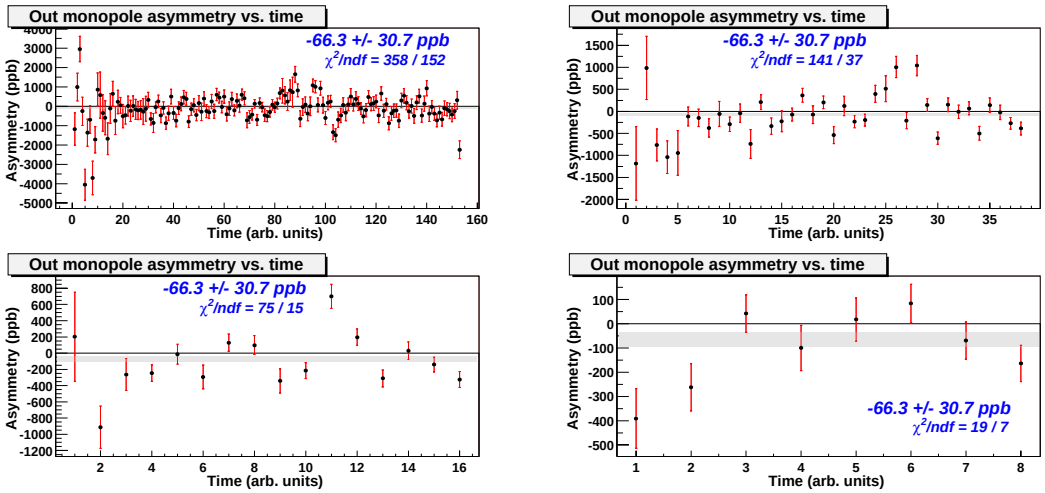

Figure 4.6: Monopole asymmetry for Møller "out" ring versus different timescales. The most useful timescale is that which results in the worse χ^2 per degree of freedom.

atic problems, and they will not be shown in the future. What timescale, then, is best when assessing the quality of the data and looking for potential systematic effects? That question is answered by a series as plots as shown in Figure 4.6. This figure again shows the "out" monopole asymmetry plotted as a function of time, but here the timescale is varied even further. First the asymmetries are plotted for every ten consecutive runs (being careful not to mix runs from different slugs). The χ^2 per degree of freedom has risen from the 1.30 it was for the run-by-run plot to 2.36. Next the asymmetries are again plotted versus slug. The χ^2 per degree of freedom has risen to 3.81. Next the asymmetries are plotted for roughly every two consecutive slugs, with more slugs sometimes being added in order to decrease the statistical error bar for a given data point. The χ^2 per degree of freedom has once more risen, this time to 5.03. Finally, the asymmetries of consecutive slugs are averaged even further before being plotted. Here each data point represents a little over one week's worth of data. The χ^2 per degree of freedom has *fallen* this time to 2.73, lower than it was for the slug plot.

Figure 4.6 sheds light on the answer to the question of which is the most appropriate timescale for the study of systematics. More and more significance is gained by the statistical outliers as the timescale is increased from runs to slugs and then to groupings of two or three consecutive slugs, corresponding to several days' worth of data. However, on longer timescales such as a week or more, the statistical outliers begin to lose their significance, implying that some form of systematics cancellation is occurring. Whatever systematics

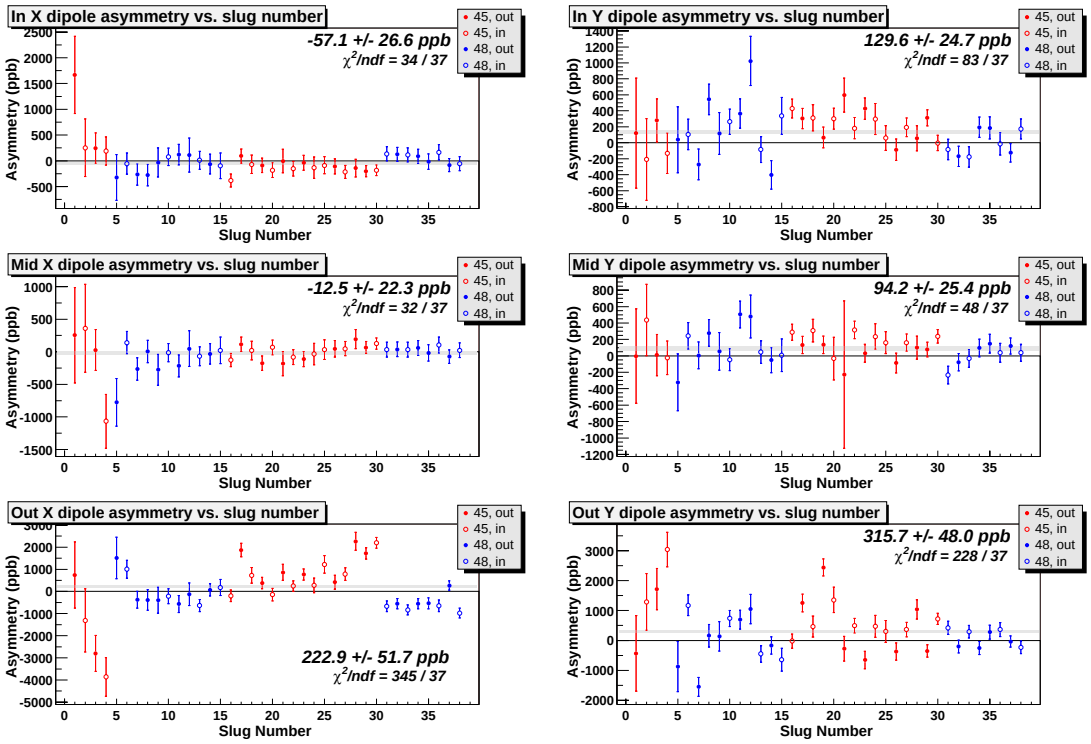

Figure 4.7: Ring dipole asymmetries for the Møller detector for Runs I and II. Dipole asymmetries are much more sensitive to beam motions and, hence, to beam-related asymmetry systematics.

there are thus seem to be most stable on the timescale of days, and eventually begin to randomize on the timescale of weeks. In the future, most quantities will therefore be plotted versus slug number, accepting that, strictly speaking, this timescale may not be the most sensitive to systematics effects, since the slug plot in Figure 4.6 does not have the maximum χ^2 per degree of freedom. Still, the slug is probably the most *relevant* timescale for most plots, since it is defined to be the longest consecutive period over which the experimental configuration is stable.

In Figure 4.7, the dipole asymmetries for each of the Møller detector rings are plotted versus slug number. These dipole combinations are extremely useful as extra-sensitive monitors for things like beam systematics, transverse polarization components, and uneven *ep* leakage, as was discussed in Section 3.4. Bearing that in mind, it is not unreasonable that some of the means are non-zero and that some of the plots have a poor χ^2 per degree of freedom. Both the "in" and "mid" rings of the Møller detector (constituting Region I) have X dipole asymmetries more or less consistent with zero, and with few statistical outliers. This can be taken as evidence, for instance, that the average vertical polarization component

is small, and that the errors stemming from the large Δx and $\Delta x'$ beam corrections for these rings are under control. The same rings, however, have significantly non-zero Y dipole asymmetries, evidence for a small horizontal polarization component due to an improperly set beam energy. In addition, there are a small number of statistical outliers in the Y dipole asymmetry plot for the "in" ring, where the χ^2 per degree of freedom is 2.2. This is perhaps evidence for a small dependence on fluctuating higher-order beam effects, though it could also simply be due to the beam energy being off by different amounts for different slugs.

In stark contrast to the relatively problem–free Region I, Region II of the Møller detector (the "out" ring) shows serious problems. In addition to the average X and Y asymmetry dipoles for Region II not agreeing with those of Region I, just as the Region II monopole does not agree with the Region I monopole, the Region II dipole asymmetries exhibit wildly erratic behavior, with many significant statistical outliers. The X and Y dipole slug plots for Region II have values of χ^2 per degree of freedom equal to 9.3 and 6.2, respectively. This observation is only one item in the long list of problems associated with Region II of the Møller detector. The three major pieces of evidence pointing towards serious problems with the "out" ring are summarized below:

1. The values of χ^2 per degree of freedom are poor for the "out" ring, for monopole as well as dipole slug plots. Either the first-order beam corrections are incorrect, or some higher-order unmeasured quantity (e.g., intra-pulse beam shape) is affecting the "out" ring. As will be seen in the next section, the first-order beam corrections appear to be correct with excellent precision. That means that higher-order beam effects are probably the cause. That the "out" ring is affected more than either the "in" or "mid"

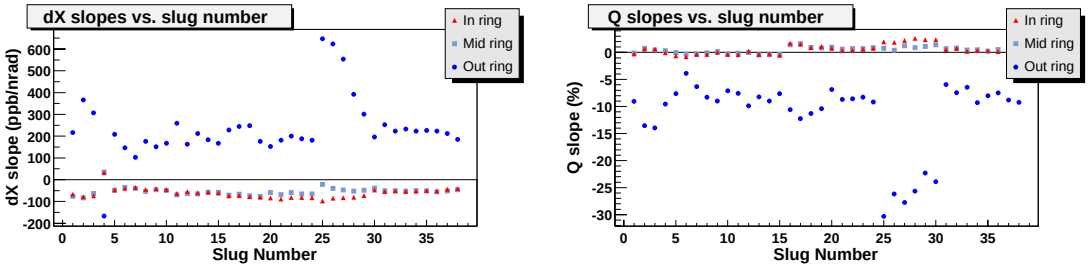

(a) (b)

Figure 4.8: Typical detector slopes for all three rings of the Møller detector for Runs I and II. Figure (a) shows the $\partial A/\partial x'$ slope (where x' is written as dX), while Figure (b) shows the charge slope. The charge slope is essentially the deviation (in percent) from perfect linearity. For instance, $\alpha_Q = 1\%$ is the same as $\alpha_Q = 0.01$ ppb/ppb.

ring is probably related to the fact that the "out" ring is typically five to ten times more sensitive to beam motions than either the "in" or "mid" ring. This can be seen in Figure 4.8(a), which shows a typical detector slope, $\partial A/\partial x'$, averaged for every slug for all three rings. Since the "out" ring is much more sensitive to beam motions, perhaps it should come as no surprise that it is more susceptible to higher-order beam effects.

2. The charge slope in the "out" ring is typically ten times bigger than the charge slopes for the "in" and "mid" rings, as shown in Figure 4.8(b). As discussed in Section 3.4.2, the charge slope absorbs some of the detector's dependence on higher-order beam effects. Charge slopes of 1% or less are acceptable given detector and toroid nonlinearities, but charge slopes of 10 or 20% as shown by the "out" ring clearly reveal an inadequacy of the correction procedure. It is also relevant that the period having the largest charge slope deviations (and with the largest slopes in general), Slugs 25 to 30, is precisely the same period showing some of the most erratic behavior in the "out" monopole and dipole plots.

3. Finally, regression against the sliced BPM data in Run III only shifts the overall asymmetry in Region I by 5 ppb. By contrast, the overall asymmetry in Region II moves by 20 ppb. Furthermore, all of the problems with the "out" plots having large values of χ^2 per degree of freedom get fixed by regression against these extra parameters in Run III. This supports the conclusion that the "out" ring possesses a large sensitivity to higher-order beam effects such as intra-pulse beam shape. Since the sliced BPM signals are not available in either Run I or II, and nothing in the data suggests that the sensitivity of the "out" ring is any less than it is for Run III, the overall asymmetry of the "out" ring for Runs I and II can only be trusted at the $\sim$20 ppb level. Combining the "out" ring's data with that of the "in" and "mid" rings would introduce a significant source of error. It is, therefore, better to drop Region II from the actual physics analysis, losing some statistical power, but at the same time avoiding a substantially increased systematic uncertainty.

From the observations listed above, it seems clear that the "out" ring suffers from a heightened sensitivity to higher-order beam effects, such as intra-pulse beam shape fluctuations. Errors stemming from such effects are described by Equation (3.25). Referring back

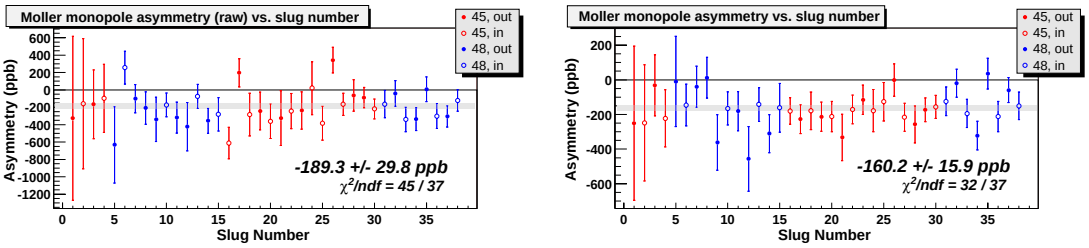

(a) (b)

Figure 4.9: Møller detector asymmetry versus slug number. Figure (a) shows the raw, uncorrected values, while Figure (b) shows the values after correction for beam effects, using linear regression.

to this equation, one notices that the size of the error produced is directly proportional to the derivatives of detector slopes with respect to the various beam parameters. Since the "out" ring's slopes are five to ten times larger than those of either the "in" or "mid" ring, it seems plausible that the derivatives of its slopes should also be relatively large. In fact, they might even be larger than one would naively expect, due to the complicated interactions of the "out" ring's signal (which sits at the very edge of the rapidly diminishing Møller signal peak, as shown in Figure 2.27) with collimators in the spectrometer system. Regardless, the "out" ring's slope derivatives are large, so that intra-pulse beam shape fluctuations due to wakefield effects produce errors via Equation (3.25) much larger in Region II than are produced in Region I. It should be noted that most of the significant statistical outliers in the slug plots of the "out" ring's monopole and dipole asymmetries occur during periods in which the beam energy is 45 GeV. The beam intensity is higher for 45 GeV than it is for 48 GeV, so that wakefield effects will be worse for 45 GeV data. Consequently, intra-pulse beam shape fluctuations will also be worse, which is consistent with there being more statistical outliers.

However it arises, the "out" ring's heightened sensitivity to higher-order beam effects means that its data should not be included in the Møller asymmetry analysis. Hence, in all future dscussions the "Møller detector" will refer only to the "in" and "mid" rings (Region I) of the main integrating calorimeter. Figure 4.9(a) shows the raw, uncorrected Møller asymmetries for every slug in Runs I and II, while Figure 4.9(b) shows the asymmetries after beam fluctuations have been removed using linear regression. Integrated over the entire course of Runs I and II, first-order beam effects can be seen to influence the Møller asymmetry result by -29.1 ppb, effectively increasing the absolute size of the underlying physics asymmetry by 18% (or nearly two statistical standard deviations). This is inferred from the

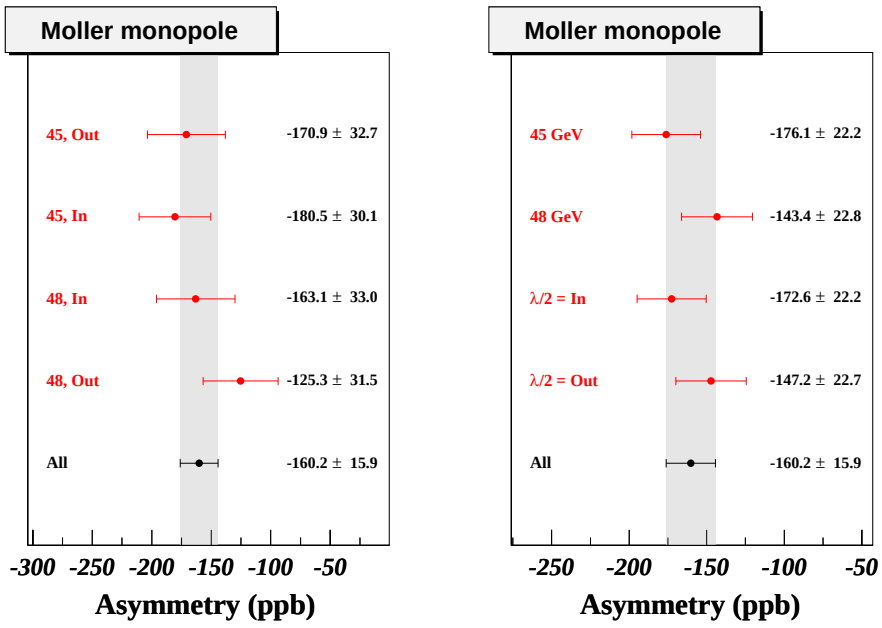

Figure 4.10: Møller detector asymmetry for various experimental configurations. Figure (a) contains four exclusive data sets, along with the overall result. Figure (b) contains four non-exclusive data sets, followed by the overall result. For instance, in Figure (b), the "45 GeV" data point contains both half-wave plate "in" as well as half-wave plate "out" data, while the "$\lambda/2 =$ in" data point contains both 45 GeV as well as 48 GeV data.

size of the correction induced by the linear regression procedure. The precision with which this correction is actually known (that is, the systematic uncertainty in the correction) is the subject of the next section. The first-order correction procedure also improves the average resolution of the Møller detector by nearly 50%, reducing it from roughly 410 ppm to 220 ppm (both numbers quoted on a per-pair basis). This in turn reduces the size of the final statistical uncertainty from nearly 30 ppb before regression to just under 16 ppb after regression. The correction procedure also improves the χ^2 per degree of freedom from 1.2 to 0.9, removing all of the statistically significant outliers.

One way to judge the overall consistency of the data and to see whether or not there are any gross systematics affecting it is to look at the plots contained in Figure 4.10. These plots give the Møller detector asymmetry for various experimental configurations (e.g., 45 vs. 48 GeV, half-wave plate in vs. out, etc.).[1] As discussed in Section 2.2.8, each configuration will in general provide sensitivity to different classes of systematics, so that the asymmetry

[1]The physics asymmetry flips sign when either the half-wave plate is inserted or the beam energy is switched to 48 GeV, and here (and everywhere else, unless stated otherwise) these factors of -1 have been taken into account.

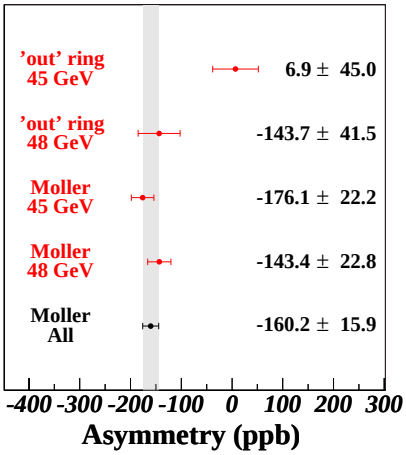

Figure 4.11: Energy-separated "out" monopole result, compared to that of the Møller detector. The 45 GeV "out" ring asymmetry disagrees with the overall Møller result, whereas the 48 GeV "out" ring asymmetry shows good agreement. This is evidence for exacerbated wakefield effects at lower beam energy (corresponding to higher beam intensity).

may potentially change depending on the configuration. For instance, it has already been noted that the problems in the "out" ring seem to be exacerbated in the 45 GeV data. This can clearly be seen from the plot contained in Figure 4.11. While the 48 GeV "out" ring asymmetry agrees with the overall Møller asymmetry (which is itself invariant with respect to beam energy), the 45 GeV "out" ring asymmetry shows a clear discrepancy, likely due to the larger wakefield effects at higher beam intensities.

If the asymmetries for two different experimental configurations disagree, then there is some unidentified systematic error that is affecting the data. Averaging the data from the two different configurations will yield some degree of cancellation, though at what level is unclear. On the other hand, invariance with respect to configuration changes (especially the energy change) is a good indication that beam-related systematics are under control. This is not a very powerful test, however, providing an upper limit on systematic contributions only at the level of the available statistics. A tighter limit can be placed by following the prescription presented in the next section.

4.5 Beam Asymmetries

Integrated over the course of both Runs I and II, the beam asymmetries average to roughly 370 ppb for charge, 0.4 keV for energy (corresponding to less than 10 ppb of the beam's total energy), less than 10 nm for position, and less than 0.35 nrad for angle, as shown

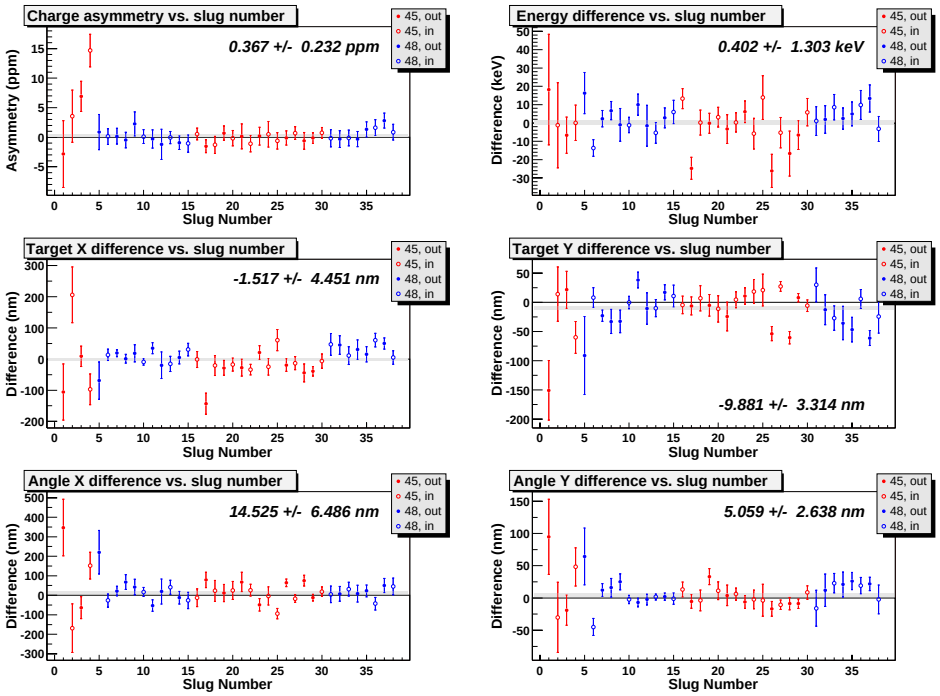

Figure 4.12: Beam asymmetries plotted against slug number for Runs I and II.

in Figure 4.12. This represents many orders of magnitude improvement over the typical starting points for beam asymmetries, which are on the order of 1000 ppm for charge, 1 MeV for energy, and 1 μm for position. As discussed in Section 2.2.7, both passive and active means of minimization are responsible for these drastic reductions in beam systematics.

In addition to beam control (i.e., controlling and minimizing beam asymmetries through the source setup), beam monitoring is an absolutely vital part of the experiment. Correcting for beam asymmetries and removing the noise in the detector caused by natural beam fluctuations relies on the ability to measure the beam parameters accurately and precisely on

Parameter	$\langle A \rangle \pm \delta A_{\text{stat}} \pm \delta A_{\text{syst}}$
Q	$367 \pm 233 \pm 3$ ppb
E	$0.4 \pm 1.3 \pm 0.2$ keV
X	$-1.5 \pm 4.5 \pm 0.4$ nm
Y	$-9.9 \pm 3.3 \pm 0.9$ nm
X'	$0.36 \pm 0.16 \pm 0.02$ nrad
Y'	$0.13 \pm 0.07 \pm 0.02$ nrad

Table 4.1: Summary of beam asymmetry data for Runs I and II. Each beam parameter has three quantities associated with it. The first gives the average right-left asymmetry (or difference), integrated over both Runs I and II. The second gives the statistical uncertainty on the measurement, which is due purely to beam jitter. The third quantity is the systematic uncertainty, resulting from the finite device resolution.

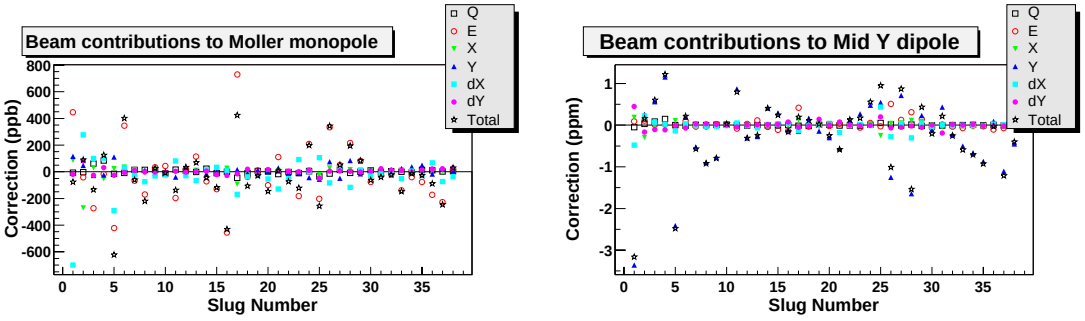

(a) (b)

Figure 4.13: Monitors exhibiting heightened sensitivity to one particular beam parameter. For example, the Møller monopole is most sensitive to the energy correction, whereas the "mid" ring Y dipole is most sensitive to the Δy correction.

a pulse-to-pulse basis. Grouping the beam monitoring devices in pairs allows for their precision to be determined by comparing the measurements of neighboring devices. Table 4.1 summarizes the beam asymmetry data for Runs I and II, showing both how well each quantity's helicity-correlated asymmetry converges to zero, compared to what is allowable given the beam jitter, and also how well each asymmetry is measured.

Of course, the results presented in Table 4.1 are impossible to interpret without knowing the detector slopes α_ξ, which parameterize the detector's overall sensitivity to each beam parameter. The sizes of the beam asymmetries themselves are not so important, nor how well they are measured, but rather the size of the false asymmetries they induce in the Møller detector, and how well these false asymmetries can be corrected for. Determining the extent to which the beam corrections work correctly is the subject of this section.

In general, the different moments (i.e., monopole, X dipole, and Y dipole) of each detector ring's asymmetry data will possess heightened sensitivity to a particular beam parameter. For instance, as shown in Figure 4.13(a), the corrections to the Møller monopole are often dominated by the energy correction. For the corrections in a given slug to be dominated by the energy correction requires two things. First, it requires that the energy slope is significantly non-zero. Second, it implies that the energy difference (integrated over the slug) is significantly non-zero. The energy slope is stable at the few percent level over time, but of course the energy differences constantly fluctuate about zero. Hence for any given slug the Møller monopole corrections may or may not be dominated by the energy correction, but *often* they are. Similarly, as pictured in Figure 4.13(b), the corrections to the mid ring's Y dipole are often dominated by the y correction. It is possible to find monitors

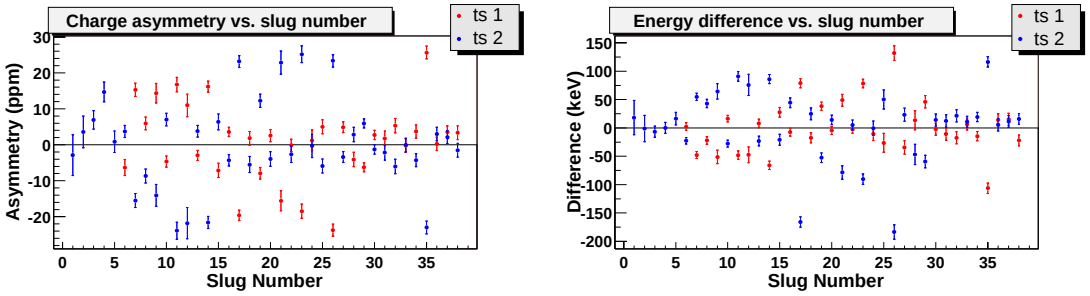

Figure 4.14: Large timeslot differences in the beam asymmetries. The intensity asymmetry (IA) feedback loop does not null the charge asymmetry for both timeslots separately, only the sum of both timeslots. Therefore, since each timeslot is sensitive to a different phase of the 60 Hz line noise, large timeslot differences in the charge asymmetry (and, by extension, all other beam asymmetries) can arise.

that are mainly sensitive to just one beam parameter in every case, with the hardest case being that of y'. Perhaps because the $\Delta y'$ differences themselves are never too large, it is rare to find a monitor whose corrections are dominated by the y' correction. However, for at least one slug's worth of data, the out ring's Y dipole does have its corrections dominated by that of y'.

The next step in estimating the uncertainties in each of the corrections is to somehow divide the data such that the corrections needed by the two subsets are very different. The two timeslots serve this purpose extremely well, since the beam asymmetries in each of the timeslots can be radically different, as shown in Figure 4.14. This makes sense, given that each timeslot is essentially a completely separate experiment, sensitive to a different phase of the considerable 60 Hz line noise. The helicity-correlated feedback loops are not programmed to pay attention to timeslot, so they null the total charge asymmetry. The two timeslots can (and do) therefore have wildly different charge asymmetries. For example, the charge asymmetry in one timeslot can be +10 ppm, while that in the other might be −10 ppm. Since every other beam parameter is influenced by charge at some level, these large timeslot differences in the charge asymmetry will tend to produce large timeslot differences in the other beam asymmetries as well.

Large timeslot differences in the beam asymmetries, it turns out, are a good thing, for they allow an accurate estimation of the uncertainty in the correction procedure. Consider a beam correction ΔA^*_ξ in a special monitor (denoted by the asterisk $*$) that is particularly sensitive to a given beam parameter ξ. If the correction $\Delta A^*_{\xi_1}$ needed in one timeslot is radically different than the correction $\Delta A^*_{\xi_2}$ needed in the other timeslot, and if the two

timeslots' asymmetries A_1^* and A_2^* agree after those very different corrections are applied, then the first-order correction procedure can be said to work very well. As discussed in Section 3.4.1, the relative error in the procedure of correcting for helicity-correlated differences in beam parameter ξ is equal to:

$$\text{relative error of } \Delta\xi \text{ correction} = \left| \frac{A_1^* - A_2^*}{\Delta A_{\xi_1}^* - \Delta A_{\xi_2}^*} \right| \tag{4.2}$$

Multiplying this fraction by the size of the correction ΔA_ξ needed by the Møller detector, one arrives at the first-order error in the Møller asymmetry due to the $\Delta\xi$ correction:

$$\text{absolute error of } \Delta\xi \text{ correction} = \left| \frac{A_1^* - A_2^*}{\Delta A_{\xi_1}^* - \Delta A_{\xi_2}^*} \right| \cdot \Delta A_\xi \tag{4.3}$$

This technically assumes that the total correction in each monitor is large and dominated by one beam systematic (the one whose correction is being "tested"). If the correction being tested is statistically insignificant, then the denominator of Equation (4.3) will be small with respect to its error. Hence the denominator's fractional error will be large. On the other hand, if the correction being tested is large, then the denominator's fractional error will be negligible compared to that of the numerator (assuming that the correction works more or less correctly at the 10% level), so that Equation (4.3) will yield an answer with a sensible error bar.

The other possibility is that the correction being tested is large, but not much larger than other corrections. In that case, correlations between parameters can arrange it so that systematic errors cancel each other out, resulting in final asymmetries that may agree despite individual beam corrections being flawed. While arguably still correct, estimating the first-order systematic uncertainty using a monitor in which multiple corrections are equally important should return a less conservative estimate of the error. All of the assumptions listed here (that the total correction in each monitor is large, dominated by one beam systematic, and works at least at the 10% level) can be tested by inspecting the result of Equation (4.2). If the calculated error turns out to be significantly non-zero, then one of the assumptions is likely faulty.

An example of this procedure can be seen in Figure 4.15(a), which shows the Møller monopole asymmetry, both its raw value as well as its corrected value, plotted versus slug

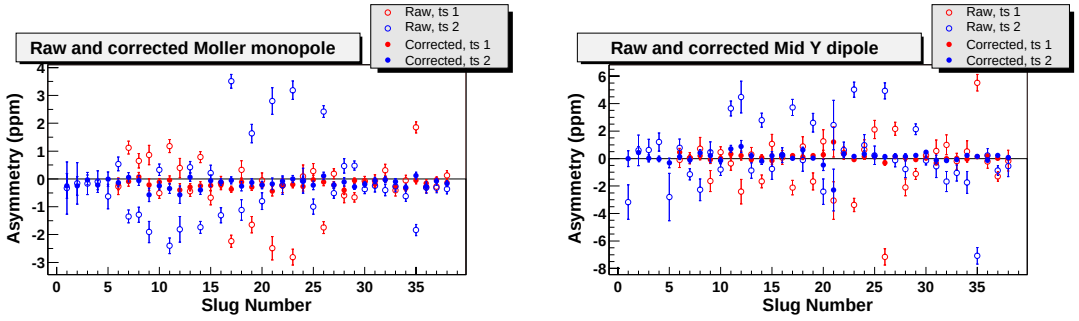

(a) (b)

Figure 4.15: Significant timeslot differences in the beam corrections can be used to assign an error to the correction procedure. If a particular monitor's raw asymmetries for either timeslot are drastically different, but its corrected asymmetries both agree, then the first-order correction procedure can be said to function properly.

for both timeslots separately. The Møller monopole is most sensitive to energy, so it can be used to obtain a relative uncertainty for the energy correction. There are many slugs with significant timeslot differences in the energy correction, Slugs 17, 23, 26 and 35 among them. For all of these slugs, however, the corrected Møller asymmetries agree very well with regard to timeslot, despite the fact that drastically different corrections are needed. For instance, using Equation (4.2), the relative error for Slug 23 can be calculated as the following:

$$\text{Slug 23}: \quad \left| \frac{(-0.21 \pm 0.11 \text{ ppm}) - (-0.003 \pm 0.13 \text{ ppm})}{-2.51 \text{ ppm} - 2.94 \text{ ppm}} \right| \longrightarrow 3.7 \pm 3.2 \ \% \tag{4.4}$$

Slugs 17, 26, and 35 produce similar estimates of 4.4 ± 2.3 %, 5.4 ± 4.6 %, and 4.5 ± 4.7 %, respectively. In fact, if this procedure is repeated for all other slugs where the assumptions can be verified to hold, the error obtained is on average five percent (4.8 ± 1.4 %), as shown in Figure 4.16(a).

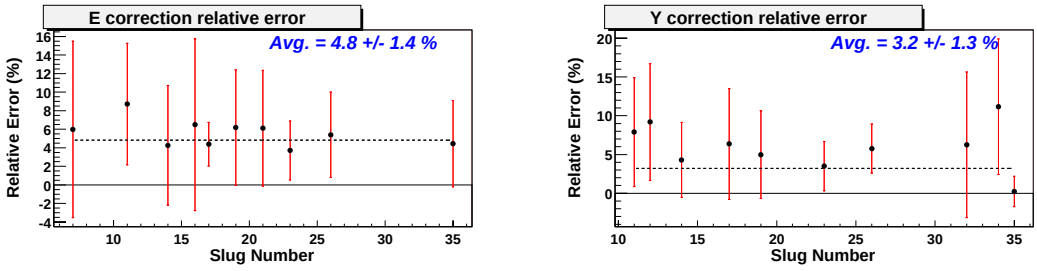

(a) (b)

Figure 4.16: Relative errors of the ΔE and Δy corrections versus slug.

Parameter	*Monitor*	*Rel. Error*	*Correction*	*Error*
Q	—	—	-0.8 ppb	—
E	Møller A	5%	17.7 ppb	0.89 ppb
X	mid A_x	1%	1.2 ppb	0.01 ppb
Y	mid A_y	3%	-5.5 ppb	0.17 ppb
X'	out A_x	3%	20.3 ppb	0.61 ppb
Y'	out A_y	6%	-3.7 ppb	0.22 ppb
TOTAL:	—	4%	29.2 ppb	1.1 ppb
Cross-check:	Møller A	5%	29.2 ppb	1.5 ppb

Table 4.2: Summary of first-order uncertainties in beam corrections. The second column gives the monitor whose corrections are most often dominated by each parameter. Here A denotes a monopole asymmetry, A_x denotes an X dipole asymmetry, and A_y denotes a Y dipole asymmetry. The fourth column gives each beam parameter's correction to the Møller monopole asymmetry. No error estimate for the charge correction is given, since this is handled by the linearity correction to be discussed later.

As another example of the procedure, Figure 4.15(b) shows the mid ring's Y dipole asymmetry, both its raw and corrected values, plotted versus slug for both timeslots separately. Using this plot in conjunction with Figure 4.13(b), one can find many slugs in which the corrections are large, have big timeslot differences, and are dominated by the Δy correction, Slugs 11, 12, 14, 23, 26, and 35 among them. Slug 35 sets what is by far the most significant limit:

$$\text{Slug 35}: \quad \left| \frac{(0.13 \pm 0.17 \text{ ppm}) - (0.16 \pm 0.16 \text{ ppm})}{5.09 \text{ ppm} - (-6.83 \text{ ppm})} \right| \quad \longrightarrow \quad 0.2 \pm 2.0 \ \% \tag{4.5}$$

The other slugs where the assumptions can be verified to hold all return errors at the few percent level. The average uncertainty for the Δy correction for all of these slugs is calculated to be 3.2 ± 1.3 %. The individual slug relative errors for the Δy correction are shown in Figure 4.16(b).

Table 4.2 summarizes the results of using Equation (4.3) to estimate uncertainties for all of the beam corrections for the Møller asymmetry. For each correction, all of the slugs for which the assumptions underlying the equation can be shown to hold are used to calculate an average relative error. In general, it makes a difference whether or not the absolute value is taken in Equation (4.3) before finding the average of all the slugs' errors. Taking the absolute value before finding the average tends to overestimate the relative error, but usually not by much. Since a conservative estimate is desirable, this is the approach used. Adding all of the individual errors in quadrature, the total uncertainty in the asymmetry

correction procedure due to first-order beam effects is 1.1 ppb, roughly 4% of the 29.2 ppb total correction.

As a cross-check, the last row of Table 4.2 shows the error on the total correction obtained by using the Møller monopole itself as the sensitive monitor. Here the strategy for picking the "best" slugs is somewhat different. Instead of picking slugs where one correction dominates, here one wants slugs in which many different corrections are important. As discussed earlier, natural correlations between beam parameters may actually cause systematic errors to cancel one another out, returning an estimate for the total error that is perhaps too small. In this case, however, as one can see by comparing the last two rows of the table, this method actually returns a relative error of 5%, slightly larger than the 4% obtained by adding the individual errors up in quadrature. This could simply be due to statistical fluctuations. To be conservative, this larger 1.5 ppb error is taken, after rounding up, as the final first-order uncertainty on the asymmetry correction procedure:

$$\text{total } 1^{\text{st}}\text{-order beam correction (Runs I \& II)} = 29.2 \pm 2.0 \text{ ppb}$$

Unfortunately, the above procedure cannot be used to place a limit on the size of the error possibly stemming from higher-order beam effects, since its methodology relies on finding slugs in which the beam contributions to a given monitor's asymmetry are dominated by a single first-order systematic whose correction works at the 10% level or better. By construction, the effects of higher-order asymmetry systematics will be highly suppressed, so that no meaningful limit on their relative importance can be made. Other methods, namely those discussed at the end of the last chapter, must be employed.

The first method involves "worst-case scenario" calculations, or applying Equation (3.25) under various conditions in an effort to reveal the order of magnitude of the effects that can be reasonably expected to appear in a given monitor. These calculations rely on three pieces of information: (a) the dependence of the detector slopes on all of the other beam parameters, (b) intra-pulse variations in the beam parameters (the beam shape), and (c) intra-pulse variations in the beam parameters' helicity correlations. The first piece of information comes from a Monte Carlo simulation, while the other two come from a mix of direct observations, measurements, and sensible assumptions. For the given inputs, one finds the effects produced in the "in" and "mid" rings generally bounded by 10 ppb, while in the "out" ring, because it is so much more sensitive to beam motions, the effects can be

five to seven times larger. These numbers are for integration over the course of the entire experiment. On shorter timescales, fluctuations might produce larger effects.

The second method for estimating the experiment's sensitivity to higher-order beam effects assumes a single dominating source of higher-order systematics common to the Møller detector and the "out" ring [105]. In this case, one can write the measured asymmetries in each of these monitors as the sum of the actual physics contributions and the contributions coming from the single dominating systematic asymmetry A_{syst}:

$$\begin{aligned} A_{meas,1} &= A_{phys,1} + \alpha_1 \cdot A_{syst} \\ A_{meas,2} &= A_{phys,2} + \alpha_2 \cdot A_{syst} \end{aligned} \tag{4.6}$$

where α_i relates how sensitive the measured asymmetry $A_{meas,i}$ (here i refers to a given monitor, either the Møller detector or the "out" ring) is to the systematic asymmetry A_{syst}. The coefficients α_i may be very different for the two monitors. The physics contributions $A_{phys,i}$ may also be different, due to the disparity in size of the ep contributions. Solving Equation (4.6) for $\alpha_1 \cdot A_{syst}$, which is defined to be the higher order systematic's contribution to the Møller asymmetry, one finds:

$$\alpha_1 \cdot A_{syst} = [(A_{meas,1} - A_{meas,2}) - (A_{phys,1} - A_{phys,2})] \cdot \frac{\alpha_1}{\alpha_1 - \alpha_2} \tag{4.7}$$

It is worth mentioning that any potential "new physics" effects causing $A_{meas,i}$ to be different from $A_{phys,i}$ will completely drop out from this equation, since such effects will contribute to the Møller detector and to the "out" ring equally. By somehow measuring (or estimating) the coefficients α_i, or at least their ratio, the systematic contributions to the two monitors can be obtained.

As a first approximation, one can use the measured charge slopes as estimates for the coefficients α_1 and α_2 in Equation (4.7). As long as the real coefficients scale with charge (which should be the case, assuming the underlying problem is due to intra-pulse variations in beam shape), one should be able to simply use α_{Q1} and α_{Q2} in Equation (4.7) instead of α_1 and α_2. In fact, only the slopes' relative sizes are important. Referring back to Figure 4.8(b), one sees that $\alpha_{Q2} \approx 10 \cdot \alpha_{Q1}$. Using $\alpha_1 = 1$ and $\alpha_2 = 10$ in Equation (4.7) (the units being unimportant), one finds the higher order systematic's contribution to the Møller asymmetry in Runs I and II to be between ± 5 and ± 10 ppb.

A second method for extracting the ratio of α_1 to α_2 works by noise analysis, recognizing that the more sensitive a monitor is to the underlying systematic, the worse its asymmetry distribution will look on a slug-by-slug basis. Essentially, then, one uses the χ^2 per degree of freedom of the two monitors' slug plots as "proxy" variables for the coefficients α_i. Again, only their ratio is important, so that to a good approximation one may write [105]:

$$\left|\frac{\alpha_1}{\alpha_2}\right| \approx \frac{\delta A_{meas,1}}{\delta A_{meas,2}} \sqrt{\frac{\chi^2_{dof_1} - 1}{\chi^2_{dof_2} - 1}} \tag{4.8}$$

where $\delta A_{meas,i}$ is the statistical uncertainty in the final asymmetry measurement for a given monitor, and $\chi^2_{dof_i}$ is the χ^2 per degree of freedom of that monitor's asymmetry slug plot. Of course, if a monitor's reduced χ^2 happens to be smaller than unity (which could happen due to statistical fluctuations), it is still possible to calculate the statistical error on the ratio $|\alpha_1/\alpha_2|$, and one should compute this error regardless. In the end, the results returned by this method are consistent with those of the previous method, namely that the size of the higher order systematic's contribution to the Møller asymmetry in Runs I and II is between 7 and 15 ppb.

It is instructive to use these methods to estimate the size of the contribution of higher-order beam effects to the Møller asymmetry in Run III, for which time-sliced BPM data is available. These methods predict that, given a single dominating source of higher-order systematics, the contribution to the Møller asymmetry in Run III should be within ± 15 ppb, which is a pretty large range. Nevertheless, it is consistent with the actual results, which show that regression against the extra sliced BPM signals shifts the final Møller asymmetry by 5 ppb.

Preliminary Run III analysis shows that all of the gross problems discussed previously, relating mainly to the "out" ring, get fixed by regression against the sliced BPM signals, giving evidence that it is indeed predominantly intra-pulse beam fluctuations that are responsible for the statistical outliers in the various monitors' asymmetry distributions. Given that this is the case, the main assumption underlying the methods just described for estimating the contributions of higher-order asymmetry systematics, that there be one dominating systematic, and that it be intra-pulse beam fluctuations, appears valid. Therefore, the limits arrived at above can be taken as fair approximations, though there is evidence that they might be overestimations, since the Run III prediction ($\Delta A_{syst} = \pm 15$ ppb) is three times

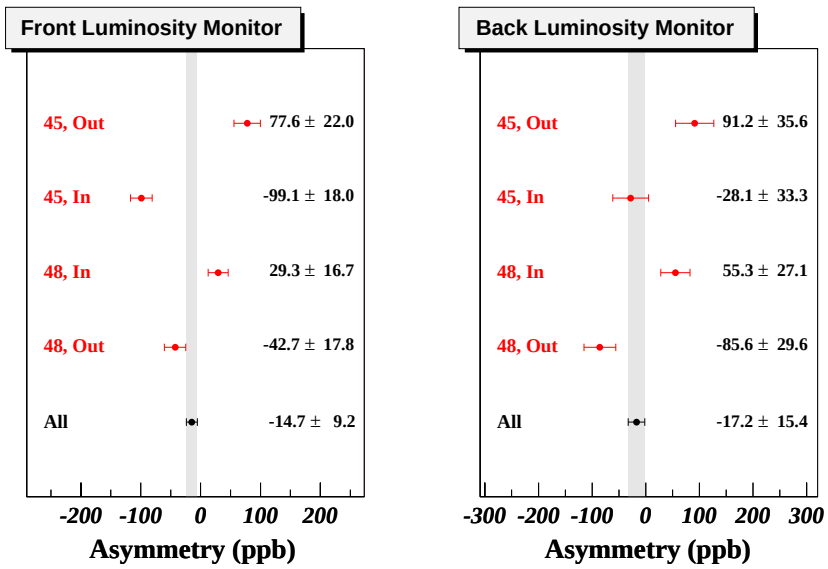

Figure 4.17: Front and back luminosity monitor asymmetries for different experimental conditions.

the observed shift in the Møller asymmetry ($\Delta A = 5$ ppb) when the sliced BPM data is regressed against. Bearing all this in mind, a best estimate for the uncertainty in the Runs I and II Møller asymmetry due to higher-order beam effects is ± 10 ppb.

4.6 Luminosity Monitor Asymmetry

The luminosity monitor is designed to function primarily as a probe for false asymmetries, including beam-related ones as well as those induced by electronics cross talk. The physics asymmetry it is sensitive to is predicted to be small (approximately -10 ppb), at least compared to the statistical precision that is realistically achievable. In addition, the luminosity monitor is very sensitive to beam motions. These two facts make it an obvious choice as a false asymmetry probe, as any significant deviation from zero in its asymmetry can be immediately interpreted as being due to a helicity correlation in some parameter of the beam.

The asymmetries observed in both rings of the luminosity monitor for Runs I and II are shown in Figure 4.17. The front and back rings can be seen to agree, the front ring's asymmetry being -14.7 ± 9.2 ppb, and the back ring's asymmetry being -17.2 ± 15.4 ppb. However, clear indications of systematics are visible in Figure 4.17, particularly in the 45 GeV data. As with the "out" ring, this could be due to increased wakefield effects

at the higher beam intensity exacerbating higher-order beam effects. Regardless of what is causing the deviations, the systematics cancellations provided by the half-wave plate toggles and the energy changes seem to work very well, reflected in the fact that the luminosity monitor's final asymmetry agrees with the physics prediction at the 10 ppb level. While this is indeed an encouraging result, its full implications for the Møller asymmetry are unclear, since the two detectors' signals are very dissimilar. That is, what affects one detector does not necessarily affect the other, and vice versa.

4.6.1 Spotsize Asymmetry Correction

One systematic source that affects both the luminosity monitor as well as the Møller detector (though by different amounts) is the so-called spotsize systematic, whereby helicity correlations in the beam's spotsize cause target density fluctuations, which in turn induce false asymmetries in the detectors. In order to set a limit on how severely spotsize-induced target density fluctuations could be affecting the detectors' asymmetries, coefficients α_σ relating how sensitive a given detector's asymmetry is to right-left differences in the beam's spotsize σ (as measured by the wire array) are determined empirically for every run in which the wire array was present in the beamline.[2] For the Møller detector, typically $|\alpha_\sigma| \approx 10 - 20$ ppm/mm^2, while for the luminosity monitor, typically $|\alpha_\sigma| \approx 30 - 100$ ppm/mm^2. The average right-left difference in the beam's spotsize for Runs I and II is $(2.8 \pm 8.2) \times 10^{-6}$ mm^2, corresponding to right-left differences in the beam's horizontal and vertical widths of 0.8 ± 1.9 nm and 0.1 ± 1.8 nm, respectively. This should be compared to a typical beam spotsize of $3 - 4$ mm^2 (roughly a round spot 2 mm in diameter). The spotsize asymmetry contribution to the Møller asymmetry for Runs I and II is found to be less than 0.5 ppb. No correction will therefore be made, but the actual uncertainty contribution will be taken as ± 1 ppb.

[2] Long exposure to the beam can damage the wire array. This was discovered during Run I, when the wire array was damaged halfway through the run. In addition, the wire array slightly increases the backgrounds due to *ep* and Rutherford scattering, which affects the asymmetry measurement by a few ppb [109]. For these reasons, it was only inserted for approximately two hours every day during Run II. The spotsize asymmetry correction is therefore based upon measurements from a subset of the data that should be representative of the whole.

4.7 Pedestal Asymmetries

The term "pedestal asymmetries" refers to asymmetries that are present in ADC channels even in the absence of any beam (i.e., even in pedestal pulse data). Such asymmetries can be an indication of electronic cross talk either from the source PMON system or from the CP cell pulsing at ± 2600 V in (as is necessary) a helicity-correlated manner. It is important to estimate at what level such false asymmetries could be affecting the Møller asymmetry measurement.

In order to set a rigorous limit on the size of potential cross talk in the toroids and detector channels, a series of runs were taken during the summer between Runs I and II with every subsystem operating nominally, but with no beam actually being delivered.[3] For instance, the source Pockels cells were pulsing at their usual voltages and the PMON system was duly determining the helicity of each "beam" pulse and broadcasting it to the rest of the experiment (after its normal one-pulse delay). In addition, dc biases were fed into a few detector ADC channels and a toroid ADC channel such that their signals (in ADC counts) were roughly the same size as they were when beam was being delivered. Also, a toroid calibrator was used to produce a signal (again, of roughly equal size as that generated by the actual beam) in two of the actual toroids in the alcove.

Roughly 20 million pairs were taken during this period of a few weeks in an effort to probe for electronics cross talk effects. Because the DAQ system was experiencing some serious problems that were to be fixed shortly thereafter, the efficiency of the DAQ–related analysis cuts was quite low for this period, so that less than 10 million pairs were successfully analyzed. Nevertheless, this data is able to place constraints on pedestal asymmetries in the detector and toroid electronics at the $3-5$ ppb level.

In terms of how such asymmetries might potentially influence the Møller asymmetry, cancellation due to the half-wave plate and energy flips must be taken into account. When either the half-wave plate is inserted or the beam energy is 48 GeV, a negative factor gets applied to all pedestal asymmetries. If equal amounts of data are taken in so-called "positive" and "negative" states, then all pedestal asymmetries will exactly cancel. The

[3] Because the electricity rates were prohibitively expensive, the linac could not be operated during this period. Therefore no beam could be delivered, making it an ideal opportunity for "false asymmetry" runs.

exact suppression factor f is calculated in the following way:

$$f = \left| \frac{(1/\delta A_+)^2 - (1/\delta A_-)^2}{(1/\delta A_+)^2 + (1/\delta A_-)^2} \right| = 0.06 \tag{4.9}$$

where δA_+ and δA_- are the statistical uncertainties of the Møller asymmetry for all "positive" and "negative" data, respectively.[4] Therefore, a pedestal asymmetry of $3-5$ ppb will only contribute to the Møller asymmetry at the $0.2-0.3$ ppb level.

Except for the few channels mentioned above into which dc biases were fed, all other ADC channels were merely monitored during this period for helicity-correlated pedestal differences, not asymmetries. Most of the channels show pedestal differences smaller than a millichannel, but some of the BPM ADC channels show pedestal differences as large as a few millichannels, significant at the two or three–sigma level. In particular, the energy- and angle-sensitive BPM's show the biggest effects, presumably because these BPM's are located furthest from their associated electronics, so that they are connected to the longest copper cables. These long copper cables make more sensitive antennas to any helicity-correlated noise broadcast from the source.

To interpret how helicity-correlated BPM pedestal differences of a few millichannels might influence the Møller asymmetry, one must convert ADC channels to real position units. Using the BPM calibration constants, one finds that a few millichannels corresponds to a few tenths of a keV in energy, a few nanometers in position, and nearly one hundred picoradians in angle. Mutliplying these differences in turn by typical Møller detector slopes of $10-20$ ppb/keV for energy, 1 ppb/nm for position, and $10-50$ ppb/nrad for angle, one finds their contributions to the Møller asymmetry to amount to a few ppb in each case. Again, however, the suppression factor of 0.06 resulting from half-wave plate flips and energy changes, calculated in Equation (4.9), applies. This means that false asymmetries due to pedestal differences of a few millichannels in the BPM's should contribute to the Møller asymmetry at the sub-ppb level, which can be neglected.

In fact, it is possible to check this analysis using the actual production dataset. Each BPM is comprised of three resonant cavities, two cavities sensitive to position (the X and Y cavities) and one sensitive to charge (the Q cavity, also called the ϕ cavity). If a BPM's Q-cavity signal is normalized to charge as measured by a toroid, all beam intensity jitter

[4] Explicitly, for 45 GeV running, all $\lambda/2$ = "out" data is "positive" and all $\lambda/2$ = "in" data is "negative". For 48 GeV running, all $\lambda/2$ = "in" data is "positive" and all $\lambda/2$ = "out" data is "negative".

is removed. The signal therefore becomes a very quiet monitor sensitive only to electronics noise. The following charge-normalized Q-cavity asymmetries A_ϕ are thus observed:

$$A_\phi = \frac{\phi_R/q_R - \phi_L/q_L}{\phi_R/q_R + \phi_L/q_L} \tag{4.10}$$

where ϕ_R and ϕ_L are the BPM Q-cavity signals, and q_R and q_L are the beam charge measurements, for a right and left helicity beam pulse, respectively. If any of the measured asymmetries A_ϕ are in fact non-zero, then false asymmetries in the electronics can only be to blame. In addition, if electronics cross talk is indeed producing significant asymmetries in these BPM Q-cavity signals, they should be stable with respect to both half-wave plate flips and energy changes.

The analysis of the charge-normalized BPM Q-cavity asymmetries in Runs I and II yields results that are consistent with those of the special false asymmetry runs conducted during the summer. Namely, a few BPM's show evidence for pedestal differences of a few millichannels, which have already been seen to be capable of producing effects in the Møller detector at the few ppb level. After the suppression factor ($f = 0.06$) due to experimental configuration changes, however, the size of the effects produced in the Møller detector gets reduced to less than one ppb, which can safely be ignored. This method has the advantage of using the actual physics data to set limits on the contributions to the Møller asymmetry coming from helicity-correlated pedestal differences in the BPM's, yet confirms the results already obtained.

4.8 Final Experimental Asymmetry

The final step in the asymmetry analysis is to correct for the effects of physics backgrounds, which may themselves contain physics asymmetries. In general, all backgrounds, whether they carry an asymmetry or not, will dilute the Møller asymmetry measurement. As long as a particular background flux is small relative to the Møller signal ($\phi_{bkg}/\phi_0 \equiv f_{bkg} \lesssim 0.1$, where ϕ_0 is the Møller signal flux and ϕ_{bkg} is the background flux), this dilution takes on a simple form: $A_{meas} = A_{\rm PV}/(1 - f_{bkg})$. Additionally, again assuming that the background flux is relatively small, an asymmetry A_{bkg} in the background flux will produce a shift in the measured asymmetry that is also easily expressed as $\Delta A_{bkg} = f_{bkg} \cdot A_{bkg}$. The effects

of a physics background can therefore be corrected for using the following formula:

$$A_{\rm PV} = \frac{A_{meas} - f_{bkg} \cdot A_{bkg}}{1 - f_{bkg}}; \quad (\Delta A_{bkg} \equiv f_{bkg} \cdot A_{bkg}) \tag{4.11}$$

Finally, the polarization and linearity corrections must be taken into account. Both enter in as scale factors, so that the final formula for determining the actual parity-violating cross-section asymmetry in Møller scattering from the measured Møller detector asymmetry can be written as:

$$A_{\rm PV} = \frac{1}{\epsilon \cdot P_{beam}} \cdot \frac{A_{meas} - f_{bkg} \cdot A_{bkg}}{1 - f_{bkg}}; \quad (\Delta A_{bkg} \equiv f_{bkg} \cdot A_{bkg}) \tag{4.12}$$

where P_{beam} is the beam polarization and ϵ is the average degree of nonlinearity exhibited by the detector signals with respect to signal flux. Since many of the background flux measurements will involve extrapolations from very small detector signal levels, it is appropriate to discuss the linearity correction first.

4.8.1 Linearity Correction

On the one hand, there has already been evidence presented that the average nonlinearity for the Møller detector is small. Referring back to Figure 4.8(b), one sees that the charge slopes α_Q for the "in" and "mid" rings, which among other things are sensitive to detector nonlinearities, rarely exceed 1%. In fact, the average charge slope for the Møller detector for Runs I and II is 0.7%. However, in addition to being sensitive to detector nonlinearities, the charge slopes are also sensitive to higher-order beam effects, so this cannot be taken as strong evidence for the absence of nonlinearities. Two more rigorous approaches, each addressing a different type of nonlinearity, will now be described.

There are two sources of nonlinearity to consider. First, for very high signal fluxes, the PMT's can saturate, through either anode or cathode effects. This will produce the type of nonlinearity that is most relevant for normal conditions. Under special conditions, however, like those discussed in the following section, the signal flux can be extremely small, less than 1% of its normal strength. In these cases, a second type of nonlinearity, resulting from the Møller detector electronics distorting signals of a few millivolts, can arise.

In order to constrain the detector's linearity under nominal conditions, special runs

were conducted during Run II in which the polarized iron foil was inserted (and the liquid hydrogen target removed) and various "filters" placed over some of the detector PMT's. These filters were aluminum foils with either 0.5-inch or 1-inch holes punched through their centers. Placing these filters over a given PMT would reduce its light level to either 15% or 50% of nominal, therefore simulating varying flux levels. From bench test observations, the operation of the PMT's at the lowest light levels (which are still high enough to avoid the distorting effects of the Møller electronics, but low enough so that neither cathode nor anode effects contribute) is known to be very linear. Comparing the asymmetries measured by a given tube with and without filters therefore yields an indication of the tube's nonlinearity at high flux, since:

$$\begin{aligned} A_{meas}^{high} &= (1-\alpha)\cdot A_{meas}^{low} \\ &= \epsilon \cdot A_{meas}^{low} \end{aligned} \tag{4.13}$$

where A_{meas}^{high} and A_{meas}^{low} are the measured asymmetries for high and low flux levels, respectively, and α is the degree of nonlinearity. The iron foil is inserted so that the asymmetries are large enough to be measured quickly with high relative precision. At the gains at which the PMT's were operated for Runs I and II, this analysis reveals the detector linearity to be $\epsilon = 99 \pm 1$ %.

The procedure described above provides an accurate means of determining the linearity of the detector in response to small to medium-sized variations in flux (in other words, flux levels remaining at least 10% of nominal). This is the scale that is relevant for the actual Møller asymmetry measurement. However, for neutral background measurements, the spectrometer quadrupole magnets are turned off, for reasons that will be discussed below. In addition, thinner targets (or no target at all) are sometimes necessary to use. Both of these conditions result in a drastically reduced total flux at the detectors, usually 1% of nominal or less. For these special runs, it is important to have a good understanding of the detector's response to very small signal fluxes.

Towards this end, a series of runs was taken with various targets (including the normal LH_2 target, a 1-cm-thick carbon target, a 4-cm-thick carbon target, the 100 μm iron foil target, the wire array, and no target at all), thereby subjecting the PMT's to various levels of incident flux. In addition, the gains for each of the detector PMT's were varied, so that the response of the detector's electronics could be probed over a wide range of input levels. Analysis of this data was able to result in the desired linearity calibration down to small

signal levels, as described more thoroughly in Reference [98]. Using this important low-level linearity calibration, it is then possible to estimate the size of the neutral background flux as well as its contribution to the Møller asymmetry.

4.8.2 Neutral Background Corrections

There are three basic types of neutral backgrounds that might be affecting the measured Møller asymmetry at some level: photonic backgrounds, neutral hadronic backgrounds, and shower-spread effects in the calorimeter. The photonic background can be further divided into synchrotron, high energy, and low energy photons. Each of these background sources can be either suppressed or enhanced by changing various conditions of the experiment, such as the target and the collimator arrangement. Estimates of the backgrounds can be obtained by studying the detector's response to these changing conditions, as is described in the sections below.

Synchrotron Photons

Synchrotron radiation, particularly from the last magnet in the dipole chicane, is mostly blocked by collimators designed for this purpose. However, multibounce photons may still present a small background in the Møller detector. This background is easily measured by taking data with all targets removed from the beamline. These measurements indicate that the synchrotron flux is 0.15 ± 0.05 % of the Møller flux, and is relatively flat with respect to the azimuth. Therefore its contribution to the dipole asymmetries can be completely ignored. However, it can still contribute to the overall Møller asymmetry, if the synchrotron flux is found to possess an asymmetry itself.

An asymmetry A_{SR} in the synchrotron flux can only result from a vertical polarization component P_y in the electron beam. For $P_y = 100\%$, the asymmetry produced in the Møller detector is estimated to be 100 ± 50 ppm, assuming a cutoff in the detector's photon efficiency at $2-3$ MeV [85]. It therefore only remains to determine P_y, and the synchrotron flux's contribution to the Møller asymmetry will be known.

The beam's vertical polarization component P_y is estimated two different ways. One estimate comes from the synchrotron light monitor (SLM) analysis, coupled with a knowledge of the corrector magnet strengths in the A-line. Simulations show that for 100% vertical polarization, the asymmetry measured by the SLM should be roughly 60 ± 30 ppm, after

correcting for beam asymmetries, particularly in energy, to which it is very sensitive [85]. The actual asymmetry measured by the SLM is approximately 0.01 ± 0.04 ppm for Runs I and II combined, implying an average vertical polarization of $P_y = 0.02 \pm 0.3$ %. This only addresses what the vertical polarization is at the location of the SLM, which is at the beginning of the A-line. The polarization at the end of the A-line may actually be slightly different, due to the effects of imperfect vertical steering through the A-line.

Any net bend in the vertical direction will induce a small vertical polarization component, as the beam electrons' spins precess about the horizontal magnetic field lines steering the beam. To a good approximation, the induced vertical polarization can be written as a function of the net vertical bend angle θ_y as follows:

$$P_y \approx -100 \times P_{beam} \times \theta_y \tag{4.14}$$

where P_{beam} is the beam's longitudinal polarization (roughly 85%). Time histories of the strengths of the vertical corrector magnets from the SLM to the end of the A-line show that $\theta_y \approx 50$ μrad, leading towards an estimate for an additional induced vertical polarization of $\Delta P_y \approx -0.4\%$. Adding this to what was measured using the SLM, the vertical polarization component at the end of the A-line is found to be $P_y \approx -0.4 \pm 0.4$ %, where the uncertainty comes from both the SLM measurement as well as the corrector strength measurements.

This estimate for the vertical polarization component can be compared to what is obtained from the dipole asymmetry analysis. Dipole asymmetries arise from two main sources: improper beam corrections and two-photon exchange processes. Assuming the latter dominates (which for the Møller detector is a fair assumption, though perhaps not for the "out" ring or the luminosity monitor, since these detectors are so much more sensitive to beam motions, particularly higher-order beam effects), one can use a calibration of this effect at some known transverse polarization to deduce the transverse polarization components at a later time, given a set of measured dipole asymmetries. This method is further explained in Section 4.8.5. A calibration reveals that 85% transverse (horizontal) beam polarization produces a Y dipole asymmetry of -3.0 ± 0.3 ppm. From Figure 4.7, one can see that the measured X dipole asymmetry for Runs I and II was -31 ± 17 ppb (combining the "in" and "mid" rings' results), implying an average vertical polarization of $P_y = -0.9 \pm 0.5$ %, which is consistent with the result obtained above.

Using $P_y = -1 \pm 1$ % as a conservative estimate for the average vertical polarization component, the synchrotron flux's contribution to the Møller asymmetry can be written as:

$$\begin{aligned} \Delta A_{SR} &= P_y \cdot f_{SR} \cdot A_{SR} && (4.15) \\ &= (-0.01 \pm 0.01) \times (0.0015 \pm 0.0005) \times (100 \pm 50 \text{ ppm}) \\ &= -2 \pm 1 \text{ ppb} && (4.16) \end{aligned}$$

Neutral Hadrons

High-energy *ep* scatters interacting in the calorimeter produce neutrons and other neutral hadrons that can penetrate the shielding encasing the PMT's, generating significant responses from the tubes' cathodes. Since the *ep* flux possesses an asymmetry that is different from that of the Møller flux, this neutral hadron background flux will dilute and shift the measured Møller asymmetry. At the end of Run I, special runs were conducted in which the detector's light guides were all covered at their ends with aluminum tape. The tubes were then referred to as "blinded." Under these conditions, only neutral hadrons capable of penetrating to the phototubes' cathodes could contribute to the detector's signal, which fell to 0.6 ± 0.4 % of its normal value. Inserting the special CM8 collimator, which suppresses the *ep* flux in the calorimeter by more than an order of magnitude, reduces the sizes of the blinded PMT signals by a factor of two. Therefore half of the neutral hadron flux hitting the PMT's is presumed to be produced by the *ep* scatters. The asymmetry of the *ep* flux was measured during Run I to be approximately -1.5 ppm. Consequently, the neutral hadron background will cause an asymmetry shift in the Møller detector equal to:

$$\Delta A_{n.h.} = (0.003 \pm 0.002) \times (-1.5 \text{ ppm}) = -5 \pm 3 \text{ ppb} \qquad (4.17)$$

Note that the neutral hadrons produced by the Møller electrons in the calorimeter will neither dilute nor shift the measured asymmetry, since they will carry the same asymmetry as that of the Møller flux. Also note that this neutral hadron background flux was completely eliminated for Run II, as the CM8 collimator reduced the total *ep* flux in the calorimeter by more than an order of magnitude.

Shower Spreading

Initially there was concern that, to some extent, electromagnetic showers produced by charged particle interactions in the *ep* detector might be leaking into the Møller detector's fibers, specifically those of the "out" ring. By turning the quadrupole magnets off, thereby more than doubling the flux of particles hitting the *ep* detector, and monitoring the "out" ring's signal size as the special CM8 collimator was inserted and removed (bearing in mind that this collimator will also suppress the neutral hadron background, which has been measured independently using the blinded tube data), shower-spread effects from the *ep* detector were found to be completely negligible. In fact, the "out" ring is not even being used for the Møller asymmetry measurement. Therefore, even if this background did contribute to the "out" ring's signal, it would not affect the Møller result.

Hard and Soft Photons

The last background to consider is that of the non-synchrotron photons. These photons can be broadly divided into two classes: high-energy ("hard") and low-energy ("soft") photons. Hard photons can be produced by bremsstrahlung processes in the target. The chances of such photons hitting the detector are greatly reduced by the photon collimators. Nevertheless, a small background may persist due to multiple scattering off collimator edges. Photons can also be produced when charged particles (pions or electrons) hit the beampipe inside the quadrupole magnets and the inside edges of the acceptance collimator. These photons may be hard or soft, depending on the energy of the charged particle. The special collimators CM1 – 7 were designed to eliminate the soft photon background produced by beampipe and collimator "splash."

In general, sensitivity to photonic backgrounds, both hard and soft, is enhanced by turning the spectrometer quadrupole magnets off. This "sweeps" the charged particle flux out of the "in," "mid," and "out" rings and almost completely into the *ep* detector (see Figure 2.27). In this configuration, any signal observed by the Møller detector could be coming either from photons (including synchrotron, hard, and soft photons) or neutral hadrons. The synchrotron and neutral hadron backgrounds have already been measured. It therefore remains to separate out the hard photonic background from the soft. This is done by comparing "quads off" data taken before and after the collimators CM1 – 7 were

Run Conditions	*Synch. γ*	*Neutral Hadrons*	*Shower Spread*	*Hard γ*	*Soft γ*
No target	$\surd$	$\times$	$\times$	$\times$	$\times$
Blinded tubes	$\times$	$\surd$	$\times$	$\times$	$\times$
CM1 $-$ 7 = out, CM8 = out	$\surd$	$\surd$	$\surd$	$\surd$	$\surd$
CM1 $-$ 7 = in, CM8 = out	$\surd$	$\surd$	$\surd$	$\surd$	$\times$
CM1 $-$ 7 = in, CM8 = in	$\surd$	$\times$	$\times$	$\surd$	$\times$
f_{bkg} (%)	0.2 ± 0.1	0.1 ± 0.1	—	0.4 ± 0.2	—
ΔA_{bkg} (ppb)	0 ± 3	-2 ± 1	—	3 ± 3	—

Table 4.3: Summary of neutral background corrections for Runs I and II. All runs have the quadrupole magnets turned off.

installed. After subtracting out the synchrotron and neutral hadronic contributions, one is left with the separate hard and soft photonic backgrounds.

By studying the "quads off" data taken with and without the special collimators in place, it was concluded that the hard photon flux is 0.4 ± 0.2 % that of the Møller flux, and that the soft photon flux is negligible. In addition, even with collimators CM1 $-$ 7 and CM8 installed, an asymmetry of roughly 2.5 ppm is observed in the Møller detector. These large-asymmetry photons are produced by either pions or electrons colliding with the beampipe and collimator edges. Part of the hard photon flux is undoubtedly due to bremsstrahlung processes in the target, and these photons are expected to carry a very small asymmetry, since they are created by very low Q^2 processes. To be conservative, however, the entire hard photon flux is assumed to be due to beampipe and collimator scattering, and to carry therefore the large asymmetry. Turning the quadrupole magnets off increases the amount of charged particles hitting the beampipe by a factor of three. Therefore, the correction to the Møller asymmetry due to the hard photon flux is estimated to be:

$$\Delta A_{\gamma_{\rm hard}} = (0.004 \pm 0.002) \times (1/3) \times 2.5 \text{ ppm} = 3 \pm 3 \text{ ppb} \tag{4.18}$$

where the uncertainty has been increased to account for the uncertainty associated with the exact source of the background.

Table 4.3 summarizes all of the neutral background information. The table displays the dilution factors f_{bkg} for each source, as well as each source's contribution ΔA_{bkg} to the

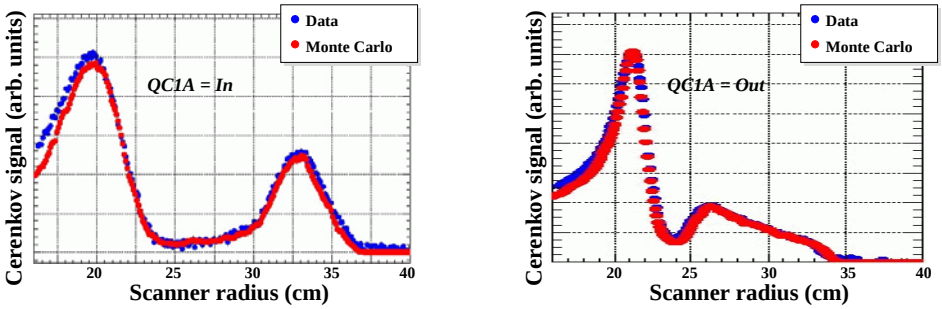

(a) (b)

Figure 4.18: Examples of flux profile scans used to calibrate the experiment's Monte Carlo. Figure (a) shows an example with the calibration collimator QC1A inserted. Under these conditions, the Møller and *ep* fluxes become widely separated, allowing for an accurate determination of the *ep* flux. The calibrated Monte Carlo simulation can then be compared to a second scan with QC1A removed, an example of which is shown in Figure (b).

Møller asymmetry measurement. In addition, the table shows which run conditions enhance or suppress each background's effect on the detector, thereby indicating the means by which each background can be measured.

4.8.3 ep Correction

The largest background in the Møller detector, both in terms of its rate as well as its overall effect on the measured asymmetry, is the *ep* background. It consists of both elastic (Mott) scatters as well as inelastic *ep* scatters. To calculate the *ep* correction to the Møller asymmetry, a detailed Monte Carlo simulation [106] of the entire experiment (from the target to the detector) was calibrated with the data obtained by the profile monitor.[5] These data were taken with the calibration collimator (QC1A) inserted. With this collimator in place, all particles must pass through very small radial apertures (typical openings are $\Delta r \approx 1$ cm, and can be seen in Figure 2.28(a)), resulting in a clean separation of the Møller and *ep* peaks at the detector. Input parameters to the Monte Carlo (e.g., spectrometer optics, collimator positions, etc.) were tuned until the simulated profiles agreed with the actual radial and azimuthal flux maps, as shown in Figure 4.18.

The flux maps reveal that approximately 7% of the Møller detector's signal is due to the *ep* background. Most of this background (about 6% of the total 7%) consists of radiative Mott's, that is, elastic *ep* scattering events that include initial and/or final state radiation.

[5] This custom-built Monte Carlo was checked by a completely independent GEANT-based simulation [107]. The two codes produced results in good agreement.

	Run I		*Run II*	
Energy	f_{ep}	ΔA_{ep}	f_{ep}	ΔA_{ep}
45 GeV	7.8 ± 0.9 %	-32.6 ± 5.0 ppb	6.4 ± 0.8 %	-27.5 ± 4.1 ppb
48 GeV	7.0 ± 0.8 %	-34.4 ± 5.4 ppb	6.1 ± 0.6 %	-29.5 ± 4.1 ppb

Table 4.4: Summary of the total *ep* contributions to the Møller asymmetry for Runs I and II.

Without this "radiative tail", in fact, the Mott contribution to the Møller flux would only be roughly 1%, equal to that of the inelastic *ep*'s. Though fewer in numbers, the inelastic *ep* scatters, because of the large asymmetry they carry, provide a larger contribution to the Møller asymmetry than do the radiative Mott's.

Once the *ep* flux distributions are known, the contributions to the Møller asymmetry can be calculated using the Monte Carlo simulation. The main complication arises from the fact that the inelastic *ep* asymmetry is not entirely well known theoretically. On the one hand, above the so-called "resonance region", corresponding to $W > 3$ GeV, where W is the total four-momentum of the recoiling system, the asymmetry is easy to calculate, using the following formula [108, 109]:

$$A_{ep,inel}(Q^2) = (0.5 \pm 0.5) \times 10^{-4}\ \text{GeV}^{-2} \cdot Q^2 \qquad (\text{for } W > 3\ \text{GeV}) \tag{4.19}$$

This formula, valid for the experimental kinematics ($Q^2 < 0.04$ GeV^2), has a large uncertainty associated with it, but this uncertainty only affects the correction to the Møller asymmetry at the few ppb level. On the other hand, in the resonance region ($M_p < W < 3$ GeV), the value of the inelastic *ep* asymmetry for low Q^2 is not well known. Its value must be derived from the *ep* detector's asymmetry measured in Run I, which is $-1.433 \pm 0.045\,(\text{stat})$ ppm for 45 GeV running and $-1.736 \pm 0.063\,(\text{stat})$ ppm for 48 GeV running. The reason there are two asymmetries is that the kinematics for the *ep* scatters (and thus the average Q^2 of those events) is very different depending on the beam energy. Thus, as will be seen, the *ep* correction to the Møller asymmetry is energy-dependent.

The total *ep* correction to the Møller asymmetry can be written as the sum of three terms [109]:

$$\Delta A_{ep} = \Delta A_{ep,elast} + \Delta A_{ep,inel,W>3} + \Delta A_{ep,inel,W<3} \tag{4.20}$$

Here $\Delta A_{ep,elast}$ is the asymmetry contribution of the elastic *ep* (Mott) scatters; $\Delta A_{ep,inel,W>3}$ is the asymmetry contribution of the inelastic *ep* scatters, for $W > 3$ GeV; and $\Delta A_{ep,inel,W<3}$

is the asymmetry contribution of the inelastic ep scatters, for $W < 3$ GeV. All but the last term ($\Delta A_{ep,inel,W<3}$) are computed entirely by the Monte Carlo simulation. The last term is extracted from the ep detector's asymmetry measurement. The total ep contribution to the Møller asymmetry, averaged appropriately over 45 GeV and 48 GeV data for both Runs I and II, is $\Delta A_{ep} = -31 \pm 5$ ppb, where the largest uncertainty comes from the uncertainty in the constant of proportionality in Equation (4.19). The results are presented in more detail in Table 4.4.

4.8.4 Pion Correction

Like the ep background, the pion background carries a large asymmetry (on the order of one ppm). However, its flux is much smaller, so its overall effect on the Møller asymmetry will not be as large. In order to determine the pion flux's contribution ΔA_π to the Møller asymmetry, as well as its dilution factor f_π, three pieces of information are required. The first is the actual asymmetry A_π carried by the pion flux. This quantity is extracted from the pion detector's asymmetry measurement, after correcting it for the dilution presented by the electron "background". The amount of electron contamination in the pion detector is measured by inducing a large asymmetry in the electron signal by inserting the polarized 100 μm iron foil into the beamline and comparing the asymmetry of the pion detector to that of the Møller. This analysis reveals that roughly 14% of the pion detector's signal comes from the electron flux. The corrected pion asymmetry is then calculated to be $A_\pi = -1.02 \pm 0.33\,(\text{stat})$ ppm.

The second piece of information needed to calculate the pion background's contribution to the Møller asymmetry is the relative pion rate n_π. The flux dilution factor f_π is related to the ratio ϵ_π of the pion detector's mean Čerenkov response for pions and electrons by the following equation: $f_\pi = \epsilon_\pi \cdot n_\pi$. The relative pion rate n_π can actually be estimated from the knowledge of the width of the pion detector's asymmetry distribution, assuming the detector's energy resolution and detection efficiency are known from simulation [94]. A typical channel's asymmetry width is 1.35%, leading to an estimate of 0.62% as the relative pion rate in the Møller detector [94].

This estimate can be compared to a more direct measurement using the concept of relative attentuation [94]. The signal level in a pion counter is monitored as increasing amounts of shielding are placed in front of it. A GEANT simulation is used to calculate,

given a certain number of electrons and pions at the start, how many will remain after a given amount of shielding, and what the response of the pion counter will be. The measured attenuation of the pion counter's response as a function of shielding length can then be used to reveal the precise mix of electrons and pions there must have been in front of the shielding. By this method, the relative number of pions to Møller electrons in front of the Møller detector is found to be $n_\pi = 0.63 \pm 0.21\%$.

Finally, the ratio ϵ_π of the pion detector's mean Čerenkov response for pions and electrons must be simulated. Simulations show that $\epsilon_\pi = 22 \pm 15$ %. This large uncertainty dominates the final uncertainty on the pion flux measurement and stems from the fact that the mean Čerenkov response of the pion detector to pions is quite sensitive to the pions' average energy, which is only known at the ± 15 GeV level. Using these three pieces of information, the dilution factor and asymmetry contribution of the pion flux to the Møller asymmetry can be calculated for Runs I and II as:

$$\begin{aligned} f_\pi &= \epsilon_\pi \cdot n_\pi = (22 \pm 15\ \%) \times (0.63 \pm 0.21\ \%) \\ &\approx 0.1 \pm 0.1\ \% \\ \Delta A_\pi &= f_\pi \cdot A_\pi = (0.001 \pm 0.001) \times (-1.02 \pm 0.33\ \text{ppm}) \\ &\approx 1 \pm 1\ \text{ppb} \end{aligned} \tag{4.21}$$

4.8.5 Transverse Asymmetry Correction

The channel-by-channel asymmetries for the "in" and "mid" rings, averaged over all of Runs I and II, are shown in Figure 4.19. The asymmetries exhibit clear sinusoidal variations. In other words, there are non-zero dipole asymmetries: combining the results of the two rings together, $A_x = -31 \pm 17$ ppb and $A_y = 112 \pm 18$ ppb. These non-zero dipole asymmetries

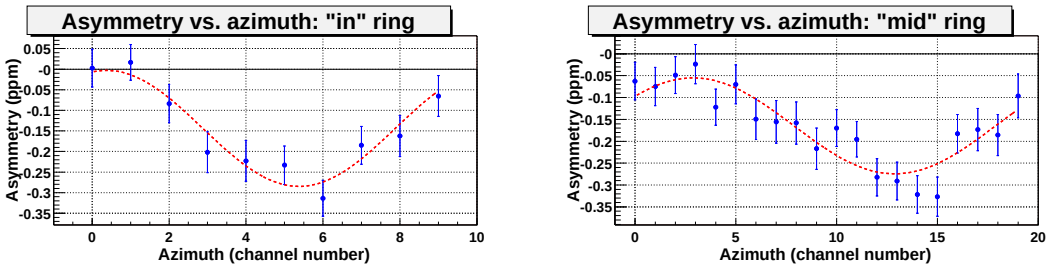

Figure 4.19: Channel-by-channel asymmetries in the "in" and "mid" rings for Runs I and II. For the "in" ring, the second channel (in01) basically corresponds to the vertical direction, whereas for the "mid" ring, the fourth channel (mid03) does. In either case, the primary component of the azimuthal variation in the asymmetry can be interpreted as a vertical dipole asymmetry.

can arise via two distinct mechanisms, improper beam corrections and two-photon exchange processes. The latter mechanism dominates in the case of the "in" and "mid" rings.

Transverse asymmetries in Møller scattering can arise from purely QED processes, namely those described by one-loop photon diagrams. Interference between the imaginary part of the two-photon exchange amplitude and the real part of the one-photon exchange amplitude results is an asymmetry that is directly proportional to the sine of the azimuthal angle ϕ [110]:

$$A(\phi) \;\propto\; \mathbf{s}\cdot(\mathbf{p}\times\mathbf{p}') \;\propto\; \sin(\phi) \tag{4.22}$$

where $\mathbf{s}$ is the spin of the incoming electron; $\mathbf{p}$ and $\mathbf{p}'$ are the four-momenta of the incoming and scattered electron, respectively; and ϕ is the azimuthal angle of the scattered electron around the beam direction, measured from the direction of $\mathbf{s}$. Normally $\mathbf{s} \parallel \mathbf{p}$, so that no dipole asymmetries (as measured by the experiment) result. However, if $\mathbf{s}$ acquires transverse polarization components (i.e., P_x or P_y), dipole asymmetries will become manifest. Specifically, P_x can be seen via Equation (4.22) to produce a vertical dipole asymmetry, while P_y will produce a horizontal dipole asymmetry.

Using the monopole weights, as is done when computing the actual Møller asymmetry, obviously suppresses these dipole effects. Given an infinite amount of running time, the transverse asymmetries will tend to contribute nothing to the Møller asymmetry. However, given a finite amount of running time, statistical fluctuations can create a measurable bias to the Møller asymmetry. This bias is easily computed, given the average channel weights w_i and X and Y dipole asymmetries, A_x and A_y:

$$\text{X dipole bias} \;=\; A_x \sum_i w_i \cdot \sin\left(\phi_i - \phi_{top}\right) \tag{4.23}$$

$$\text{Y dipole bias} \;=\; A_y \sum_i w_i \cdot \cos\left(\phi_i - \phi_{top}\right) \tag{4.24}$$

Here the sums run over both the "in" and "mid" rings; ϕ_i is the phase of channel i, for a given ring; and ϕ_{top} is the phase of the top of the ring. Applying these formulae, one arrives at a total contribution of -6 ± 3 ppb to the Møller asymmetry for both Runs I and II.

As has been mentioned previously, the transverse asymmetry can be a useful method for determining transverse polarization components. During Runs I and II, a portion of the running time was dedicated to running "off-energy" at 43 and 46 GeV. At these en-

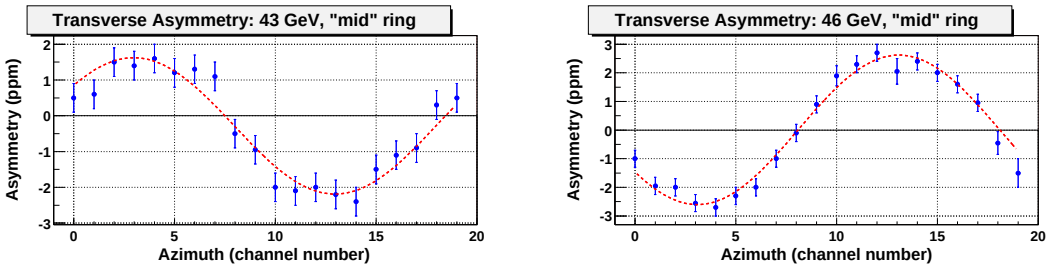

Figure 4.20: Channel-by-channel asymmetries in the "mid" ring for 43 and 46 GeV data. The azimuthal variations are evidence for a transverse asymmetry caused by two-photon exchange processes. Since the horizontal beam polarization changes sign going from 43 GeV to 46 GeV, the asymmetry is also expected to change sign, as is observed.

ergies, because of the $g-2$ spin precession in the A-line, the beam's polarization should be 85% horizontal instead of longitudinal (43 GeV and 46 GeV correspond to opposite horizontal polarizations). Measuring the azimuthal asymmetry variation at these energies therefore provides a clean calibration of the size of the asymmetry dipoles that the two-photon exchange processes can produce. The results of the 43 and 46 GeV data are shown in Figure 4.20. The clear reversal of the behavior with respect to energy, as predicted by the physics, is evident. The data indicate a total dipole asymmetry of -3.0 ± 0.3 ppm (divided between a $\sim$1.5 ppm X dipole asymmetry and a $\sim$2.5 ppm Y dipole asymmetry), which agrees with a calculation that has been made for the experimental conditions [110]. Though it is impossible to induce 85% vertical polarization, the effect would be entirely analogous. Any dipole asymmetry measurements from the actual physics dataset can therefore be turned into transverse polarization measurements, using the following formulae:

$$P_x = 33.3 \pm 3.3 \ \%/\text{ppm} \times A_y \tag{4.25}$$

$$P_y = 33.3 \pm 3.3 \ \%/\text{ppm} \times A_x \tag{4.26}$$

This is precisely the method used to provide one estimate for the vertical polarization component P_y in Section 4.8.2, which was necessary in order to calculate the synchrotron flux's contribution to the Møller asymmetry.

4.8.6 Final Asymmetry Calculation

All of the corrections and backgrounds discussed above are summarized in Table 4.5. Results for both Runs I and II have been combined, using appropriate weighting. For instance, the

Correction or Background	*Runs I and II*	
	f (%)	ΔA (ppb)
Beam corrections (1^{st} order)	—	-29 ± 2
Beam corrections (higher order)	—	0 ± 10
Beam spotsize	—	0 ± 1
Transverse pol'n	—	-6 ± 3
ep (elastic + inelastic)	6.8 ± 0.7	-31 ± 5
High energy photons	0.4 ± 0.2	3 ± 3
Synchrotron photons	0.2 ± 0.1	-2 ± 1
Neutral hadrons	0.1 ± 0.1	-2 ± 2
Pions	0.1 ± 0.1	1 ± 1
TOTAL	7.6 ± 0.7	-66 ± 12

Table 4.5: Summary of all corrections and background contributions to the measured asymmetry for Runs I and II. It is important to remember that the +29 ppb correction to the Møller asymmetry (to account for the effects of beam asymmetries) has already been taken into account by the regression procedure.

neutral hadron contribution is different for Runs I and II because of the installation of the additional collimator CM8 before Run II. The two runs' flux measurements f_1 and f_2 should therefore be combined according to each run's statistical weight:

$$f_{tot} = \frac{f_1 \cdot w_1 + f_2 \cdot w_2}{w_1 + w_2} \tag{4.27}$$

where w_1 and w_2 are the statistical weights for Runs I and II, equal to $1/\delta(A_1)^2$ and $1/\delta(A_2)^2$, respectively, where $\delta(A_1)$ and $\delta(A_2)$ are the statistical uncertainties on the two runs' asymmetry measurements. The same type of statistical-weighted average is then taken to find the uncertainty on the combined flux measurement:

$$\delta(f_{tot}) = \frac{\delta(f_1) \cdot w_1 + \delta(f_2) \cdot w_2}{w_1 + w_2} \tag{4.28}$$

The same method can be applied to the asymmetry contributions ΔA. For example, in Run I the neutral hadron contribution is -5 ± 3 ppb, while in Run II it is zero. The statistical uncertainties of the Møller asymmetry measurements for Runs I and II are 23 and 22 ppb, respectively. Using these numbers in the above equations yields -2 ± 2 ppb as the average neutral hadron asymmetry contribution. It should be noted that this treatment is exactly equivalent to computing a separate asymmetry for each run individually, and then combining the results according to their statistical weight. It should also be noted

Normalization	*Runs I and II* n (%)
Beam polarization	84 ± 5
Linearity	99 ± 1
TOTAL	83 ± 5

Table 4.6: Summary of normalization factors for Runs I and II.

that averaging the systematic uncertainties in this way assumes that they are not statistics limited, which in some cases may be a conservative approach.

For the *ep* background, further averaging is necessary. This is because, as Table 4.4 shows, the *ep* flux and its asymmetry are different for the two beam energies (due to the slightly different kinematics at either energy). Thus, to find the average *ep* contribution (or the average *ep* background flux) for Runs I and II combined, one must average over each run's energy subsets separately, in a manner entirely analogous to Equation (4.27) above, weighting by the statistical uncertainty of each subset's asymmetry measurement.

Aside from corrections and backgrounds, normalization factors (i.e., beam polarization and linearity) must be taken into account. These results are summarized in Table 4.6. Using all of these corrections, dilutions, and normalizations in Equation (4.12), the final result for the parity-violating asymmetry in Møller scattering at $Q^2 = 0.026$ GeV^2 is:

$$\text{Runs I and II:} \quad A_{\text{PV}}(Q^2 = 0.026 \text{ GeV}^2) = -160 \pm 21\,(\text{stat}) \pm 16\,(\text{syst}) \text{ ppb}$$

Of the total 16 ppb systematic uncertainty, 12 ppb comes from the asymmetry corrections (mainly the uncertainty due to higher-order beam effects), 10 ppb from the normalizations (mainly the uncertainty in the beam polarization measurements), and 1 ppb from the dilutions (mainly from the *ep* background).

Figure 4.21 shows the parity-violating asymmetry A_{PV} for every slug of data for both Runs I and II. Only the statistical error bars have been plotted. The true physics sign of each slug (i.e., the factor of -1 to account for the status of the half-wave plate as well as for the beam energy) has been suppressed, so that the asymmetry can be seen to alternate stably over time as the half-wave plate is inserted and removed and as the beam energy is switched from 45 to 48 GeV. Such a plot reinforces the fact that the asymmetry observed in the detector is indeed due to a helicity-dependent term in the Møller scattering cross section, and not due to, for instance, electronics cross talk. Such a plot is convincing because changing the beam energy from 45 to 48 GeV, or inserting a half-wave plate in the

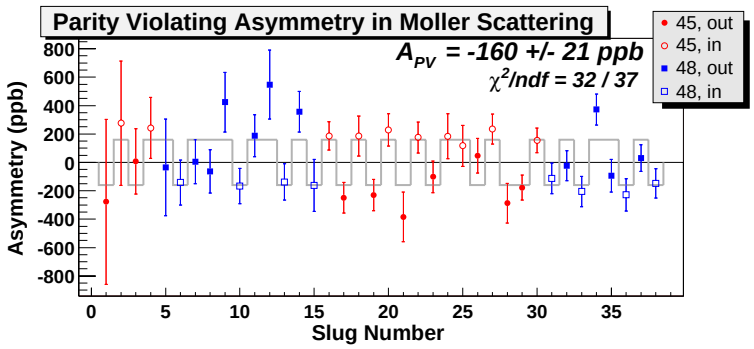

Figure 4.21: Parity-violating cross section asymmetry A_{PV} in Møller scattering, plotted versus slug number, for Runs I and II.

path of the laser beam at the source, manages to change, in a very passive way, the beam helicity at the target, while affecting as little else in the experiment as possible.

When computing the significance of this first observation of parity violation in Møller scattering, the true measured asymmetry, that is, the asymmetry given by Equation (4.11), which takes into account all corrections and dilutions, but not any normalizations, must be compared to its total uncertainty. The true measured asymmetry is -133 ppb, while its uncertainty contains the following contributions: 17.2 ppb from statistics, 12.4 ppb from corrections, and 1.0 ppb from dilutions. The significance of the observation of parity violation in Møller scattering is therefore:

$$\text{Significance of parity violation observation} = \frac{133}{\sqrt{17.2^2 + 12.4^2 + 1.0^2}} = 6.3\,\sigma$$

Normalizations (i.e., beam polarization and linearity) are necessary in order to extract the true value of the underlying physics asymmetry, which may then be interpreted in terms of Standard Model parameters (e.g., the weak mixing angle). This facilitates comparisons to other experiments and theoretical models, as will be discussed in the next chapter. However, for merely setting a limit on the signficance of parity violation in the process of Møller scattering, the actual value of the underlying physics asymmetry is of no consequence; only the deviation of the measured asymmetry from zero matters.

Chapter 5

Conclusions

5.1 Calculation of the Weak Mixing Angle

The preceding chapter gave the final result for the value of the parity-violating asymmetry $A_{\rm PV}$ in Møller scattering, averaged over the specific experimental kinematics. This asymmetry can be related to a particular value of $\sin^2\theta_W$ via the following equation [29]:

$$\begin{aligned} A_{\rm PV} &= -\frac{2\sqrt{2}G_F Q^2}{\pi\alpha}\cdot\mathcal{F}_{brem}\cdot\frac{1+\cos\Theta}{(3+\cos^2\Theta)^2}\cdot\left(1-4\sin^2\theta_W(Q^2)_{\overline{\rm MS}}\right) \\ &= \mathcal{AP}\cdot\mathcal{F}_{brem}\cdot\left(1-4\sin^2\theta_W(Q^2)_{\overline{\rm MS}}\right) \end{aligned} \tag{5.1}$$

The factor $\mathcal{F}_{brem}$ takes into account initial and final state bremsstrahlung radiation effects [111], whose Feynman diagrams are depicted in Figure 5.1. These effects are similar to the electroweak radiative corrections discussed in Section 1.2.1, but must be calculated for the specific experimental spectrometer acceptance values, $13 < E' < 24$ GeV; therefore they have not been mentioned until now. Likewise, the effective analyzing power $\mathcal{AP}$ absorbs the asymmetry's dependence on Q^2 and Θ and depends on the precise experimental

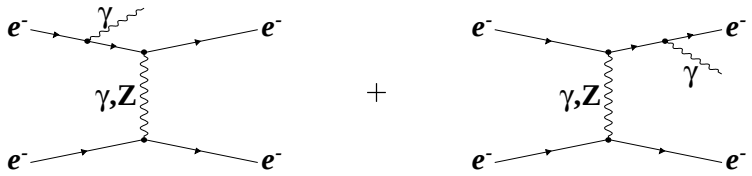

(a) (b)

+ Inverted + crossed diagrams

Figure 5.1: Bremsstrahlung radiation diagrams modifying the tree level expression for $A_{\rm PV}$. The effects of such diagrams must be carefully evaluated for the given experimental kinematics. The diagrams factorize and largely cancel in the asymmetry ratio [29], which is enhanced by 1% [111].

$\langle \mathcal{F}_{brem} \rangle$	$\langle \mathcal{AP} \rangle$	$\langle Q^2 \rangle$	$\langle \cos\Theta \rangle$	$\langle \theta_{lab} \rangle$
1.01 ± 0.01	-3.34 ± 0.02 ppm	0.026 GeV^2	-0.2	5.7 mrad

Table 5.1: Average kinematic values for Runs I and II. Only the errors on $\mathcal{F}_{brem}$ and $\mathcal{AP}$ influence $\sin^2\theta_W$. The quantity $\mathcal{F}_{brem}$ is obtained from calculating the effects of the diagrams in Figure 5.1, while $\mathcal{AP}$ is obtained from a detailed Monte Carlo simulation of the experiment.

kinematics and geometry. It is calculated using the same detailed Monte Carlo simulation that is used to determine the *ep* correction to the Møller asymmetry [107].

Since the kinematics of the scattered electrons reaching the detector's acceptance can change depending on the beam energy, in general both $\mathcal{F}_{brem}$ and $\mathcal{AP}$ will be different for the two energy states. The average values for these parameters, along with the average values for Q^2, $\cos\Theta$ and θ_{lab} to which they correspond, are given in Table 5.1. One can use the values from this table to find the average value for the weak mixing angle:

$$\sin^2\theta_W(Q^2 = 0.026 \text{ GeV}^2)_{\overline{\text{MS}}} = \frac{1}{4} \times \left(1 - \frac{\langle A_{\text{PV}} \rangle}{\langle \mathcal{AP} \rangle \cdot \langle \mathcal{F}_{brem} \rangle}\right) \quad (5.2)$$

$$= 0.2381 \pm 0.0015 \text{ (stat)} \pm 0.0014 \text{ (syst)} \quad (5.3)$$

This low-energy weak mixing angle measurement is shown in Figure 5.2(a), which also shows the results of the Z-pole measurements as well as those of the NuTeV experiment and the Boulder atomic parity violation group.

In order to more directly compare the results of the various experiments to one another, all measurements are evolved to the Z-pole. The contributions of electroweak radiative corrections, which modify $\sin^2\theta_W$ for $Q \neq m_Z$, must therefore be removed. The full set of electroweak radiative corrections relevant for Møller scattering at $Q^2 \approx 0.03$ GeV^2 has been calculated in Reference [29]. The dominant corrections, produced by the γ-Z mixing diagrams as well as the W loop contribution to the anapole moment, have been discussed in Section 1.2.1. The net effect of all one-loop radiative corrections is to make the following substitution for $\sin^2\theta_W$ [29]:

$$\sin^2\theta_W(Q^2 = 0.026 \text{ GeV}^2) = (1.0301 \pm 0.0025) \times \sin^2\theta_W(m_Z) + 0.0004 \pm 0.0003 \quad (5.4)$$

The uncertainty of ± 0.0025 comes from the interpretation of the $e^+e^- \rightarrow$ hadrons data that is necessary in order to properly handle the vacuum polarization loops involving quarks. The uncertainty of ± 0.0003 comes from the box diagrams involving one photon and one

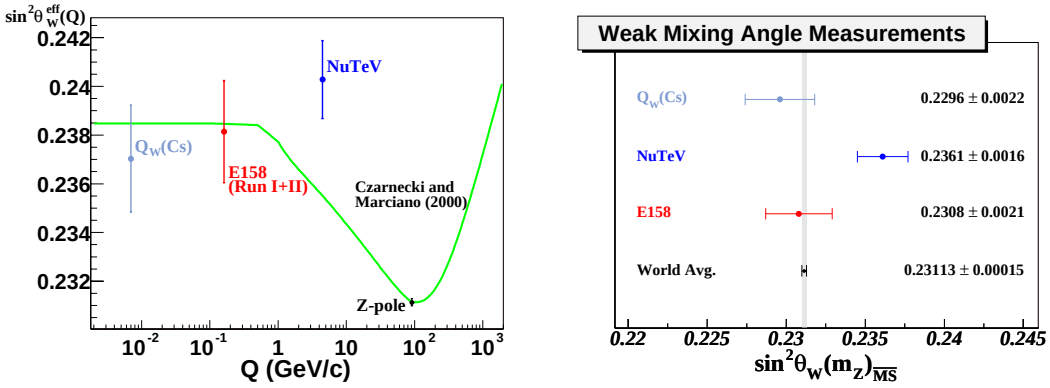

Figure 5.2: Run I and II results for $\sin^2\theta_W$. Figure (a) shows the approximate evolution of $\sin^2\theta_W$ as a function of energy, superimposing its precise measurements at different energy scales. Figure (b) evolves all measurements to the Z-pole for a more direct comparison.

heavy boson, as well as the Z-loop contribution to the anapole moment. In order to evaluate these diagrams' contributions, a value for $\sin^2\theta_W$ must be chosen. Instead of carrying out a proper analysis in order to determine which value ($\sin^2\theta_W(m_Z)$ or $\sin^2\theta_W(0)$) to use, the average value is chosen. The spread is then used to estimate a theoretical uncertainty. Use of Equation (5.4) thus yields the following value for the weak mixing angle:

$$\sin^2\theta_W(m_Z)_{\overline{\rm MS}} = 0.2308 \pm 0.0015\ (\text{stat}) \pm 0.0015\ (\text{syst}) \tag{5.5}$$

Of the ± 0.0015 systematic uncertainty, ± 0.0014 comes from experimental systematics (raw asymmetry corrections, beam polarization, analayzing power, etc.), while ± 0.0006 comes from the theoretical uncertainties in the electroweak radiative correction calculations discussed above. This value for $\sin^2\theta_W(m_Z)$ is compared to those of the other measurements in Figure 5.2(b). The central values and error bars are current as of the time of this writing. However, the central values and/or error bars of both the NuTeV and Q_W(Cs) results may change in the future, due to the reasons discussed in Section 1.2.2. The "World Avg." data point reflects the average of all the world's electroweak data (current as of July 2002), including the "direct" Z-pole measurements at SLAC and CERN [18].

Experiment	$X \pm \delta(X)$
World average (1996), Ref. [36]	0.13 ± 0.51
Q_W(Cs), Ref. [57]	0.22 ± 0.29
NuTeV, Ref. [44]	-0.60 ± 0.21
This measurement	$\mathbf{0.05 \pm 0.27}$
New world average	-0.19 ± 0.14

Table 5.2: Limits on the extended oblique corrections parameter X. Off-shell measurements of the weak mixing angle are particularly sensitive to this parameter, as it is directly proportional to the "running" of $\sin^2\theta_W$.

5.2 Physics Implications

From Figure 5.2, it is clear that no significant discrepancy from the Standard Model prediction has been observed in this measurement of the parity-violating Møller scattering asymmetry. This result can therefore tighten the constraints on many "new physics" models, including those involving new heavy fermions, new massive gauge bosons, and lepton compositeness. The impact on each of these constraints will now be discussed.

5.2.1 Oblique Corrections

Current electroweak data, primarily from the e^+e^- colliders, can tightly constrain the Peskin–Takeuchi parameters S, T, and U used to parameterize oblique corrections from new particles with $m_{new} \gg m_Z$ [35]. However, if the mass constraint is relaxed, so that $m_{new} \approx \mathcal{O}(100)$ GeV, the "extended" oblique corrections parameters V, W, and X need to be introduced [37]. Low-energy measurements of the weak mixing angle are particularly sensitive to the X parameter, since its value can be related directly to the running of the weak mixing angle [29]:

$$\frac{\sin^2\theta_W(0)}{\sin^2\theta_W(m_Z)} \simeq 1.032 - 0.033X \tag{5.6}$$

This equation merely states that, according to the Standard Model, $\sin^2\theta_W(0)$ should be roughly 3% larger than $\sin^2\theta_W(m_Z)$ (here $\sin^2\theta_W(m_Z)$ refers to the world average value of 0.23113 ± 0.00015). Any deviation from this behavior could signal the influence of new oblique corrections (modifying the $\Pi_{\gamma Z}$ vacuum polarization function) whose relevant energy scale is roughly equal to that of electroweak symmetry-breaking. The $Q \approx 0$ measurement of the weak mixing angle reported on here finds that $\sin^2\theta_W$ rises by 3.0 ± 0.9 %, corresponding to the constraint $X = 0.05 \pm 0.27$. This limit can be compared to those obtained from other

electroweak measurements in Table 5.2.

5.2.2 New Neutral Gauge Bosons

New heavy neutral gauge bosons, collectively referred to as Z' bosons, are a fairly ubiquitous feature of many extensions to the Standard Model currently being pursued, including grand unified theories (GUT's), supersymmetric theories, models with extra dimensions, and superstring theories. In general, each theory specifies its own definitions for the couplings of the new Z' boson to the various other particles of the Standard Model, resulting in very model-dependent predictions for how a given observable will be affected. It is therefore hard to discuss limits on $m_{Z'}$ in a general, model-independent fashion. However, the implications for a "typical" model, namely an E_6-inspired GUT [112], will be discussed below.

The additional U(1) gauge symmetries contained in E_6 models require two additional neutral gauge bosons, Z_χ and Z_ψ. In general these additional bosons mix with one another as well as with the Z. However, from the Z-pole measurements constraining the Peskin-Takeuchi parameters S and T, the degree of mixing with the Z is known to be very small. The angle β specifies the degree to which the Z_χ and Z_ψ mix with each other in order to form the following mass eigenstates [34]:

$$\begin{aligned} Z_\beta &= Z_\chi \cos\beta + Z_\psi \sin\beta \\ Z'_\beta &= -Z_\chi \sin\beta + Z_\psi \cos\beta \\ & \quad -\frac{\pi}{2} \le \beta \le \frac{\pi}{2} \end{aligned} \tag{5.7}$$

The symmetries in the theory completely specify the couplings of these eigenstates to the chiral fermion states of the Standard Model. Parity-violating terms do exist in the interaction Lagrangian. In fact, one finds that $A_{\rm PV}$ is enhanced by the following factor [29]:

$$1 + 7\left\{\frac{m_Z^2}{m_{Z_\beta}^2}\left(\cos^2\beta + \sqrt{\frac{5}{3}}\sin\beta\cos\beta\right) + \frac{m_Z^2}{m_{Z'_\beta}^2}\left(\sin^2\beta - \sqrt{\frac{5}{3}}\sin\beta\cos\beta\right)\right\} \tag{5.8}$$

By comparing this prediction to the actual measurement of $A_{\rm PV}$, limits can be placed on m_{Z_β} (which is defined to be lighter than $m_{Z'_\beta}$). For example, considering the case where $\beta = 0$, the theory becomes effectively like an SO(10) GUT with a single extra Z_χ boson. The overall agreement between the $A_{\rm PV}$ measurement and the Standard Model prediction,

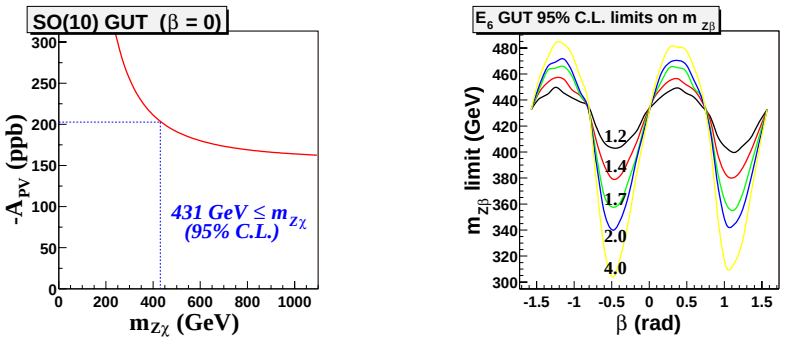

(a) (b)

Figure 5.3: Mass exclusion limits for extra neutral gauge bosons. Figure (a) shows how a single SO(10) Z_χ would modify $A_{\rm PV}$. The fact that the measured value agrees with the Standard Model sets a 95% C.L. limit of $m_{Z_\chi} > 430$ GeV. Figure (b) shows the resulting mass limits for an E_6 model, in which two additional neutral bosons Z_χ and Z_ψ exist. These mix (the mixing is parameterized by the angle β) to form the mass eigenstates Z_β and Z'_β (with $m_{Z_\beta} < m_{Z'_\beta}$). The data curves correspond to various boson mass fractions m_{Z_ψ}/m_{Z_χ}.

however, imposes the constraint $m_{Z_\chi} > 430$ GeV at the 95% confidence level. This is illustrated in Figure 5.3(a). The structure of the more general E_6 theory can be probed by plotting the 95% confidence level limit on m_{Z_β} as a function of the mixing angle β, as shown in Figure 5.3(b). When constructing such a plot, a value for $m_{Z'_\beta}$ must be chosen, subject only to the constraint that $m_{Z_\beta} < m_{Z'_\beta}$. The physically observable mass eigenvalues can be written as:

$$\begin{aligned} m^2_{Z_\beta} &= m^2_{Z_\chi} + \left(m^2_{Z_\psi} - m^2_{Z_\chi}\right)\sin^2\beta \\ m^2_{Z'_\beta} &= m^2_{Z_\chi} + \left(m^2_{Z_\psi} - m^2_{Z_\chi}\right)\cos^2\beta \end{aligned} \tag{5.9}$$

Requiring m_{Z_β} to be lighter than $m_{Z'_\beta}$ is therefore equal to requiring that $m_{Z_\chi} < m_{Z_\psi}$, as long as $|\beta| < \pi/4$. Outside of this range of β, the relationship between m_{Z_χ} and m_{Z_ψ} is simply reversed. The data curves in Figure 5.3(b) therefore all show cyclic behavior with period $\pi/2$, with each data curve corresponding to a different boson mass fraction m_{Z_ψ}/m_{Z_χ} (or m_{Z_χ}/m_{Z_ψ} for $|\beta| > \pi/4$). As was seen before, for $|\beta| \to 0$ the model becomes more SO(10)-like, and a limit of 430 GeV for the mass of Z_χ (which essentially becomes the same as Z_β) is obtained. For $\beta \neq 0$, however, the mass limit on m_{Z_β} becomes very sensitive to the particular choice of input parameters. It is interesting to note that for some input values, two additional neutral gauge bosons are actually more compatible with the data than is a single extra gauge boson (i.e., "simple" SO(10) is less favorable than E_6). For example, for $m_{Z_\psi} \approx 2m_{Z_\chi}$ and $\beta \approx -\pi/6$, a Z_β boson could exist as low as 340 GeV and still not generate effects that would show up in the present measurement of $A_{\rm PV}$. Accommodating

a Z_β boson with a mass much lower than this requires a large discrepancy between the masses of the underlying Z_χ and Z_ψ bosons and fine-tuning the mixing angle β.

As with the X parameter analysis, these mass limits should be compared to those already obtained, either from e^+e^- collider data, direct searches at $p\overline{p}$ colliders, or precision low-energy measurements. Data from the e^+e^- colliders at SLAC and CERN have already ruled out the possibility of additional light neutral bosons (with masses less than m_Z) of almost any sort [18]. Also, LEP-II data places a constraint on the mass of an SO(10) Z_χ at roughly 680 GeV [42]. Direct searches at the Tevatron have yielded limits not quite as high, $m_{Z_\chi} > 595$ GeV [18]. For the LEP-II data, adding a second gauge boson, as in an E_6 model, loosens the constraints somewhat, depending on the particulars of the model. Two specific models are model ψ ($\beta = \pi/2$) and model η ($\tan\beta = -\sqrt{5/3}$). This latter model arises naturally in some superstring theories. The limits from LEP-II are roughly 450 GeV for these two models [42], whereas the Tevatron is able to do better, excluding new bosons up to 600 GeV [18]. An additional SO(10) Z_χ would modify Q_W(Cs) in the following manner [34]:

$$Q_W(^{133}\mathrm{Cs}) = -73.10 \pm 0.03 + 65.4\frac{m_Z^2}{m_{Z_\chi}^2} \tag{5.10}$$

The difference between the atomic parity violation group's measurement of Q_W and the Standard Model prediction is 0.39 ± 0.49, corresponding to an impressive $\sim$750 GeV limit on the mass of the Z_χ, if the theoretical uncertainties in the atomic transition amplitudes are taken at face value.

5.2.3 Lepton Compositeness

Contact interactions provide a general method of characterizing the effects of new high energy physics on low energy observables [41, 113]. Examples of new high energy physics include new interactions (such as the additional neutral gauge bosons discussed above) and fermionic substructure. The possibility of fermionic substructure, in which quarks and leptons, instead of being pointlike as has been assumed until now, are actually bound states of more fundamental constituents, is an intriguing feature of many possible extensions to the Standard Model. The energy at which the constituent binding forces become important is parameterized by the compositeness scale Λ_{ff}, where f denotes the fermion type

Model	$(\eta_{LL},\ \eta_{RR},\ \eta_{LR})$	*LEP-II* Λ^+_{ee} / Λ^-_{ee} (TeV)	***Møller*** Λ^+_{ee} / Λ^-_{ee} (TeV)
LL	$(\pm1,\ 0,\ 0)$	7.1 / 9.0	**7.0 / 7.9**
RR	$(0,\ \pm1,\ 0)$	7.0 / 8.9	**7.9 / 7.0**
VV	$(\pm1,\ \pm1,\ \pm1)$	15.9 / 18.0	— / —
AA	$(\pm1,\ \pm1,\ \mp1)$	11.3 / 11.5	— / —

Table 5.3: Limits on the electron–electron compositeness scale Λ_{ee} for various models. For each model, two one-sided 95% C.L. lower limits are quoted. The first corresponds to the upper sign of the model parameters η, while the second corresponds to the lower sign. The LEP-II limits come from a combined analysis of the ALEPH, DELPHI, L3 and OPAL Bhabha scattering data at $\sqrt{s} \approx 200$ GeV [42].

being considered. For example, well below the electron compositeness scale Λ_{ee}, the Lagrangian describing the new physics can be written as an effective four-electron contact interaction [41]:

$$\mathcal{L}_{ee} = \frac{4\pi}{2\Lambda^2_{ee}} \left[\eta_{LL}(\overline{\psi}_{_L}\gamma_\mu\psi_{_L})^2 + \eta_{RR}(\overline{\psi}_{_R}\gamma_\mu\psi_{_R})^2 + 2\eta_{LR}(\overline{\psi}_{_L}\gamma_\mu\psi_{_L})(\overline{\psi}_{_R}\gamma^\mu\psi_{_R})\right] \tag{5.11}$$

In most typical models, the helicity amplitudes are defined to be $\eta_{ij} = \pm 1$. In particular, this measurement of $A_{\rm PV}$ is sensitive to the left-handed (LL) and the right-handed (RR) models, which define $(\eta_{LL},\ \eta_{RR},\ \eta_{LR})$ to be $(\pm1,\ 0,\ 0)$ and $(0,\ \pm1,\ 0)$, respectively. The contact Lagrangian in Equation (5.11) leads to the following modification of $\sin^2\theta_W$:

$$\sin^2\theta_W(\text{meas.}) - \sin^2\theta_W(\text{SM}) = \frac{\pi}{G_F\sqrt{2}} \frac{(\eta_{LL} - \eta_{RR})}{\Lambda^2_{ee}} \tag{5.12}$$

where $\sin^2\theta_W(\text{SM})$ is the Standard Model prediction for $\sin^2\theta_W$ and $\sin^2\theta_W(\text{meas.})$ is what is actually measured. The limits this measurement sets on Λ_{ee} for the LL and RR models are compared in Table 5.3 to those obtained from the LEP-II e^+e^- collider data. (Note that the other low-energy weak mixing angle measurements considered here, namely the NuTeV experiment and the atomic parity violation experiment, are insensitive to Λ_{ee}. The atomic parity violation experiment, however, is sensitive to Λ_{eq}, which parameterizes contact interactions between quarks and electrons.) For each model, two lower bounds are given, corresponding to the two different cases $\eta_{LL(RR)} = +1$ and $\eta_{LL(RR)} = +1$. This experiment is not sensitive to the vector (VV) or axial-vector (AA) models, since they define $\eta_{RR} = \eta_{LL}$.

This discussion has treated lepton compositeness within the context of an effective four-

Fermi contact interaction. In general, this same type of treatment can extend to any new class of physics involving very high energy scales. For instance, limits on the masses of new neutral gauge bosons can be obtained following this prescription, as long as one properly defines the helicity amplitudes (η_{LL}, η_{RR}, η_{LR}) for the particular model being considered (e.g., SO(10), E_6, etc.) [113]. An example is the brief hint in 2001 of a statistically significant finite compositeness scale for the AA model. Specifically, a combined analysis of LEP-II data found that $7.2 \leq \Lambda^{+}_{ee}(\mathrm{AA}) \leq 15.8$ TeV (95% C.L.) [114]. Among other things, this could have been interpreted as possible evidence for a new gauge boson [115]. However, in 2003, this preliminary result was repudiated by a more complete analysis [42].

Another example of how this formalism can be interpreted in terms of specific models is for the case of doubly-charged Higgs bosons, which abound in certain extensions of the Standard Model, such as left-right supersymmetric models [116, 117]. Assuming $m_{\Delta^{--}}$ is the mass of the doubly-charged Higgs boson and g_{ee} is the $ee\Delta^{--}$ coupling strength, the following ratio can be solved for using the $\Lambda^{\pm}_{ee}$ limits given above [118]:

$$\frac{g^2_{ee}}{m^2_{\Delta^{--}}} \leq \frac{4\pi}{(\Lambda^{\pm}_{ee})^2} \tag{5.13}$$

This leads to the limits $g^2_{ee}/m^2_{\Delta^{--}} \leq 0.017\ G_F$ or $0.022\ G_F$, depending on the helicity structure one assumes. This is roughly a factor of five times better than the current limits for $g^2_{\ell\ell}/m^2_{\Delta^{--}}$, which come from a variety of sources, including studies of muon decay and Bhabha scattering data [118]. One can choose to convert these limits into mass limits, if one assumes a particular value for g_{ee}.

5.3 Future Prospects

The results presented here have summarized the analysis of Runs I and II of SLAC experiment E-158. Since this comprises roughly one-half of the total data set, one may expect a decrease of approximately $1/\sqrt{2}$ in the statistical uncertainty on the $\sin^2\theta_W$ measurement. What effect might this increased statistical resolution have on the final weak mixing angle measurement? This question is explored in Table 5.4, which essentially reiterates the discussion at the end of Chapter 4, but in terms of $\delta(\sin^2\theta_W)$ as well as $\delta(A_{\mathrm{PV}})$. This table also attempts to project the future likely final uncertainty on $\sin^2\theta_W$, based on

Source of uncertainty	*Runs I & II* $\delta(A_{\rm PV})$ (ppb) / $\delta(\sin^2\theta_W)$	*Runs I, II & III* $\delta(A_{\rm PV})$ (ppb) / $\delta(\sin^2\theta_W)$
Statistics	± 21 / ± 0.0015	$\pm\mathbf{14}$ / $\pm\mathbf{0.0010}$
Asymmetry corrections	± 12 / ± 0.0011	$\pm\mathbf{9}$ / $\pm\mathbf{0.0008}$
Background dilutions	± 1 / ± 0.0001	± 1 / ± 0.0001
Normalizations	± 10 / ± 0.0007	$\pm\mathbf{6}$ / $\pm\mathbf{0.0004}$
Analyzing power ($\mathcal{AP}$)	— / ± 0.0001	— / ± 0.0001
Bremsstrahlung ($\mathcal{F}_{brem}$)	— / ± 0.0001	— / ± 0.0001
Electroweak r.c.'s	— / ± 0.0006	— / $\pm\mathbf{0.0004}$
TOTAL	± 26 / ± 0.0021	± 18 / ± 0.0015

Table 5.4: A listing of the contributions to $\delta(A_{\rm PV})$ and $\delta(\sin^2\theta_W)$. The right-most column gives the projected estimates for the final errors, once the Run III data is included. Note that some sources only contribute to $\delta(\sin^2\theta_W)$, not to $\delta(A_{\rm PV})$. The dominant sources of uncertainty are shown in bold. These represent the areas on which similar parity-violating electron scattering experiments should focus in order to drive $\delta(\sin^2\theta_W)$ to the smallest possible value.

the best information available at the present time. One can see that there are four dominant sources of uncertainty: statistics, asymmetry corrections (primarily the beam corrections), normalizations (primarily the beam polarization), and the theoretical uncertainty in the calculation of the electroweak radiative corrections. This latter uncertainty stems from the handling of the low-energy quark loop contributions, which must be interpreted from $e^+e^- \rightarrow$ hadrons data. This uncertainty contributes whenever one tries to compare the $\sin^2\theta_W$ measurement to any Standard Model or "new physics" prediction, and so should be counted as a real uncertainty. It has been noted that, with a more careful analysis, this uncertainty *could* be reduced by perhaps a factor of two [29].

Assuming that the reductions in the uncertainties projected in Table 5.4 actually occur, the discovery reach for the various classes of new physics discussed in the previous sections can be reevaluated. With a total uncertainty of ± 0.0015 on the final measurement of $\sin^2\theta_W$, the X parameter constraint would decrease to $\delta X = \pm 0.2$. In addition, the extra Z' models would be probed at the $400 - 600$ GeV level with 95% confidence, depending on the specifics of each particular theory and how the mean value of $\sin^2\theta_W$ moves. Finally, assuming the mean value of $\sin^2\theta_W$ remains roughly the same, the 95% C.L. limits on Λ^+_{ee}(LL) and Λ^-_{ee}(LL) would increase to 8.2 and 9.6 TeV, respectively.

This experiment builds upon the work of a long line of fixed target parity-violating electron scattering experiments, beginning in 1978 with the pioneering efforts of SLAC experiment E-122, whose results were so crucial in establishing the GWS theory as the "standard

model" of particle physics. That experiment measured the parity-violating asymmetry in inclusive electron–deuterium scattering, which was of order 100 ppm for the experimental kinematics, ultimately achieving a precision of roughly 10 ppm [11]. Since then, technology has improved considerably, so that the current experiment measures an asymmetry of order 100 ppb with a precision of 26 ppb, and a precision of nearly 15 ppb could potentially be within reach. Future parity-violating electron scattering experiments (see below) plan to measure $\mathcal{O}(100)$ ppb asymmetries with precisions of roughly 10 ppb. This experiment has clearly shown the feasibility of making such measurements. However, how far can the basic techniques of such measurements be pushed? Table 5.4 may help to illuminate an answer.

Disregarding the statistical uncertainty in Table 5.4, the only significant uncertainties that remain are those due to asymmetry corrections, normalizations, and electroweak radiative corrections. The likelihood of the reduction in the theoretical uncertainties connected with the computation of the electroweak radiative corrections for this experiment has already been commented upon. In any precision parity violation experiment, one-loop electroweak radiative corrections are essential to compute, and in fact provide one of the primary motivations for conducting such experiments, allowing the Standard Model to be tested at a very deep level. Unfortunately, a general discussion of their significance is hindered by the fact that they are naturally very process-specific. As this experiment has shown, however, even in the theoretically clean leptonic sector the uncertainties associated with such calculations can be substantial, contributing at the few percent level. Hadronic processes will tend to suffer from even larger uncertainties. Clearly, then, such uncertainties must be taken into account if precisions of a few percent are desired.

For any parity-violating electron scattering experiment, beam polarization is likely to contribute signficantly to the overall uncertainty. In the case of the present experiment, the uncertainty reduction projected in Table 5.4 amounts to improving the beam polarization measurement error from 5% to 3%. This should be possible, assuming various refinements in the polarization analysis procedure. In general, beam polarization measurements are easily capable of achieving precisions of 3%, and in fact even better precisions are possible [119]. Thus, while beam polarization should be a primary concern for any future parity violation experiment, it does not present an insurmountable obstacle.

The last issue relates to the asymmetry corrections. The most troublesome asymmetry corrections for the present experiment are those due to higher-order beam effects, such as

intra-pulse beam fluctuations. As the preliminary analysis of the Run III data has shown, splitting the BPM signals into time slices provides valuable information, and should definitely be considered for any future, high-precision parity violation experiment. Backgrounds must be considered as they relate to each specific experiment. For the present experiment, the most significant background is presented by the ep flux. Correcting for the sizeable asymmetry it induces in the Møller detector, however, can be accomplished such that the uncertainty introduced to $A_{\rm PV}$ is only 5 ppb. Each experiment must be designed with the relevant backgrounds in mind, eliminating or greatly reducing them whenever possible, and also providing the necessary means (for example, with insertable/removable components) for sensitive background measurements.

As an example of a future high-precision parity-violating electron scattering experiment, the Q_{weak} collaboration at the Jefferson National Accelerator Facility (JLab) plans to determine the weak charge of the proton (which is directly proportional to $\sin^2\theta_W$) by measuring the parity-violating asymmetry in elastic ep scattering at a kinematic point of $Q^2 = 0.03$ GeV^2 [120]. The experiment will use a 6 GeV continuous electron beam ($\sim$180 μA) with $\sim$80% polarization, scattering it off a 35 cm LH_2 target. The parity-violating ep asymmetry $A_{\rm PV}^{ep}$ can be related to $\sin^2\theta_W$ in the following way [120]:

$$A_{\rm PV}^{ep} = \frac{-G_F Q^2}{4\pi\alpha\sqrt{2}}\left\{1 - 4\sin^2\theta_W + Q^2 B(Q^2)\right\} \tag{5.14}$$

where $B(Q^2)$ is a nucleon structure function defined in terms of neutron and proton electromagnetic and weak form factors. While these quantities have been measured very precisely by previous JLab experiments [121, 122, 123], nucleon structure is expected to be the largest source of theoretical uncertainty, contributing to $\delta(\sin^2\theta_W)$ at the ± 0.0002 level. In all, the experiment plans to measure $A_{\rm PV}^{ep} \approx -280$ ppb to a precision of nearly 10 ppb, corresponding to a low-energy measurement of $\sin^2\theta_W$ with a total uncertainty of roughly ± 0.0007. This very precise, off-shell measurement of the weak mixing angle will search for many of the same types of new physics as SLAC experiment E-158, such as Z' bosons and new heavy particles in vacuum polarization loops. However, since it employs a semi-leptonic process, in general it will not be sensitive to the exact same types of effective contact interactions. The compositeness scale Λ_{eq} will be probed at above the TeV level.

A potential 12 GeV upgrade to the JLab beam will make possible two other exciting

opportunities to test the Standard Model with increasing precision [124]. First, a repeat of the present parity-violating Møller scattering experiment, but with subtle differences in the kinematics, is being considered. The upgraded 12 GeV, ∼80% longitudinally polarized, 100 μA continuous electron beam will be scattered off a 75 cm LH_2 target. A spectrometer and collimator system will be used to isolate the Møller scatters from the other backgrounds. At 12 GeV, and at the optimal kinematic operating point of $Q^2 \approx 0.008$ GeV^2, the parity-violating Møller asymmetry $A_{\rm PV}$ will be four times smaller, roughly 40 ppb, and the measurement will necessarily be harder and take longer. However, with essentially one full year of running, it is calculated that $A_{\rm PV}$ could be measured with a precision approaching 1 ppb, corresponding to an uncertainty in $\sin^2\theta_W$ of ± 0.0004. Since this proposed experiment would observe exactly the same physical process as the present experiment, it would probe the same types of new physics as discussed in this thesis, though with increased sensitivity. Electroweak-scale oblique corrections would be probed at four times the present level, corresponding to $\delta X = \pm 0.05$. Extra Z_χ bosons in the SO(10) model, for instance, would be sought up to nearly 1 TeV. In fact, a Z_χ boson with $m_{Z_\chi} \approx 800$ GeV, which is well outside the current limits set by LEP, would enhance $A_{\rm PV}$ by nearly 10%. A 1 ppb asymmetry measurement would therefore be able to detect such a particle at the impressive 3.6σ level. In addition, electron compositeness would be probed at the 17 TeV level with 95% confidence.

A second proposed experiment for the upgraded 12 GeV JLab facility is called the DIS-parity experiment [124]. It would measure the parity-violating asymmetry in electron–deuterium deep inelastic scattering, much like the first neutral current asymmetry measurement, SLAC experiment E-122. However, the upgraded JLab beam would be able to provide enough luminosity to allow for a 0.6% statistical measurement of the $\mathcal{O}(100)$ ppm asymmetry. Systematic uncertainties would actually dominate the measurement, resulting in a final precision of roughly 1.3% (compare this to the ∼10% measurement of 26 years ago). This would translate into a ± 0.0016 measurement of $\sin^2\theta_W$. One of the interesting aspects of the DIS-parity experiment is that its average Q^2 of 2.9 GeV^2 is very nearly equal to that of the NuTeV experiment. Thus, it offers the possibility of being able to resolve the 3σ deviation presently reported by that collaboration.

On the end of the spectrum opposite to the high-precision parity-violating electron scattering experiments described above, there are the collider programs currently in operation

or else under construction. For example, the Tevatron continues to explore the high energy frontier in its search for the Higgs particle. In several years, perhaps as soon as 2007, the Large Hadron Collider (LHC) at CERN is scheduled to begin taking data. The LHC consists of a ring 8.6 km in diameter. Inside this ring, two counter-rotating 7 TeV proton beams will collide with each other at four interaction points. The nearly 17-mile-long tunnel will house thousands of two-in-one superconducting magnets, able to produce the ± 8 T fields needed to steer the two proton beams in opposite directions. Many of the physics goals motivating the construction of the LHC relate to the dynamics of electroweak symmetry-breaking. Is there indeed a Higgs particle? If so, what is its mass? Does supersymmetry play a role? The LHC will have a profound impact on the entire field of particle physics, if not by what it finds, then by precisely what it does not find.

5.4 Summary

From the earliest days of the Standard Model, parity-violating scattering asymmetries have played a crucial role in determining the precise way in which the elementary particles interact. The experiment described in this thesis has measured the parity-violating cross-section asymmetry $A_{\rm PV}$ in Møller scattering, obtaining the following result:

$$A_{\rm PV} = -160 \pm 21\ ({\rm stat}) \pm 16\ ({\rm syst})\ {\rm ppb} \tag{5.15}$$

This is for the the kinematics $Q^2 = 0.026\ {\rm GeV}^2$. This result provides the first-ever confirmation of parity-violation in electron–electron interactions at the $6.3\,\sigma$ level. The measurement can be interpreted in terms of the weak mixing angle in the $\overline{\rm MS}$ renormalization scheme, yielding the following results:

$$\sin^2\theta_W(Q^2 = 0.026\ {\rm GeV}^2) = 0.2381 \pm 0.0015\ ({\rm stat}) \pm 0.0014\ ({\rm syst}) \tag{5.16}$$

$$\sin^2\theta_W(m_Z) = 0.2308 \pm 0.0015\ ({\rm stat}) \pm 0.0015\ ({\rm syst}) \tag{5.17}$$

These measurements have significant implications for several classes of new physics, including new gauge bosons, new heavy fermions, and electron compositeness. In the SO(10) Z_χ model, for instance, no evidence for new bosons is found up to approximately 430 GeV at the 95% confidence level. The constraint on the parameter X describing electroweak-scale

oblique corrections, including the possibility of new fermions of $\mathcal{O}(100)$ GeV mass, is improved by almost a factor of two. In addition, the lepton compositeness parameters Λ_{ee}^{+} and Λ_{ee}^{-} are probed at the 7 TeV and 8 TeV level, respectively. These limits are found to be comparable to those provided by the high-energy colliders, and will only grow more interesting as both the statistical and systematic uncertainties in Equations (5.15) continue to be reduced by further analysis of additional data.

The techniques used by polarized electron scattering experiments have advanced considerably since their advent in the 1970's, allowing for the measurement of exceedingly small asymmetries [125]. This experiment has clearly demonstrated the feasibility of measuring $\mathcal{O}(100)$ ppb asymmetries with precisions approaching 10 ppb. A future generation of experiments proposes to push the technologies even further by working to achieve $\mathcal{O}(1)$ ppb uncertainties on parity-violating asymmetries. Such measurements provide rigorous tests of the Standard Model at the quantum-loop level and are an essential part of the larger physics program dedicated towards a fundamental understanding of the inner workings of nature.

Bibliography

[1] J. Goldstone, Nuovo Cimento **19**, 154–164 (1961).

[2] J. Goldstone, A. Salam and S. Weinberg, Phys. Rev. **127**, 965–970 (1962).

[3] P. W. Higgs, Phys. Rev. Lett. **13**, 508–509 (1964).

[4] S. Weinberg, Phys. Rev. Lett. **19**, 1264–1266 (1967).

[5] S. L. Glashow, Nucl. Phy. **22**, 579–588 (1961).

[6] A. Salam and J. C. Ward, Phys. Lett. **13**, 168–171 (1964).

[7] S. A. Bludman, Nuovo Cimento **9**, 433 (1958).

[8] F. J. Hasert *et al.* (Gargamelle Neutrino Collaboration), Phys. Lett. **B46**, 138–140 (1973).

[9] A. De Rújula, H. Georgi and S. L. Glashow, Annals Phys. **109**, 258 (1977).

[10] C. Y. Prescott *et al.*, Phys. Lett. **B77**, 347–352 (1978). SLAC-PUB-2148.

[11] C. Y. Prescott *et al.*, Phys. Lett. **B84**, 524 (1979). SLAC-PUB-2319.

[12] R. N. Cahn and F. J. Gilman, Phys. Rev. **D17**, 1313–1322 (1978).

[13] E. Fermi, Nuovo Cimento **11**, 1–19 (1934).

[14] G. Arnison *et al.*, Phys. Lett. **B122**, 103 (1983).

[15] G. Arnison *et al.*, Phys. Lett. **B126**, 398 (1983).

[16] B. Pietrzyk (for the LEP collaborations and the LEP Electroweak Working Group), "LEP asymmetries and fits to the Standard Model," hep-ex/9406001 (1994).

[17] F. Abe *et al.* (CDF Collaboration), Phys. Rev. Lett. **74**, 2626–2631 (1995).

[18] Particle Data Group, K. Hagiwara *et al.*, Phys. Rev. **D66**, 010001 (2002).

[19] A. Heister *et al.* (ALEPH Collaboration), Phys. Lett. **B526**, 191–205 (2002). hep-ex/0201014.

[20] P. P. Janot, "Higgs searches at LEP and at the Tevatron," 21st Physics in Collisions Conference (PIC 01), Seoul, Korea, 28–30 Jun 2001. hep-ex/0110076.

[21] M. Spira, A. Djouadi, D. Graudenz and P. M. Zerwas, Nucl. Phys. **B453**, 17–82 (1995). hep-ph/9504378.

[22] W. J. Marciano, Ann. Rev. Nucl. Part. Sci. **41**, 469–509 (1991).

[23] H. Georgi, H. R. Quinn and S. Weinberg, Phys. Rev. Lett **33**, 451–454 (1974).

[24] H. Georgi and S. L. Glashow, Phys. Rev. Lett. **32**, 438–441 (1974).

[25] S. Weinberg, Phys. Rev. **D19**, 1277–1280 (1979).

[26] S. Dimopoulos, S. Raby and F. Wilczek, Phys. Rev. **D24**, 1681–1683 (1981).

[27] E. Derman and W. J. Marciano, Ann. Phys. **121**, 147–207 (1979).

[28] C. Møller, Ann. Phys. (Leipzig) **14**, 531 (1933).

[29] A. Czarnecki and W. J. Marciano, Phys. Rev **D53**, 1066–1072 (1996).

[30] A. Sirlin, Phys. Rev. **D22**, 971–981 (1980).

[31] P. Gambino and A. Sirlin, Phys. Rev. **D49**, R1160–R1162 (1994).

[32] W. J. Marciano, Annu. Rev. Nucl. Part. Sci. **41**, 469 (1991).

[33] W. J. Marciano and A. Sirlin, Phys. Rev. **D27**, 552–556 (1983).

[34] W. J. Marciano, "Spin and precision electroweak physics," in *Spin Structure in High Energy Processes*, Proceedings of the 21st SLAC Summer Institute, Stanford, California, 1993, edited by L. De Porcel and C. Dunwoodie (SLAC Report No. 444, Stanford, 1994).

[35] M. Peskin and T. Takeuchi, Phys. Rev. Lett. **65**, 964–967 (1990).

[36] J. L. Hewett, T. Takeuchi and S. Thomas, "Indirect probes of new physics," in *Electroweak Symmetry Breaking and New Physics at the TeV Scale*, ed. T. Barklow *et al.*, pp. 548–649, World Scientific (1996). hep-ph/9603391

[37] I. Maksymyk, C. P. Burgess and D. London, Phys. Rev. **D50**, 529–535 (1994).

[38] J. Rosner, Comm. Nucl. Part. Phys **14**, 229 (1985).

[39] B. Holdom, Phys. Lett. **B166**, 196–205 (1986).

[40] S. Godfrey, Phys. Rev. **D51**, 1402–1405 (1995).

[41] E. J. Eichten, K. D. Lane and M. E. Peskin, Phys. Rev. Lett. **50**, 811–814 (1983).

[42] ALEPH, DELPHI, L3, OPAL, and SLD collaborations, "A combination of preliminary electroweak measurements and constraints on the Standard Model," hep-ex/0312023 (2003).

[43] A. Czarnecki and W. J. Marciano, Int. J. Mod. Phys. **A15**, 2365–2376 (2000). hep-ph/0003049.

[44] G. P. Zeller *et al.*, Phys. Rev. Lett. **88**, 091802 (2002).

[45] G. P. Zeller *et al.*, Phys. Rev. Lett. **90**, 239902 (2003).

[46] C. S. Wood *et al.*, "Measurement of parity nonconservation and an anapole moment in cesium", *Science* **275**, 1759–1763 (1997).

[47] E. A. Paschos and L. Wolfenstein, Phys. Rev. **D7**, 91–95 (1973).

[48] D. Bardin and V. A. Dokuchaeva, Report No. JINR-E2-86-260 (1986).

[49] K.-P. O. Diener, S. Dittmaier and W. Hollik, "Electroweak radiative corrections in deep–inelastic neutrino scattering — implications for NuTeV?", hep-ph/0310364 (2003).

[50] S. Davidson, S. Forte, P. Gambino, N. Rius and A. Strumia, "Old and new physics interpretations of the NuTeV anomaly", JHEP **02**, 037 (2002). hep-ph/0112302.

[51] W. C. Haxton and C. E. Wieman, Ann. Rev. Nucl. Part. Sci. **51**, 261–293 (2001). nucl-th/0104026.

[52] S. C. Bennett and C. E. Wieman, Phys. Rev. Lett. **82**, 2484–2487 (1999).

[53] A. Derevianko, Phys. Rev. Lett. **85**, 1618–1621 (2000).

[54] A. Derevianko, Phys. Rev. **A65**, 012106 (2001).

[55] V. A. Dzuba, V. V. Flambaum and J. S. M. Ginges, Phys. Rev. **D66**, 076013 (2002).

[56] O. P. Sushkov, Phys. Rev. **A63**, 042504 (2001).

[57] M. Yu. Kuchiev and V. V. Flambaum, Phys. Rev. Lett. **89**, 283002 (2002).

[58] A. Derevianko, B. Ravaine and W. R. Johnson, "Relaxation effect and radiative corrections in many-electron atoms," physics/0401043 (2004).

[59] Ya. B. Zel'dovich, Sov. Phys. JETP **6**, 1184 (1958).

[60] P. Mastromarino *et al.*, IEEE Trans. Nucl. Sci. **49**, Issue 3, 1097–1105 (2002). SLAC-PUB-9071.

[61] K. S. Kumar *et al.*, "A precision measurement of the weak mixing angle in Møller scattering," SLAC-Proposal-E-158 (1997). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[62] D. T. Pierce *et al.*, Phys. Lett. **51A**, 465 (1975).

[63] D. T. Pierce *et al.*, Appl. Phys. Lett. **26**, 670 (1975).

[64] C. K. Sinclair *et al.*, AIP Conf. Proc. **35**, 424 (1976). SLAC-PUB-2109.

[65] E. L. Garwin, D.T. Pierce and H.C. Siegmann, "Polarized photoelectrons from optically magnetized semiconductors," Swiss Physical Society Meeting, Bern, Apr 26, 1974. SLAC-PUB-1576.

[66] T. Maruyama *et al.*, Nucl. Instrum. Meth. **A492**, 199 (2002). SLAC-PUB-9133.

[67] D. A. Luh *et al.*, AIP Conf. Proc. **675**, 1029 (2003). SLAC-PUB-9603.

[68] G. Mulhollan *et al.*, SPIN 96, 706 (1996). SLAC-PUB-7346.

[69] R. Mair *et al.*, Phy. Lett. **A212**, 231 (1996). SLAC-PUB-6836.

[70] J. E. Clendenin *et al.*, 7th International Workshop on Polarized Gas Targets and Polarized Beams, Urbana, IL, 18–22 Aug 1997, unpublished. SLAC-PUB-7619.

[71] M. Woods *et al.*, J. Appl. Phys. **73**, 8531 (1993). SLAC-PUB-5894.

[72] T. Maruyama *et al.*, AIP Conf. Proc. **570**, 976 (2001). SLAC-REPRINT-1997-085.

[73] R. Alley *et al.*, Nucl. Instrum. Meth. A **365**, 1 (1995). SLAC-PUB-6489.

[74] A. Brachmann *et al.*, Proc. SPIE Int. Soc. Opt. Eng. **4632**, 211 (2002). SLAC-PUB-9145.

[75] T. B. Humensky *et al.*, Nucl. Instrum. Meth. A **521**, Issues 2–3, 261–298 (2004). SLAC-PUB-9381.

[76] E. Hecht, *Optics*. Pearson Addison Wesley, Reading, Massachusetts (2001).

[77] G. D. Cates *et al.*, Nucl. Instrum. Meth. A **278**, 293 (1989).

[78] K. S. Kumar. *Parity Violation in Elastic Electron Carbon Scattering*. PhD thesis, Syracuse University, unpublished (1990).

[79] P. Horowitz and W. Hill, *The Art of Electronics*. Cambridge University Press, Cambridge (1980). The bit shift algorithm described on pp. 437 – 442 is used to generate the pseudo-random helicity sequence.

[80] H. Wiedemann, *Particle Accelerator Physics I*. Springer-Verlag, Berlin (1998).

[81] R. B. Neal, *The Stanford Two-Mile Accelerator*. W. A. Benjamin, New York, New York (1968).

[82] F. J. Decker *et al.*, *Particle Accelerator, Vol. 2*, New York, 774–776 (1999). SLAC-PUB-8113.

[83] D. H. Whittum and Yu. Kolomensky, Rev. Sci. Instrum. **70**, 2300-2313 (1999). SLAC-PUB-7846.

[84] M. S. Griffo, "E-158 Technical Note 04: Design for Electron Beam Energy Monitors Using Synchrotron Light," (2000). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[85] M. Woods, "E-158 Technical Note 39: Synchrotron Background Analysis For Run I," (2003). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[86] J. Gao *et al.*, Nucl. Instrum. Meth. A **498**, 90 (2003). SLAC-PUB-9565.

[87] J. G. Weisend, II *et al.*, Adv. Cryo. Engr. **47A**, 171 (2002). SLAC-PUB-8917.

[88] J. G. Weisend, II *et al.*, 19th International Cryogenic Engineering Conference (ICEC 19), Grenoble, France, 22-26 Jul 2002, unpublished. SLAC-PUB-9294.

[89] M. Woods, "E-158 Technical Note 40: Spotsize Analysis For Run I," (2003). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[90] I. Younus, "E-158 Technical Note 47: E158 Calorimeter," (2003). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[91] P. Gorodetzky *et al.*, Nucl. Instrum. Meth. A **361**, 161 (1995). CERN-PPE-94-226.

[92] L. Keller, "E-158 Technical Note 28: Lumi Asymmetry and Acceptance as a Function of Transverse and Angular Displacement at the Target," (2001). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[93] G. M. Jones *et al.*, 10th International Conference on Calorimetry in High Energy Physics (CALOR 2002), Pasadena, California, 25-30 Mar 2002. SLAC-PUB-10173.

[94] D. Lhuillier, "E-158 Technical Note 42: Pion Detector Analysis," (2003). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[95] V. B. Berestetskii, E. M. Lifshitz and L. P. Pitaevskii, *Quantum Electrodynamics, Vol. 4.* Pergamon, Oxford (1982).

[96] G. Haller *et al.*, "Integrating ADC VME Module for E158," (2000). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[97] M. S. Cooke, "E-158 Technical Note 48: Probing Moller Detector Linearity with the Polarized Iron Foil and Light Filters," (2003). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[98] L. J. Kaufman and K. S. Kumar, "E-158 Technical Note 44: Neutral Backgrounds in the Moller Detector," (2003). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[99] P. Anthony and Z. Szalata, "A Flexible High Performance UME Based Data Acquisition System," unpublished (1996). SLAC-PUB-7201.

[100] M. Woods *et al.*, "E-158 Technical Note 45: Run I Systematic Error for Beam Asymmetry Corrections," (2003). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[101] J. Farrell, "E-158 Technical Note 58: Characterization of E158 Beam," (2003). SLAC-TN-03-010. Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[102] G. A. Loew *et al.*, 7th Int. Conf. on High Energy Accelerators, Yerevan, USSR, 28 Aug–2 Sep, 1969. SLAC-PUB-0659.

[103] J. Seeman, "Linear Collider Accelerator Physics Issues Regarding Alignment," *Stanford 1989, Proceedings, Accelerator Alignment*, pp. 257–262 (1989).

[104] E. Chudakov, "E-158 Technical Note 43: Polarimetry, Run I," (2003). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[105] Yu. Kolomensky, "E-158 Technical Note 56: Higher Order Asymmetry Systematics," (2003). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[106] P. Bosted, "E-158 Technical Note 23: Radiative Corrections for E158," (2002). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[107] B. Tweedie, "E-158 Technical Note 37: The GEANT3 Simulation of E-158," (2003). Available at: http://www.slac.stanford.edu/exp/e158/doc.html.

[108] P. Souder and K. Kumar, "E-158 Technical Note 30: Systematic Corrections for E-158," (2003). Available at:
http://www.slac.stanford.edu/exp/e158/doc.html.

[109] C. Arroyo, "E-158 Technical Note 33: ep Inelastic Scattering Asymmetry Correction," (2003). Available at:
http://www.slac.stanford.edu/exp/e158/doc.html.

[110] L. Dixon and M. Schreiber, "Radiative corrections to the azimuthal asymmetry in transversely polarized Møller scattering," (2004). SLAC-PUB-10345, hep-ph/0402221.

[111] V. A. Zykunov, Yad. Phys. **66**, annot. (2003).

[112] D. London and J. L. Rosner, Phys. Rev. **D34**, 1530–1546 (1986).

[113] M. J. Ramsey-Musolf, Phys. Rev. **C60**, 015501 (1999).

[114] D. Bourilkov, Phys. Rev. D **64**, 071701 (2001).

[115] V. I. Demchik, A. V. Gulov, V. V. Skalozub and A. Yu. Tischenko, "Signals of the Abelian Z' boson within the analysis of the LEP2 data," (2003). hep-ph/0302211.

[116] J. F. Gunion, J. Grifols, A. Mendez, B. Kayser and F. Olness, Phys. Rev. **D40**, 1546–1561 (1989).

[117] B. Dutta and R. N. Mohapatra, Phys. Rev. D **59**, 015018 (1999).

[118] M. L. Swartz, Phys. Rev. **D40**, 1521–1528 (1989).

[119] F. Marie, E. Burtin, C. Cavata, S. Escoffier, D. Lhuillier *et al.*, AIP Conf. Proc. **675**, 846–851 (2003). Prepared for 15th International Spin Physics Symposium (SPIN 2002), Long Island, New York, 9–14 Sep 2002.

[120] D. Armstrong *et al.*, "The Q_{weak} Experiment: A Search for New Physics at the TeV Scale via a Measurement of the Proton's Weak Charge," JLab-Proposal-PR02-020 (2001).

[121] K. A. Aniol *et al.* (HAPPEX Collaboration), Phys. Rev. Lett. **82**, 1096–1100 (1999). nucl-ex/9810012.

[122] K. A. Aniol *et al.* (HAPPEX Collaboration), submitted to Phys. Rev. **C** (2004). nucl-ex/0402004.

[123] L. Lee (for the G0 Collaboration), "The G0 experiment at Jefferson Lab," AIP Conf. Proc. **675**, 272–278 (2003). Prepared for 15th International Spin Physics Symposium (SPIN 2002), Long Island, New York, 9–14 Sep 2002.

[124] Thomas Jefferson National Accelerator Facility, "Pre-Conceptual Design Report (pCDR) for the Science and Experimental Equipment for the 12 GeV Upgrade of CEBAF," Draft 12.1 (2003).

[125] K. S. Kumar and P. A. Souder, Prog. Part. Nucl. Phys. **45**, S333–S395 (2000).